

z/OS



DFSMS: Using DFSMSdfp in the z/OS V1R5 Environment

z/OS



DFSMS: Using DFSMSdfp in the z/OS V1R5 Environment

Note

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 43.

First Edition, March 2004

This edition applies to Version 1 Release 5 of z/OS[®] (5694-A01), Version 1 Release 5 of z/OS.e (5655-G52), and to all subsequent releases and modifications until otherwise indicated in new editions.

Order publications through your IBM[®] representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

IBM welcomes your comments. A form for readers' comments may be provided at the back of this publication, or you may address your comments to the following address:

International Business Machines Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY 12601-5400
United States of America

FAX (United States & Canada): 1+845+432-9405

FAX (Other Countries):

Your International Access Code +1+845+432-9405

IBMLink[™] (United States customers only): IBMUSM10(MHVRCFS)

Internet e-mail: mhvrcfs@us.ibm.com

World Wide Web: <http://www.ibm.com/servers/eserver/zseries/zos/webqs.html>

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:

- Title and order number of this document
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About This Document. v

Required product knowledge.	v
Referenced documents	v
Accessing z/OS DFSMS documents on the Internet	vi
Using LookAt to look up message explanations . . .	vi
Accessing z/OS licensed documents on the Internet	vi

Chapter 1. Introduction 1

How can z/OS DFSMSdfp V1R5 help me?	1
What do I need to do to use the z/OS DFSMSdfp V1R5 enhancements?	2

Chapter 2. Using the BPAM enhancements 5

Application programming	5
Setting up BPAM access to UNIX files	5
Verifying and setting search and read authority.	6
Increasing the concatenation limit for existing programs	7

Chapter 3. Using the catalog enhancements 9

Administering	10
Enabling JOBCAT and STEPCAT DD statements	10
Disabling JOBCAT and STEPCAT DD statements	10
Specifying automatic autoadd enablement	11
Setting the catalog usage threshold	12
Diagnosing	12
Ending a service task abnormally	12
Disabling and enabling the SYMREC function.	13

Chapter 4. Using the enhanced data integrity function for physical sequential data sets 15

Administering	15
Enabling the enhanced data integrity function.	16
Modifying the enhanced data integrity function	17
Disabling the enhanced data integrity function	18
Application programming	18
Excluding your applications from enhanced data integrity processing.	18

Chapter 5. Using up to 65 535 file sequence numbers 21

Administering	22
Creating new data sets with file sequence numbers greater than 9999	22

Processing existing data sets using file sequence numbers greater than 9999	23
Application programming	24
Modifying applications that use IECDUCBT	24
Modifying applications that use IEFJFCBN	24

Chapter 6. Using the name-hiding enhancement 25

Administering	26
-------------------------	----

Chapter 7. Using the re-indexing of online volumes enhancement 27

Customizing	27
Operating	28

Chapter 8. Using the Storage Management Subsystem (SMS) enhancements 29

Administering	31
Disabling GDS reclaim processing.	31
Enabling GDS reclaim processing	33
Setting the multi-tiered storage group option	33

Chapter 9. Using the sequential data set striping enhancement. 37

Administering	37
-------------------------	----

Chapter 10. Using the VSAM message enhancements 39

Application Programming	39
Implementing the message improvements	39

Appendix. Accessibility 41

Using assistive technologies	41
Keyboard navigation of the user interface	41

Notices 43

Programming interface information	44
Trademarks	44

Glossary 45

Index 49

About This Document

This document helps you to understand the z/OS V1R5 DFSMSdfp enhancements. It includes only information about the z/OS V1R5 DFSMSdfp enhancements; appropriate cross references are provided throughout the document if you need more background information about a specific concept or task.

If you are an experienced system programmer, application programmer, storage administrator, or member of the information technology team responsible for understanding and using the DFSMSdfp enhancements, this document provides you with enough information to understand and apply the DFSMSdfp enhancements at your site. This information is also helpful for anyone who wants a quick explanation of the tasks associated with these enhancements.

For information about accessibility features of z/OS, for users who have a physical disability, please see "Accessibility," on page 41.

Required product knowledge

To use this book effectively, you should have an in-depth knowledge of your current DFSMSdfp™ installation, including the programming, configuration, and procedures used at your site. Use this document in combination with *z/OS DFSMS Migration*, which covers each enhancement as well as the evaluation, planning, and migration tasks required for each enhancement.

Referenced documents

The following publication has additional detailed information about all aspects of z/OS® V1R5 DFSMS:

Publication Title	Order Number
<i>z/OS DFSMS Migration</i>	GC26-7398

The following publications are referenced in this document:

Title	Order Number
<i>z/OS DFSMS Migration</i>	GC26-7398
<i>z/OS DFSMSdfp Storage Administration Reference</i>	SC26-7402
<i>z/OS DFSMS: Using Data Sets</i>	SC26-7410
<i>z/OS DFSMS Access Method Services for Catalogs</i>	SC26-7394
<i>z/OS DFSMS: Managing Catalogs</i>	SC26-7409
<i>z/OS Security Server RACF General User's Guide</i>	SA22-7685
<i>z/OS Security Server RACF Security Administrator's Guide</i>	SA22-7683
<i>z/OS DFSMS Installation Exits</i>	SC26-7396
<i>z/OS DFSMS: Using the Interactive Storage Management Facility</i>	SC26-7411
<i>z/OS DFSMS: Implementing System-Managed Storage</i>	SC26-7407
<i>z/OS MVS JCL User's Guide</i>	SA22-7598
<i>z/OS UNIX System Services User's Guide</i>	SA22-7801

Title	Order Number
<i>z/OS MVS Initialization and Tuning Reference</i>	SA22-7592
<i>z/OS DFSMS: Using Magnetic Tapes</i>	SC26-7412
<i>z/OS MVS Data Areas, Vol 2 (DCCB-ITZYRETC)</i>	GA22-7582
<i>Device Support Facilities User's Guide and Reference</i>	GC35-0033
<i>z/OS Planning for Multilevel Security</i>	GA22-7509

Accessing z/OS DFSMS documents on the Internet

In addition to making softcopy documents available on CD-ROM, IBM provides access to unlicensed z/OS softcopy documents on the Internet. To view, search, and print z/OS documents, go to the z/OS Internet Library:
<http://www.ibm.com/eserver/zseries/zos/bkserv/>

Using LookAt to look up message explanations

LookAt is an online facility that lets you look up explanations for most of the IBM® messages you encounter, as well as for some system abends and codes. Using LookAt to find information is faster than a conventional search because in most cases LookAt goes directly to the message explanation.

You can use LookAt from the following locations to find IBM message explanations for z/OS elements and features, z/VM®, and VSE:

- The Internet. You can access IBM message explanations directly from the LookAt Web site at <http://www.ibm.com/eserver/zseries/zos/bkserv/lookat/>.
- Your z/OS TSO/E host system. You can install code on your z/OS or z/OS.e systems to access IBM message explanations, using LookAt from a TSO/E command line (for example, TSO/E prompt, ISPF, or z/OS UNIX® System Services running OMVS).
- Your Windows® workstation. You can install code to access IBM message explanations on the *z/OS Collection (SK3T-4269)*, using LookAt from a Windows DOS command line.
- Your wireless handheld device. You can use the LookAt Mobile Edition with a handheld device that has wireless access and an Internet browser (for example, Internet Explorer for Pocket PCs, Blazer, or Eudora for Palm OS, or Opera for Linux handheld devices). Link to the LookAt Mobile Edition from the LookAt Web site.

You can obtain code to install LookAt on your host system or Windows workstation from a disk on your *z/OS Collection (SK3T-4269)*, or from the LookAt Web site (click **Download**, and select the platform, release, collection, and location that suit your needs). More information is available in the LOOKAT.ME files available during the download process.

Accessing z/OS licensed documents on the Internet

z/OS licensed documentation is available on the Internet in PDF format at the IBM Resource Link™ Web site at:
<http://www.ibm.com/servers/resourceLink>

Licensed documents are available only to customers with a z/OS license. Access to these documents requires an IBM Resource Link user ID and password, and a key code. With your z/OS order you received a Memo to Licensees, (GI10-0671), that includes this key code. ¹

To obtain your IBM Resource Link user ID and password, log on to:

<http://www.ibm.com/servers/resourceLink>

To register for access to the z/OS licensed documents:

1. Sign in to Resource Link using your Resource Link user ID and password.
2. Select **User Profiles** located on the left-hand navigation bar.

Note: You cannot access the z/OS licensed documents unless you have registered for access to them and received an e-mail confirmation informing you that your request has been processed.

Printed licensed documents are not available from IBM.

You can use the PDF format on either **z/OS Licensed Product Library CD-ROM** or IBM Resource Link to print licensed documents.

1. z/OS.e customers received a Memo to Licensees, (GI10-0684) that includes this key code.

Chapter 1. Introduction

If you are the person responsible for implementing the z/OS DFSMSdfp V1R5 enhancements at your site, this document is written for you. It uses a "cookbook" approach to using these enhancements. First, you will find out what is included in each enhancement — the "recipe". Then, you will find out how to implement these enhancements — the step-by-step procedures or "directions" for creating this "recipe".

This chapter includes information about the z/OS DFSMSdfp V1R5 enhancements and how they are presented in this document. It also includes a description of the tasks you need to perform to take advantage of these enhancements.

How can z/OS DFSMSdfp V1R5 help me?

The z/OS DFSMSdfp V1R5 enhancements are intended to help you use DFSMSdfp more effectively and efficiently at your site. Each chapter contains a description of one of the following enhancements for z/OS DFSMSdfp V1R5:

- **Basic partitioned access method (BPAM) enhancements** BPAM enhancements include read-only access of UNIX files and an increase in the BPAM concatenation limit. See Chapter 2, "Using the BPAM enhancements," on page 5 for more information.
- **Catalog enhancements** Catalog enhancements allow you to perform the following tasks:
 - Disable or enable JOBCAT and STEPCAT DD statements
 - End a service task even if the task is in recall
 - Receive additional information in SYMREC records
 - Disable and enable the creation of SYMREC records
 - Experience improved catalog performance on your system
 - Enable the autoadd option during initial program load (IPL)
 - Set a usage extent threshold for a catalog

See Chapter 3, "Using the catalog enhancements," on page 9 for more information.

- **Data integrity enhancements** Data integrity is enhanced so that multiple users no longer have concurrent output or update access to a shared physical sequential (PS) data set on direct access storage device (DASD), unless concurrent access is allowed by the application. See Chapter 4, "Using the enhanced data integrity function for physical sequential data sets," on page 15 for more information.
- **File sequence number enhancements** The maximum for file sequence number for tapes and bypass label processing is increased from 9999 to 65535. See Chapter 5, "Using up to 65 535 file sequence numbers," on page 21 for more information.
- **Resource Access Control Facility (RACF®) name-hiding enhancements** RACF name-hiding allows you to restrict user access to the names of data sets. When name-hiding is in effect, only users who have read access to the data sets can retrieve the names of data sets. See Chapter 6, "Using the name-hiding enhancement," on page 25 for more information.

- **Re-indexing of online volume enhancements** Re-indexing of online volumes allows you to re-build a volume index at any time without having to take the volume offline. See Chapter 7, “Using the re-indexing of online volumes enhancement,” on page 27 for more information.
- **Storage Management Subsystem (SMS) enhancements** SMS enhancements allow you to:
 - Set generation data set (GDS) reclaim processing on or off as needed
 - Save temporary changes made to the active configuration data set by saving the active control data set (ACDS) as the source control data set (SCDS)
 - Automatically receive a high-allocation-threshold-exceeded message for a pool storage group
 - Use the multi-tiered storage group enhancement to control SMS volume selection based on the storage group sequence order listed in the automatic class selection (ACS) routine
 - Automatically receive an error message if SMS end-of-volume (EOV) processing fails

See Chapter 8, “Using the Storage Management Subsystem (SMS) enhancements,” on page 29 for more information.

- **Striping enhancements** Prevent storage-related abnormal endings due to extend failures for large sequential data sets by increasing the stripes from 16 to 59. See Chapter 9, “Using the sequential data set striping enhancement,” on page 37 for more information.
- **Virtual storage access method (VSAM) message enhancements** Diagnose problems more easily and quickly because of improved multi-line virtual storage access method (VSAM) messages. See Chapter 10, “Using the VSAM message enhancements,” on page 39 for more information.

If you want information about z/OS V1R5 enhancements for other DFSMS features, such as DFSMSHsm, DFSMSdss, and DFSMSrmm, see the documentation for those features.

What do I need to do to use the z/OS DFSMSdfp V1R5 enhancements?

Before you can implement the z/OS DFSMSdfp V1R5 enhancements at your site, you need to understand the enhancement and how it will fit with your current configuration. Then you might need to perform some tasks in order to start using the enhancement.

In each chapter, you will see an overview about one enhancement, and information about how the enhancement works. There is also a roadmap of the tasks associated with the enhancement. Each chapter includes the tasks that you must perform to use this enhancement and the associated procedures. Not every enhancement requires that you perform every type of task for full use of the enhancement at your site.

These tasks are organized into five types based on the usual tasks performed by DFSMSdfp users like you. The types of tasks include the following:

Administering

Managing the data processing resources used with an IBM program to meet the planned processing goals of an enterprise. Resources include such items as processor cycles, real and virtual storage, networks, nodes, communication paths, programs, data terminals, and queues.

Operating

Starting and stopping programs, checking and controlling programs, recording the status of programs and data, and reacting to abnormal events.

Customizing

Tailoring a program to suit the needs of your installation. Enhancing or extending an IBM program by using services and built-in facilities provided by IBM.

Application programming

Designing, coding, compiling, executing, debugging, and testing application programs. Application programs put your computing system to work to meet the specific needs of your business. All other programming supports the tasks of installing, administering, operating, customizing, or diagnosing.

Diagnosing

Identifying the IBM program that is the source of a programming problem. Describing the problem, comparing it to similar known problems, reporting a new problem, and correcting the problem. It does not include diagnosing hardware problems or user errors (debugging).

Each chapter that follows provides you with the information that you need to understand the new enhancement and begin analyzing how you can use it most effectively on your system. If you need to brush up on your basic knowledge about a specific DFSMSdfp function, you can read the related background information suggested in the chapter. After you have a good understanding of the enhancement, look at the specific tasks you need to complete in order to use this enhancement effectively.

This document does not describe every enhancement included in z/OS V1R5 DFSMS. It only describes those enhancements that require you to perform one of the five types of tasks mentioned above.

Read the *z/OS DFSMS Migration* to find out about all of the new enhancements for z/OS V1R5 DFSMS. You will also see details about the following additional types of tasks you need to perform for each of the enhancements:

Evaluating

Judging the applicability of an IBM program for your installation and deciding whether or not to install the program at your site. It includes deciding which program options are useful for your site, what data processing resources are needed to support the program, and what skills need to be taught to users.

Planning

Making fundamental decisions about the options a program offers. These decisions guide, set limits for, and identify requirements for installation, customization, operation, administration, application programming, and diagnosis tasks. Planning is an ongoing task; decisions are made before installation, evaluated after installation, and revised as appropriate.

Installing

Making a program ready to do useful work. This task includes generating a program, initializing a program, and applying program temporary fixes (PTFs) to a program.

Chapter 2. Using the BPAM enhancements

The basic partitioned access method (BPAM) enhancements in this release provide you with the following benefits:

- You can use BPAM for read-only access of UNIX files including zSeries File System (zFS), hierarchical file system (HFS), Network File System (NFS), and temporary file system (TFS).

UNIX files and directories are protected by security rules provided by z/OS Security Server Resource Access Control Facility (RACF), or other security products. Each UNIX file has permission bits. Permission bits allow you to specify read authority, write authority, or search authority for a directory. They also allow you to specify read, write, or execute authority for a file. You need UNIX read authority for each file that you want to access using BPAM.

- You can also take advantage of the increase in the BPAM concatenation limit to exactly 255 partitioned data set (PDS) extents, PDSE data sets, and UNIX directories in any combination. In previous releases, the concatenation limit was about 120 PDS extents and PDSE data sets.

Before you begin: To understand the concepts related to using UNIX files with DFSMS, see "Processing z/OS UNIX Files" in *z/OS DFSMS: Using Data Sets*.

The following table lists the types of tasks and associated procedures that you must complete in order to fully use this enhancement.

Tasks	Procedures you must perform:
Application programming	<ul style="list-style-type: none">• Set up BPAM access for existing z/OS programs• Verify and set search and read authority• Take advantage of the increased concatenation limit for existing programs

Application programming

In order to use the BPAM enhancements, complete the following tasks:

- Set up BPAM access for existing programs
- Verify and setting search and read authority
- Take advantage of the increased concatenation limit for existing programs

Setting up BPAM access to UNIX files

If you have data in PDS or PDSE format that you also want to make available to UNIX programs, you can consider moving the PDS and PDSE members to a UNIX directory. For example, you might have a macro library or data library in a PDS or PDSE that you must periodically copy to UNIX files for use by another program. Or, you might have UNIX files that you must periodically copy to a PDS or PDSE for use by other programs. If you set up BPAM access of UNIX files for PDS and PDSE data sets, the data is available to other programs.

Before you begin: To understand the basics of modifying JCL, see *z/OS MVS JCL Reference*. To understand more about switching from using PDSs or PDSEs to using

z/OS UNIX files, see *z/OS DFSMS: Using Data Sets*. For information on UNIX commands, see *z/OS UNIX System Services Command Reference*. For information on programming methods to use with UNIX files, see *z/OS UNIX System Services Programming: Assembler Callable Services Reference*. For general information about UNIX, see *z/OS UNIX System Services User's Guide*.

To make existing UNIX files in a directory available to programs that read with BPAM, complete the following steps:

1. Identify the directories and ensure that the file names that you want to use do not exceed eight bytes. If they contain assembler source code, the file names must be upper case. Directory names can be mixed case.

-
2. Create a DD statement or the dynamic allocation equivalent similar to the following for each directory:

Example:

```
//SYSLIB DD PATH='/usr/applics/paytime',  
//          FILEDATA=TEXT,PATHOPTS=ORDONLY
```

PATHOPTS=ORDONLY is required; you can use either FILEDATA=TEXT or FILEDATA=BINARY on each DD statement that uses the PATH keyword. The default value is FILEDATA=BINARY. For more information, see *z/OS DFSMS: Using Data Sets*.

Requirement: In a concatenation, the DD name for each DD after the first one is blank.

-
3. Change your methods of updating the former PDS or PDSE members so that you update UNIX files instead. You can use such methods as a UNIX editor, the UNIX OPUT command, or FTP.

Verifying and setting search and read authority

Before you begin: To understand the basics about security for UNIX files, see "Handling Security for Your Files" in *z/OS UNIX System Services User's Guide*.

Complete the following steps to verify and set search and read authority:

1. Enter the UNIX list command to display the permissions for the UNIX files in your working directory.

Example:

```
ls -W
```

-
2. Verify whether the permissions are set correctly.
 - a. If the permissions are set correctly, you have completed this task.
 - b. If the permissions are not set correctly, continue to the next step.

-
3. To change the permissions for any UNIX file, you must be the owner of the file. If you are the owner of the file, use the chmod command to change the permissions.

Example:

```
chmod a+r MAINSAVE
```

This command gives read permission for the MAINSAVE file to all users and assumes that MAINSAVE is in the current directory.

Increasing the concatenation limit for existing programs

Before you begin: To read about the basic concepts and specific uses of concatenating UNIX files and directories, see "Processing z/OS UNIX System Services Files" in *z/OS DFSMS: Using Data Sets*.

In most cases, you can take advantage of the increased concatenation limits for your application programs by including more DD statements in the JCL or the dynamic allocation equivalent for your application programs.

Complete the following steps to take advantage of the increased concatenation limit for your application programs:

1. If your program depends on the maximum number of data sets in a partitioned concatenation, then you can perform one of the following steps.
 - Ensure that you do not exceed the existing limit of the program, or
 - Change the program to allow for the increased concatenation limit.
-
2. If you need to change these programs to reflect the new partitioned concatenation limit of 255, refer to the documentation for these programs. The required steps depend on the specific programs you are using at your site.
-

Chapter 3. Using the catalog enhancements

The catalog enhancements available with z/OS V1R5 DFSMS provide you with these benefits:

- You can disable or enable JOBCAT and STEPCAT data definition (DD) statements, even on a temporary basis. The default for JOBCAT and STEPCAT DD statements is that they are disabled during processing. You can enable JOBCAT and STEPCAT DD statements temporarily so that you can run those jobs that use those statements rather than having to change the job control language (JCL).
- You can end a service task even if the task is in recall. For example, sometimes a service task does not end correctly when the task was processing a catalog request and the task itself ended. It might be necessary to end this service task because it might be holding a resource needed by another catalog request.
- You receive information in SYMREC records that can help you resolve certain types of catalog problems without having to contact the IBM Support Center. This means that you do not have to produce an SVC dump to get additional information which is both more time-consuming and disruptive. You can disable and enable the creation of SYMREC records if necessary. For example, assume a problem is causing repeated creation of SYMREC records and this is disrupting the management of the SYMREC target data set. Now, you can disable the SYMREC records until the problem is resolved.
- Your catalog performance is improved because unnecessary purges of virtual lookaside facility (VLF) cache for shared catalogs are eliminated.
- You can enable the autoadd option during initial program load (IPL) by using an option in SYS1.PARMLIB rather than manually entering an operator command to enable autoadd after IPL completes. When the first connection is made to the coupling facility, the coupling facility supports enhanced catalog sharing (ECS) for eligible catalogs; ECS requires the coupling facility. By setting the autoadd option in the SYS1.PARMLIB, the option is not lost if the coupling facility structure is lost.
- You can set a usage extent threshold for a catalog so that a warning is displayed when a catalog reaches a specified percentage of this usage extent threshold. With this information, you can perform the necessary tasks to avoid the loss of data that can occur if a catalog becomes full.

Before you begin: Understand the following concepts and how they apply to your system:

- How JOBCAT and STEPCAT DD statements affect your system; see "Planning a Configuration of Catalogs" in *z/OS DFSMS: Managing Catalogs*.
- What a service task does; see "The Catalog Address Space" in *z/OS DFSMS: Managing Catalogs*.
- How the coupling facility and enhanced catalog sharing work together; see "Planning a Configuration of Catalogs" in *z/OS DFSMS: Managing Catalogs*.

For general catalog concepts, see *z/OS DFSMS: Managing Catalogs*.

The following table lists the types of tasks and associated procedures that you must complete in order to fully use this enhancement.

Tasks	Procedures you must perform:
Administering	<ul style="list-style-type: none"> • Enabling JOBCAT and STEPCAT DD statements • Disabling JOBCAT and STEPCAT DD statements • Specifying automatic autoadd enablement • Setting the catalog usage threshold
Diagnosing	<ul style="list-style-type: none"> • Ending a service task abnormally • Disabling and enabling the SYMREC function

Administering

You can implement the functions associated with the catalog enhancements by completing the following tasks:

- Enabling JOBCAT and STEPCAT DD statements
- Disabling JOBCAT and STEPCAT DD statements
- Specifying automatic autoadd enablement
- Setting the catalog usage threshold

Enabling JOBCAT and STEPCAT DD statements

The default for JOBCAT and STEPCAT DD statements is that they are disabled during processing. You can enable them by using the keyword, `JOBSTEP`, for the `MODIFY CATALOG,ENABLE` command. If you enable the JOBCAT and STEPCAT DD statements, they are disabled the next time you re-start the system.

Before you begin: Determine whether JOBCAT and STEPCAT DD statements are in use on your system.

If so, complete these steps to enable the JOBCAT and STEPCAT DD statements temporarily.

1. Enter the following DFSMS command:

```
MODIFY CATALOG,ENABLE(JOBSTEP)
```

2. Press enter.

Result: Both the JOBCAT and STEPCAT DD statements are enabled during processing.

Disabling JOBCAT and STEPCAT DD statements

The default for JOBCAT and STEPCAT DD statements is that they are disabled during processing. If you enabled the JOBCAT and STEPCAT DD statements, you can disable the JOBCAT and STEPCAT DD statements without re-starting the system.

If the JOBCAT and STEPCAT DD statements are enabled on your system, complete these steps to set the JOBCAT and STEPCAT DD statements to the default, disabled.

1. Enter the following DFSMS command:

```
MODIFY CATALOG,DISABLE(JOBSTEPCAT)
```

2. Press enter.

Result: Both the JOBCAT and STEPCAT DD statements are disabled during processing.

Specifying automatic autoadd enablement

You can specify the autoadd option so that coupling facility support of enhanced catalog sharing is automatically available for eligible catalogs when the coupling facility is first connected.

Before you begin: Determine whether autoadd should be automatically enabled. If so, use the SYSCATxx member of SYS1.NUCLEUS or the LOADxx member of SYS1.PARMLIB. For more information on these two methods, see "Defining the Catalog Configuration" in *z/OS DFSMS: Managing Catalogs*.

Specifying autoadd through SYS1.NUCLEUS(SYSCATxx)

Before you begin: For information on how to create the SYS1.NUCLEUS member, SYSCATxx, see *z/OS MVS Initialization and Tuning Reference*.

To specify the autoadd option when you IPL your system, complete these steps:

1. Set the value in column 63 of the SYSCATxx member of SYS1.NUCLEUS to Y.

Example:

```
//SYSCAT EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=A
//SYSIN DD DUMMY
//SYSUT2 DD DSN=SYS1.NUCLEUS(SYSCATLG),DISP=SHR,DCB=(RECFM=U)
//SYSUT1 DD *
SYSRES11 SYS1.MASTERA.ICFCAT Y
/*
```

2. IPL your system.

Result: Autoadd is automatically enabled on your system when the first connection is made to the coupling facility. The catalog address space begins using enhanced catalog sharing.

Specifying autoadd through SYS1.PARMLIB(LOADxx)

Before you begin: For information on how to create the SYS1.PARMLIB member, LOADxx, see *z/OS MVS Initialization and Tuning Reference*.

To specify the autoadd option when you IPL your system, complete these steps:

1. Set the value in column 72 of the LOADxx member of SYS1.PARMLIB to Y.

Example:

```
SYSCAT SYSRES11 SYS1.MASTERA.ICFCAT Y
```

Result: AUTOADD is automatically enabled on your system when the first connection is made. The catalog address space begins using enhanced catalog sharing.

2. IPL your system.
-

Setting the catalog usage threshold

You can specify a catalog usage threshold for the maximum percentage of extents that can be used before a warning message is displayed. If you set this threshold to 0 or 100, the catalog usage threshold option is disabled. By default, the catalog usage threshold is set to 80%.

Before you begin: Determine the appropriate threshold for the catalogs at your site. For example, if the average catalog uses 10% of the space every week, it will take ten weeks for the average catalog to use up the space. If you need two weeks to plan and reorganize the catalog, you could set the usage threshold at 75% or 80%. When the catalog reaches the usage threshold, you are notified so that you have more than two weeks to reorganize the catalog before it is out of space. For more information about defining catalogs, see "Working with the Catalog Address Space" in *z/OS DFSMS: Managing Catalogs*.

To set the catalog usage threshold, complete these steps:

1. Enter the MODIFY CATALOG command.

Example:

```
MODIFY CATALOG,NOTIFYEXTENT(xxx)
```

2. Press enter.

Result: The catalog usage threshold is set.

Diagnosing

You can implement two functions that will help you with system problems by completing the following tasks:

- Specify the keyword on the MODIFY CATALOG command to end a service task.
- Perform an abnormal ending of a service task.

Before you begin: Determine the conditions under which you would end a service task abnormally. For example, you can enter the MODIFY CATALOG, LIST command to see if any task has been running for an extended period of time. Then, determine whether the task is still in process. If not, use one of the following commands to force the end of the task. For more information about service tasks, see "Working with the Catalog Address Space" in *z/OS DFSMS: Managing Catalogs*.

Ending a service task abnormally

To end a service task abnormally, complete these steps:

1. Enter the MODIFY CATALOG command.

Example:

```
MODIFY CATALOG,END(xx),{REDRIVE|NOREDRIVE}[ ,FORCE]
```

or

Enter the MODIFY CATALOG command.

Example:

```
MODIFY CATALOG,ABEND(xx),[,FORCE]
```

2. Press enter.

Result: The service task is ended.

Disabling and enabling the SYMREC function

You can disable the overall SYMREC function temporarily and then enable the SYMREC function when you want to use it again. The value for the SYMREC function is that it is enabled.

Before you begin: Determine the conditions under which you would disable the overall creation of SYMREC records. For example, if an error is filling up your logrec data set, you can disable the SYMREC function temporarily until you correct the error.

To disable the SYMREC function, complete these steps:

1. Enter the MODIFY CATALOG command.

Example:

```
MODIFY CATALOG,DISABLE(SYMREC)
```

2. Press enter.

Result: The SYMREC function is disabled.

To enable the overall SYMREC function once it has been disabled, complete these steps:

1. Enter the MODIFY CATALOG command:

Example:

```
MODIFY CATALOG,ENABLE(SYMREC)
```

2. Press enter.

Result: The setting is changed to the default, enabled.

Chapter 4. Using the enhanced data integrity function for physical sequential data sets

Data integrity is enhanced for physical sequential data sets. If users attempt to access the same data set at the same time, it is considered a violation of integrity. When a violation occurs, you can choose to have a warning message displayed but still allow access or end the program abnormally. Concurrent access for multiple users is allowed for those applications that can control and ensure data integrity on their own.

Restriction: Only z/OS V1R5 systems can detect concurrent output or update access to physical sequential data sets. Concurrent access is not detected on lower level systems.

You enable this function by creating a SYS1.PARMLIB member, IFGPSEDI. In this SYS1.PARMLIB member, you specify whether you want the program to display a warning message or abnormally end when there is a violation. You can also specify which data sets you want to exclude from enhanced data integrity processing. This allows those applications that control sharing on their own to continue to operate.

After you create the SYS1.PARMLIB member, IFGPSEDI, the data integrity function is automatically enabled during an initial program load (IPL), or you can enable it by issuing the start IFGEDI command. Once the function is enabled, you can modify IFGPSEDI if necessary and then enter the start IFGEDI command to enable the revised SYS1.PARMLIB member.

Hint: For more information about the syntax for the SYS1.PARMLIB member, IFGPSEDI, see *z/OS MVS Initialization and Tuning Guide*.

Before you begin: Understand shared sequential data sets and how they are used on your system; see *z/OS DFSMS: Using Data Sets* for more information. Understand how to create a SYS1.PARMLIB member; see *z/OS DFSMS: Using Data Sets* and *z/OS MVS Initialization and Tuning Guide* for more information.

The following table lists the types of tasks and associated procedures that you must complete in order to fully use this enhancement.

Tasks	Procedures you must perform:
Administering	<ul style="list-style-type: none">• Enabling the enhanced data integrity function• Disabling the enhanced data integrity function
Application programming	<ul style="list-style-type: none">• Excluding your applications from enhanced data integrity processing

Administering

By default, the enhanced data integrity function is not enabled. In order to use the enhanced data integrity function, you must enable the enhanced data integrity function on your system. You can also disable the enhanced data integrity function on a temporary basis.

Enabling the enhanced data integrity function

Before you enable the enhanced data integrity function, determine which of the following actions should occur in response to a data integrity violation on your system:

- A warning message displays indicating that a violation has occurred.
- The program abnormally ends.

You also need to determine which data sets to exclude from the enhanced data integrity function.

Before you begin: Understand the basics of creating a SYS1.PARMLIB member. See *z/OS MVS Initialization and Tuning Guide* and *z/OS MVS Initialization and Tuning Reference* for information on how to create a SYS1.PARMLIB member.

To enable the enhanced data integrity function on your system, complete the following steps:

1. Create a new SYS1.PARMLIB member, IFGPSEDI.

Example:

```
{MODE(WARN |ENFORCE |DISABLE)}[,DSN(datasetname)][,DSN(datasetname)]....
```

2. Specify the parameters to set the mode and to exclude the data sets that you do not want processed.

- a. Set the mode using the DFSMS MODE(keyword)

Example:

```
MODE(WARN)
```

Specifying WARN causes a warning message to be displayed, but does not end the program.

Example:

```
MODE(ENFORCE)
```

Specifying ENFORCE causes the program to abnormally end.

Recommendation: To reduce possible data integrity violations, use the MODE(WARN) setting initially. If a violation occurs, you receive warning messages indicating which data sets have a violation. Determine which data sets should be excluded from enhanced data integrity processing. Add those data set names to the IFGPSEDI exclude list, or, modify the application so that it does not violate the enhanced data integrity function.

- b. Exclude data sets that should not be processed using the DSN(keyword).

Example:

```
MODE(WARN),DSN(SYS1.OPX.*)
```

Example:

```
MODE(ENFORCE),DSN(SYS1.PAGEDUMP.*)  
DSN(SYS1.OVERLIB.)
```

Hint: You can specify any number of data set name (DSN) parameters in the list of excluded data sets. However, you should specify a minimum amount of data set names because a large number could affect performance.

If you do not specify any data set names, all shared sequential data sets are processed using the data integrity function.

-
3. Make IFGPSEDI available to all systems in the sysplex, or create an identical SYS1.PARMLIB member, IFGPSEDI, for every system in the sysplex.

Restriction: The enhanced data integrity function cannot be used for data sets that are shared across multiple sysplexes.

-
4. To enable the enhanced data integrity function, enter the S IDFGEDI command or perform an initial program load (IPL) on your system.
-

Modifying the enhanced data integrity function

After you have enabled the enhanced data integrity function on your system, you can modify how this function works by changing the mode. For example, you might want to modify the list of data set names that are being excluded from the enhanced data integrity function. Or, you might want to exclude certain data sets from enhanced data integrity processing temporarily because jobs are abending incorrectly.

Recommendation: Do not update IFGPSEDI to add a data set name to the exclude list while that data set is open for output processing. If the data set is open when you specify that name in the exclude list and start the task, your changes will not take effect.

To modify the enhanced data integrity function on your system, complete the following steps:

1. Update the list of data set names to be excluded from enhanced data integrity processing.

Example:

```
MODE(WARN),DSN(PAYROLL.EMPLNO.*)
```

Example:

```
MODE(FORCE),DSN(PAYROLL.EMPLNO.*)
```

-
2. Enable the enhanced data integrity changes using the start IFGEDI command or perform an IPL on your system.

Example:

```
S IFGEDI
```

Result: The enhanced data integrity function is enabled with the updated SYS1.PARMLIB member, IFGPSEDI, in effect.

Restriction: The enhanced data integrity function does not check the volume serial number of the data sets. Therefore, if data sets on different volumes have the same names, opening a data set on one volume might interfere with opening a data set with the same name on another volume.

Disabling the enhanced data integrity function

After you have enabled the enhanced data integrity function on your system, you can disable the function if necessary.

To disable the enhanced data integrity function on your system, complete the following steps:

1. Change the mode in SYS1.PARMLIB member, IFGPSEDI, to DISABLE.

Example:

```
MODE(DISABLE),DSN(SYS1.PAGEDUMP.*)  
          DSN(SYS1.OVERLIB.)
```

Result: The enhanced data integrity function is temporarily disabled on your system.

-
2. To put the enhanced data integrity function in the disabled state, use the start IFGEDI command or perform an IPL on your system.

Example:

```
S IFGEDI
```

Result: The enhanced data integrity function is disabled with the updated SYS1.PARMLIB member, IFGPSEDI, in effect.

Application programming

After you have enabled the enhanced data integrity function on your system, you might need to perform some specific tasks with regard to your applications.

Excluding your applications from enhanced data integrity processing

Before you begin: Determine which applications require the use of enhanced data integrity support. For more detailed information, see "Sharing Non-VSAM Data Sets" in *z/OS DFSMS: Using Data Sets*.

To exclude data sets from enhanced data integrity processing, specify the list of sequential data sets that you want to exclude from this processing in the SYS1.PARMLIB member, IFGPSEDI. See "Enabling the enhanced data integrity function" on page 16.

Recommendation: Do not update IFGPSEDI to add a data set name to the exclude list while that data set is open for output processing. If the data set is open when you specify that name in the exclude list and start the task, your changes will not take effect until the data set is closed and reopened with the enhanced data integrity function active.

To exclude your applications from enhanced data integrity processing, modify the application so that it sets the DCBEEXPS exclude flag located in the DCBE. Setting this flag prevents OPEN processing from failing.

Requirement: The DCBEEXPS flag is ignored in the DCBE unless the application meets one of the following criteria:

- The application is running in supervisor state.
- The application is running in a system key (0–7).
- The application is APF-authorized.

As an alternative to setting the DCBEEEXPS flag for applications that are APF-authorized, you can specify the NODSI option in one of the following ways:

- NODSI is specified in the program properties table (PPT).
- NODSI is specified when the application is dynamically allocating the data set.

For more information about the NODSI option, see *z/OS MVS Programming: Authorized Assembler Services Guide*.

Chapter 5. Using up to 65 535 file sequence numbers

This enhancement provides you with the ability to use up to 65 535 file sequence numbers (a great increase from the previous maximum of 9999) for the following media:

- Standard label (SL) tapes including standard user label tape (SUL) and leading tape mark (LTM)

Exception: If you use a data set sequence number between 1 and 9999, the format of the data set sequence number field in the HDR1/EOV1/EOF1 label for SL tapes is a 4-byte numeric value in EBCDIC format. If you use a data set sequence number between 10000 and 65535, the first byte in this field is the EBCDIC character '?'; the next three bytes are the data set sequence number in binary format.

- Unlabeled (NL) tapes
- Bypass label processing (BLP)

Requirement: In order to use this enhancement, you must have z/OS DFSMS V1R5 installed.

In order to use file sequence numbers that are greater than 9999 during input access for an uncataloged data set or during creation of a new data set, use the RDJFCB MACRO to get a copy of the JFCB. Update the file sequence number in the job file control block (JFCB) and use the OPEN,TYPE=J macro to access the file. You cannot specify file sequence numbers larger than 9999 in the JCL.

Or, you can use the dsname allocation function of DYNALLOC so that file sequence numbers greater than 9999 can be used during input access for an uncataloged data set, or during creation of a new data set. You must include the DALDSSEQ text unit to specify the relative position of the tape data set up to a maximum of 65535. For more information about using the DYNALLOC macro, see *z/OS MVS Programming: Authorized Assembler Services Guide*.

In order to use file sequence numbers that are greater than 9999 for a cataloged data set, you use only the OPEN macro because the catalog provides the file sequence number.

Before you begin: For more information on the JFCB and to understand how and why to use a certain file sequence, see *z/OS DFSMS: Using Magnetic Tapes*. For more information on the OPEN macro and to understand how and why to use a certain file sequence, see *z/OS DFSMS: Using Data Sets*. For more information on the OPEN,TYPE=J macro and to understand how and why to use a certain file sequence, see *z/OS DFSMSdfp Advanced Services*.

The following table lists the types of tasks and associated procedures that you must complete in order to fully utilize this enhancement.

Topic	Procedures you must perform:
Administering	<ul style="list-style-type: none">• Creating new data sets with file sequence numbers greater than 9999• Processing existing data sets using file sequence numbers greater than 9999

Topic	Procedures you must perform:
Application programming	<ul style="list-style-type: none"> • Modifying applications that use IEFUCBOB because of a change to IECDUCBT • Modifying applications that use IEFJFCBN

Administering

Use only one of the two sets of tasks that follow, depending on whether you are creating a new data set or reading an existing data set.

Creating new data sets with file sequence numbers greater than 9999

Before you begin: Before you implement this enhancement for a new data set, know how to create a new data set; see *z/OS DFSMS: Using Data Sets* for more information. You should also know how to use the RDJFCB macro; see *z/OS DFSMSdfp Advanced Services* for more information.

To create a new data set with file sequence numbers up to 65 535, complete these steps. An excerpt of the code for each step is provided in the example that follows the steps:

1. Use the RDJFCB macro to obtain the JFCB.
2. Store the file sequence number in the JFCB.
3. Use the OPEN,TYPE=J macro to send the updated JFCB to OPEN.

Example:

```
//DD1 DD DSN=DATASET1,UNIT=TAPE,VOL=SER=TAPE01,DISP=(NEW,CATLG),
// LABEL=(1,SL)
...
*
* CREATE A TAPE DATA SET WITH FILE SEQUENCE NUMBER 10011.
*
L 6,=F'10011' CREATE FILE 10011
RDJFCB (DCBAD) READ JFCB
STCM 6,B'0011',JFCBFLSQ STORE NEW FILE SEQ NUMBER IN
JFCB
*
OPEN (DCBAD,(OUTPUT)),TYPE=J CREATE FILE N
...
CLOSE DCBAD CLOSE FILE
...
*
* RETURN RETURN TO SYSTEM CONTROL
DCBAD DCB DDNAME=DD1,DSORG=PS,EXLST=LSTA,MACRF=PM,BLKSIZE=80,RECFM=F
LSTA DS 0F RJFCB EXIT LIST
DC X'87'
DC AL3(JFCBAREA) PTR TO JFCB AREA
JFCBAREA DC 50F'0' JFCB AREA
END
```

Recommendation When you create a new data set with a file sequence number greater than 9999, put the data set in a catalog. This ensures that existing applications can process a data set with a file sequence number greater than 9999.

Processing existing data sets using file sequence numbers greater than 9999

In order to process existing data sets using file sequence numbers greater than 9999, you must determine whether a data set is uncataloged or cataloged. You can use the Data Set Application option in the Interactive Storage Management Facility (ISMF); see *z/OS DFSMS: Using the Interactive Storage Management Facility* for more information about how to use this DFSMS interface.

Processing existing cataloged data sets

To allow up to 65 535 file sequence numbers for an existing cataloged data set, use the OPEN macro to access the data set. The file sequence number is supplied by the catalog; you do not need to perform any additional tasks

Processing existing uncataloged data sets

Before you begin: In order to use file sequence numbers that are greater than 9999 for an uncataloged existing data set, know how to update the file sequence number in the job file control block (JFCB) and use the OPEN,TYPE=J macro. For more information on the JFCB, see *z/OS DFSMS: Using Magnetic Tapes*; for more information on the OPEN,TYPE=J macro, see *z/OS DFSMSdfp Advanced Services*.

To use up to 65 535 file sequence numbers for uncataloged data sets, complete the following steps:

1. Use the RDJFCB macro to obtain the system JFCB.
2. Change the file sequence number in the JFCBFLSQ field in the JFCB to the required sequence number.
3. Use the OPEN,TYPE=J macro to send the updated JFCB to OPEN.

Example:

```
//DD1 DD DSN=DATASET1,UNIT=TAPE,VOL=SER=TAPE01,DISP=(OLD,CATLG),
//      LABEL=(1,SL)
...
*
* CREATE A TAPE DATA SET WITH FILE SEQUENCE NUMBER 10011.
*
      L      6,=F'10011'          CREATE FILE 10011
RDJFCB (DCBAD)                  READ JFCB
      STCM  6,B'0011',JFCBFLSQ   STORE NEW FILE SEQ NUMBER IN
*                                JFCB
      OPEN  (DCBAD,(OUTPUT)),TYPE=J CREATE FILE N
      ...
      CLOSE DCBAD                CLOSE FILE
*
*
* RETURN                          RETURN TO SYSTEM CONTROL
DCBAD DCB DDNAME=DD1,DSORG=PS,EXLST=LSTA,MACRF=PM,BLKSIZE=80,RECFM=F
LSTA  DS    0F                    RJFCB EXIT LIST
      DC    X'87'
      DC    AL3(JFCBAREA)         PTR TO JFCB AREA
JFCBAREA DC 50F'0'                JFCB AREA
      END
```

Application programming

You might have to modify some of your applications programs if they use either of the following macros:

- IECDUCBT
- IEFJFCBN

Before you begin: To understand the layout of the fields in the IECDUCBT and the IEFJFCBN macros, see *z/OS MVS Data Areas, Vol 2 (DCCB-ITZYRETC)*.

Modifying applications that use IECDUCBT

Make any modifications that might be required by your application because of the following changes to the IECDUCBT macro:

- UCBFSEQ is defined as a 16 bit unsigned field.
- UCBFSCT is defined as a 16 bit unsigned field.

Modifying applications that use IEFJFCBN

Make any modifications that might be required by your application because the JFCBFLSQ is defined as a 16 bit unsigned field in the IEFJFCBN macro.

Chapter 6. Using the name-hiding enhancement

Name-hiding is an enhancement which uses RACF to check a user's authority to see a data set name when a generic request is made. With this enhancement, you can restrict access to information about DFSMS data sets to only those users who have read authority. If name-hiding is enabled, users who request a list of data set names using Interactive System Productivity Facility (ISPF) or Interactive Storage Management Facility (ISMF) can see only data set names for which they have read authority; if the user does not have read authority to a specific data set, the name of the data set is not included in the list. If name-hiding is enabled and the user enters the complete data set name, RACF does not check a user's authority to see a data set name.

Name-hiding applies to a request to list a data set name when the request is provided in ways such as the following:

- LISTVTOC command
- IDCAMS LISTCAT command
- CVAF access macros (CVAFDIR; CVAFSEQ; CVAFFILT)
- DADSM OBTAIN macro

With the additional name-hiding security, a user can obtain the name of a data set or information about a data set only when one of the following criteria are met:

- The user specifies a fully qualified data set name.
- The user has the appropriate access authority to the data set.

And, users can obtain a list of data set names if they use one of the following ways:

- The user requests a list of data set names for which the user is authorized by using an operation such as the IDCAMS LISTCAT command. All of the data sets for which the user is authorized will be included in the list. If the user is not authorized for any of the data sets, no data set names will be returned in the list.
- The user accesses the VTOC and is authorized to use the RACF Facility Class STGADMIN.IFG.READVTOC.*volser* resource.

Restrictions: The catalog search interface (CSI) treats fully qualified data set names like generic names. Therefore, if you use the CSI to request a fully-qualified data set name with name-hiding active, the data set name is hidden unless you have at least READ access to the data set.

Neither the CVAF macro or DADSM OBTAIN macro provides the name-hiding function to calling programs that are APF-authorized or running in supervisor state or key zero when name-hiding is active. In the name-hiding environment, these programs must provide the name-hiding function to unauthorized users of the program.

Requirement: In order to use the name-hiding enhancement, you must have z/OS DFSMS V1R5 installed.

Before you begin: To understand the concepts related to system security, see *z/OS Security Server RACF Security Administrator's Guide*, *z/OS Security Server RACF Command Language Reference*, and *z/OS V1R5 Planning for Multilevel Security*.

The following table lists the types of tasks and associated procedures that you must complete in order to fully use this enhancement.

Tasks	Procedures you must perform
Administering	<ul style="list-style-type: none">• Activate name-hiding• Restrict user access

Administering

In order to implement name-hiding, you must work with the person responsible for RACF administration at your site. The RACF administrator activates the name-hiding function using the z/OS Security Server RACF command, SETROPTS MLNAMES, to set options in the STGADMIN.IFG.READVTOC.volser.

Attention: You can implement name-hiding without making any changes to the CVAF macros.

To use the name-hiding function at your site, complete the following steps:

1. Provide the RACF administrator with a complete list of the data set names that require restricted access.

Recommendation: To ensure that your company's information is protected, enable name-hiding for all data sets that contain critical information.

2. Determine which groups of users require access to these restricted data sets.
3. Create a list of the groups of users that require access to these restricted data sets. Give the list to the RACF administrator who will ensure that only these users have access to these restricted data sets.
4. Notify the affected groups of users about any access changes.

Chapter 7. Using the re-indexing of online volumes enhancement

When a volume table of contents (VTOC) index becomes disabled, you should re-build the volume index in order to restore volume performance. Prior to this release, you had to vary the volume offline from all other sharing systems to ensure that the rebuilding of the disabled index would succeed. This enhancement provides you with the ability to enter this command at any time without taking the volume offline. Because you can re-build the index while the volume is still online, applications that use this volume on other systems can continue to operate without an outage. You can also refresh the VTOC and expand the VTOC and INDEX while the volume remains online.

Before you begin: Understand the BUILD IX command including the requirements, restrictions, and value definitions. For more information about ICKDSF, see *Device Support Facilities User's Guide and Reference*.

The following table lists the types of tasks and associated procedures that you must complete in order to fully use this enhancement.

Tasks	Procedures you must perform:
Customizing	<ul style="list-style-type: none">• Ensure appropriate PTFs are installed• Verify level of ICKDSF• Update automation programs
Operating	<ul style="list-style-type: none">• Rebuild the volume index• Refresh the VTOC

Customizing

In order to use this feature across all of the systems at your site, complete the following tasks:

1. For systems that are on a lower level of software than z/OS DFSMS Version 1 Release 5, install the appropriate PTFs:
 - UW95817 for systems using z/OS V1R3 or lower
 - UW95815 for systems using OS/390 Release 10
 2. If you submit a job that rebuilds the index or performs the reformatting operation, verify that ICKDSF Version 1 Release 7 is installed on the system you use to run the job.
 3. If you use any automation programs to restore the INDEX VTOC when certain error conditions occur, make any necessary updates to these programs to use this feature.
-

Requirement: If the appropriate PTFs are not installed on the shared system that includes the volume for which you must re-build the index, ensure that you vary that system offline before you re-build the index. Otherwise, the new index might become disabled.

Operating

Before you begin: For more specific details about the ICKDSF commands, see *Device Support Facilities User's Guide and Reference*.

If a VTOC becomes disabled, complete the following steps to restore the volume performance:

1. Create a new index for the volume using the following ICKDSF command:

```
BUILDX IX DDNAME(ddname)
```

Result: The volume index is rebuilt.

-
2. Refresh the VTOC and expand the VTOC and INDEX using the following ICKDSF command:

```
REFORMAT REFVTOC DDNAME(ddname)
```

Result: If the volume has an active index VTOC, then both the VTOC and INDEX data sets are refreshed. If the volume has an OSVTOC, the VTOC is refreshed.

Chapter 8. Using the Storage Management Subsystem (SMS) enhancements

These SMS enhancements provide you with the following benefits:

Generation data set (GDS) reclaim processing

Generation data sets in the deferred roll-in state are reused when a job that creates these generations is rerun after a failure. Now you can control whether a generation is saved and not overlaid, such as when another job is run that also creates a new generation. Set generation data set (GDS) reclaim processing on or off as needed using one of the following methods:

- Update the GDS_RECLAIM keyword in the IGDSMSxx member of SYS1.PARMLIB and then enter SET SMS,xx command; the update value takes effect without having to do an initial program load (IPL) of your system.
- Enter the SETSMS GDS_RECLAIM(YES|NO) command but the change is retained only until the next IPL is performed.

By default, GDS reclaim processing occurs. To change GDS_RECLAIM to NO in a sysplex, you should make the change on every system in the sysplex. If you change the value for GDS reclaim processing, it does not affect jobs that are already running.

Active control data set (ACDS)

If the source control data set (SCDS) and its backups are lost, you can save the active control data set (ACDS) as the SCDS using a new keyword, SAVESCDS, on the SETSMS command. If you make changes to an active configuration using the SET SMS command, you can save the changes; you do not have to reenter the changes when a new configuration is activated or make changes to the SCDS for these changes by saving the active configuration in a SCDS. This means that you don't have to redefine your base configuration using the interactive storage management facility (ISMF), which is a long and tedious process.

High allocation threshold

You now receive a high-allocation-threshold-exceeded message in the printed log when a pool storage group has exceeded its high allocation threshold after a data set is successfully allocated to it. This message can be the signal, for you as the storage administrator or for an automated program, to perform space management.

For example, if a storage group has 85% defined as the high allocation threshold and that percentage is exceeded after an allocation is performed, a message is displayed on the printed log. This message is only displayed the first time a storage group exceeds its high allocation threshold. An informational message is also displayed when the storage group goes below 80% of its high allocation threshold.

Multi-tiered storage groups

You can now indicate the preferred SMS volume selection of the storage groups, based on the storage group sequence order listed in the automatic class selection (ACS) routine using the multi-tiered storage group enhancement.

A new attribute called *multi-tiered storage group* can be set in one or more storage classes. Assign this attribute to data sets in the storage class ACS routine to instruct SMS allocations to choose the ACS routine assigned storage group sequence order when selecting volumes.

Previously, you could structure your storage groups based on data set sizes such as small, medium, and large. To ensure that small data sets did not fragment large data sets, which can result in allocation or extension failures. Now you can prevent such fragmentation by defining the storage group ACS routine to assign all small data sets to the Small storage group; the Medium or Large storage groups are only used when there is no more space in Small storage group.

For example, assume you have three data sets that are small, medium, and large in size. Each data set has its own storage group; the small data set is in the Small storage group, the medium data set is in the Medium storage group, and the large data set is in the Large storage group. You can define your storage group ACS routine so that when volumes in the Small storage group are not available, SMS selects volumes in the Medium storage group. If the volumes in the Medium storage group are not available, SMS then selects volumes in the Large storage group.

End-of-volume (EOV) failure messages

You now receive an error message if SMS end-of-volume (EOV) processing fails. Those EOV failure messages that were created during SMS extend-to-new-volume processing are added to the joblog and the printed log. Therefore, if an X37 abend occurs, there might be one or more SMS messages that might help you identify the reason for the abnormal ending.

Before you begin: Be familiar with information about the following topics:

- Generation data sets (GDS); see "Processing Generation Data Groups" in *z/OS DFSMS: Using Data Sets*.
- The source control data set (SCDS) and the active control data set (ACDS); see "Preparing for the Storage Management Subsystem" in *z/OS DFSMSdfp Storage Administration Reference*.
- Storage groups and the allocation threshold attribute; see "Defining Storage Groups" in *z/OS DFSMSdfp Storage Administration Reference*.
- Storage groups, storage classes, volume selection, and the automatic class selection (ACS) routine; see "Defining Storage Classes" in *z/OS DFSMSdfp Storage Administration Reference*.
- The Interactive Storage Management Facility (ISMF) interface; see *z/OS DFSMS: Using the Interactive Storage Management Facility*.

The following table lists the types of tasks and associated procedures that you must complete in order to fully use this enhancement.

Tasks	Procedures you must perform
Administering	<ul style="list-style-type: none"> • Disabling GDS reclaim processing • Enabling GDS reclaim processing • Setting the multi-tiered storage group option

Administering

In order to use the following enhancements to the SMS features, complete the tasks described in the sections that follow:

- Disabling GDS reclaim processing
- Enabling GDS reclaim processing
- Setting the multi-tiered storage group option

Disabling GDS reclaim processing

By default, GDS reclaim processing occurs. You can choose to disable the GDS reclaim processing option depending on the needs of your system, using one of two methods:

- Disable GDS reclaim processing by changing the PARMLIB member, IGDSMSxx, if you want to disable GDS reclaim processing indefinitely.
- Disable GDS reclaim processing by issuing an operator command if you want to disable GDS reclaim processing only until the next initial program load (IPL).

Before you begin: Determine if GDS reclaim processing should be disabled on your system.

Disabling GDS reclaim processing by changing the PARMLIB IGDSMSxx member

Before you can disable GDS reclaim processing, you must understand how to complete the following tasks:

- Update the SYS1.PARMLIB; see *z/OS DFSMSdfp Storage Administration Reference*.
- Use the access method services (AMS) DELETE and ALTER commands; see *z/OS DFSMS Access Method Services for Catalogs*.

To disable GDS reclaim processing by changing the PARMLIB member, IGDSMSxx:

1. Change the SMS keyword GDS_RECLAIM(YES) in the IGDSMSxx member of SYS1.PARMLIB.

Example:

```
SMS GDS_RECLAIM(NO)
```

-
2. Enter the following SMS command:

Example:

```
SET SMS=xx
```

Result: GDS reclaim processing is disabled on this SMS subsystem even if you perform an initial program load (IPL).

-
3. Ensure that the generation data set (GDS) that did not get rolled in is either deleted, renamed, or rolled in by using one of the methods that follows.

Before you begin: To ensure that you use all of the required options correctly for these Access Method Services (AMS) commands, see *z/OS DFSMS Access Method Services for Catalogs*.

- Use the AMS DELETE command to delete the GDS:

Example:

```
DELETE MY.DATASET.GxxVyy
```

Result: The deferred generation data set, MY.DATASET.GxxVyy, is deleted.

- Use the AMS ALTER RENAME to rename the GDS:

Example:

```
ALTER MY.DATASET.GxxVyy NEWNAME(MYDATASET.SAVE1)
```

Result: The deferred generation data set, MY.DATASET.GxxVyy, is renamed to MYDATASET.SAVE1.

- Use the AMS ALTER ROLLIN command to roll in the GDS if its status is deferred rolled-in or rolled-off:

Example:

```
ALTER MY.DATASET.GxxVyy ROLLIN
```

Result: The deferred generation data set, MY.DATASET.GxxVyy, is rolled into the generation data group, MY.DATASET.

4. Verify that all systems in the sysplex that are sharing GDSs are using this same setting for GDS reclaim processing.

Restriction: If different systems in a sysplex are at different levels of SMS or specify different settings for GDS reclaim processing, the GDSs should not be shared.

Disabling GDS reclaim processing by issuing an operator command

Before you can disable GDS reclaim by issuing an operator command, you must understand how to complete the following tasks:

- Update the SYS1.PARMLIB; see *z/OS DFSMSdfp Storage Administration Reference*.
- Use the access method services (AMS) DELETE and ALTER commands; see *z/OS DFSMS Access Method Services for Catalogs*.

To disable GDS reclaim processing by issuing an operator command:

1. Enter the following SMS operator command:

Example:

```
SETSMS GDS_RECLAIM(NO)
```

Result: GDS reclaim processing is disabled on this SMS subsystem until you perform an initial program load (IPL).

2. Ensure that the generation data set (GDS) that did not get rolled in is either deleted, renamed, or rolled in by using one of these methods.

- Use the access method services (AMS) DELETE command to delete the GDS:

Example:

```
DELETE MY.DATASET.GxxVyy
```

Result: The deferred generation data set, MY.DATASET.GxxVyy, is deleted.

-
- Use the AMS ALTER RENAME to rename the GDS:

Example:

```
ALTER MY.DATASET.GxxVyy NEWNAME(MYDATASET.SAVE1)
```

Result: The deferred generation data set, MY.DATASET.GxxVyy, is renamed to MYDATASET.SAVE1.

-
- Use the AMS ALTER ROLLIN command to roll in the GDS if its status is deferred rolled-in or rolled-off:

Example:

```
ALTER MY.DATASET.GxxVyy ROLLIN
```

Result: The deferred generation data set, MY.DATASET.GxxVyy, is rolled into the generation data group, MY.DATASET.

-
3. Verify that all systems in the sysplex that are sharing GDSs are using this same setting for GDS reclaim processing.

Restriction: If different systems in a sysplex are at different levels of SMS or specify different settings for GDSG reclaim processing, the GDSs should not be shared.

Enabling GDS reclaim processing

To enable GDS reclaim processing after you have disabled it, complete the following steps:

1. Use the SETSMS command.

Example:

```
SETSMS GDS_RECLAIM(YES)
```

Result: GDS reclaim processing is enabled on this SMS subsystem. The next time an IPL is performed, the default setting for GDS reclaim processing is restored.

Restriction: Changes to settings for GDS reclaim processing do not take effect for jobs that are already running.

-
2. Verify that all systems in the sysplex that are sharing GDSs are using this same setting for GDS reclaim processing.

Restriction: If different systems in a sysplex are at different levels of SMS or specify different settings for GDS reclaim processing, then the GDSs should not be shared.

Setting the multi-tiered storage group option

Before you begin: Because the multi-tiered storage group option is an attribute of storage class, you must determine which storage classes should use SMS volume selection based on the storage group sequence order listed in the automatic class

selection (ACS) routine. For more information on storage classes and ACS routines, see *z/OS DFSMSdfp Storage Administration Reference*. To understand how to use ISMF, see *z/OS DFSMS: Using the Interactive Storage Management Facility*. To understand how to activate a new or changed SMS configuration, see *z/OS DFSMSdfp Storage Administration Reference*, "Activating Storage Management Subsystem Configurations".

To use the multi-tiered storage group option, complete one of the following tasks:

- Set the multi-tiered storage group option for an existing storage class
- Define a new storage class with the multi-tiered storage group option enabled

Setting the multi-tiered storage group option for an existing storage class

To set the multi-tiered storage group option for an existing storage class, complete the following steps:

1. Use ISMF to access the Storage Class application.

-
2. Select the Storage Class Alter option. Enter Y in the *Multi-Tiered SG* field.

Example:

Multi-Tiered SG (Y)

-
3. Activate the updated configuration using one of the following methods:

- a. On the CDS APPLICATION SELECTION panel in ISMF, complete one of the following actions:

Select the Activate option

or

Enter the ACTIVATE operator command on the command line.:

Example:

ACTIVATE SCDS dsname

- b. From the operator console, enter the SETSMS command.

Example:

SETSMS SCDS dsname

Result: The multi-tiered storage group option is enabled for this storage class.

-
4. Examine the ACS storage group selection routine to ensure the selection sequence order is appropriate. If necessary, you can modify the ACS storage group selection routine to allow storage groups to borrow space from each other.

For example, assume the storage group ACS routines assign three storage groups in the following order for a particular allocation: Small, Medium, Large. If the Small storage group space is low, SMS attempts to allocate space to the second storage group listed, Medium. And, if the Medium storage group is low on storage space, the third choice, Large, is selected. If volumes in all three storage groups are low on space, SMS selects volumes in the Small storage group.

Defining a new storage class with the multi-tiered storage group option enabled

To define a new storage class with the multi-tiered storage group option enabled, complete the following steps:

1. Use ISMF to access the Storage Class application.

2. Select the Storage Class Define option to create a new storage class. Enter Y in the *Multi-Tiered SG* field.

Example:

```
Multi-Tiered SG (Y)
```

3. Complete the remaining entries on the Storage Class Define option panel.

4. Activate the new storage class using one of the following methods:

- a. On the CDS APPLICATION SELECTION panel in ISMF, complete one of the following actions:

Select the Activate option

or

Enter the ACTIVATE command on the command line.:

Example:

```
ACTIVATE SCDS dsname
```

- b. From the operator console, enter the SETSMS command.

Example:

```
SETSMS SCDS dsname
```

Result: The new storage class is activated.

5. Examine the ACS storage group selection routine to ensure the selection sequence order is appropriate. If necessary, you can modify the ACS storage group selection routine to allow storage groups to borrow space from each other.

For example, assume the storage group ACS routines assign three storage groups in the following order for a particular allocation: Small, Medium, Large. If the Small storage group space is low, SMS attempts to allocate space to the second storage group listed, Medium. And, if the Medium storage group is low on storage space, the third choice, Large, is selected. If volumes in all three storage groups are low on space, SMS selects volumes in the Small storage group.

Chapter 9. Using the sequential data set striping enhancement

With sequential data set striping increased from 16 to 59 stripes, your data can now be allocated to as many as 59 volumes which can prevent storage-related abnormal endings. Physical sequential striped data set sizes are limited by the number of extensions possible on any of the striped volumes. Previously, large data sets were limited by the striping count of 16; now you can stripe large sequential data sets on up to 59 volumes. Striping allows you to spread data across different DASD volumes and controllers. The number of stripes is the number of volumes on which the data set is initially allocated. Striped data sets must be system-managed and must be in an extended format. The sustained data rate (SDR) has an effect only for extended format striped data sets.

Before you begin: To understand more about sequential data sets, see *z/OS DFSMS: Using Data Sets*. To understand the concept of striping and how it works to make more efficient use of your storage, see *z/OS DFSMSdfp Storage Administration Reference*.

The following table lists the types of tasks and associated procedures that you must complete in order to fully use this enhancement.

Topic	Procedures you must perform
Administering	<ul style="list-style-type: none">• Set the correct sustained data rate

Administering

Before you begin: For information on how to set the sustained data rate, see "Defining Sustained Data Rate" in *z/OS DFSMSdfp Storage Administration Reference*. See *z/OS DFSMS: Using the Interactive Storage Management Facility* for information about Interactive Storage Management Facility (ISMF).

Requirement: If your system is using an earlier release than z/OS V1R5 to access extended format sequential data sets with more than 16 stripes, you must install the following PTFs:

- For z/OS V1R3, PTF UW87192
- For OS/390 V2R10, PTF UW87191

To use the sequential data set striping enhancement effectively on your site, complete the following steps:

1. Review the current setting for the sustained data rate in the Sustained Data Rate field on the ISMF panel, Defining Storage Class Attributes.

Example:

Sustained Data Rate (MB/sec) . . . (0 to 9999 or blank)

Hint: If the value in the Sustained Data Rate field is greater than 1, it is divided by 4 for 3390s, or by 3 for 3380s to determine the stripe count, even if the 3390 or 3380 is being emulated by another device such as an IBM Enterprise Storage Server.

-
2. Modify the sustained data rate as appropriate for your system. The larger the value you specify for the sustained data rate, the greater the number of stripes used by your sequential data sets.
-

Chapter 10. Using the VSAM message enhancements

VSAM messages include the following enhancements to help you diagnose problems:

- Messages with more than one line are displayed as single multi-line messages. No other information is displayed between the beginning and end of the VSAM message; all of the information for a single message appears on adjacent lines. This makes it easier for the operator to relate the various pieces of information to each other and interpret the information more quickly.
- When a sharing option violation occurs, the name of the first job causing the violation is provided in the VSAM OPEN message, IEC161I 052. This information is included in the subfunction information (*sfi*) field and can help you diagnose problems.
- When a catalog locate request fails, VSAM messages, IEC070I 032 and IEC070I 034, contain a value in the *sfi* field. You can use this value to aid you with problem determination.

The following table lists the types of tasks and associated procedures that you must complete in order to fully use this enhancement.

Tasks	Procedures you must perform
Application programming	<ul style="list-style-type: none">• Review applications that process VSAM messages• Modify applications as required

Application Programming

For VSAM messages, all of the information for one message is displayed on adjacent lines; no other information is displayed between the beginning and end of a single message.

Implementing the message improvements

In order to implement the message improvements with your application programs, complete the following steps:

1. Review any application or automation program that processes VSAM messages to ensure that it is not affected by the single multi-line message format.

Example:

```
JOB00037 00000010      WTOTEST   READY2GO   WTORPGM   0000
JOB00037 00000010      IECYUK0   START OF WTO TEST
JOB00037 00000010      IECYUK2   INSERTED MSG 1
JOB00037 00000010      IECYUK2   INSERTED LINE 2
JOB00037 00000010      IEC070I   195(00000004,00000009)-220,STRPOCE2,TESTIT,DD3,, 636
        636 00000010      IEC070I   IBMUSER.XXX1.KSDS,IBMUSER.XXX1.KSDS.DATA,
        636 00000010      IEC070I   SYS1.MVSRES.MASTCAT
JOB00037 00000010      IECYUK3   INSERTED LINE 3
JOB00037 00000010      IECYUK4   WTO END OF TEST
JOB00037 00000010      IECYUK5   WTO END OF THE TEST CASE - GO HOME !!!
JOB00037 00000010      WTOTEST   STEP1     WTOTEST   0B10
JOB00037 00000010      $HASP395 WTOTEST   ENDED
        00000010      $HASP309 INIT 1   INACTIVE ***** C=ABJ
```

If any application or automation program is affected by the single multi-line message format, contact the person responsible for that program.

2. Review any application that processes VSAM messages to ensure that it is not affected by the additional optional information in the optional *sfi* field.
If any application or automation program is affected by the additional optional information, contact the person responsible for that program.
-

Appendix. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS enable users to:

- Use assistive technologies such as screen-readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

Using assistive technologies

Assistive technology products, such as screen-readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using it to access z/OS interfaces.

Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. Refer to *z/OS TSO/E Primer*, *z/OS TSO/E User's Guide*, and *z/OS ISPF User's Guide Volume I* for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Programming interface information

This publication documents information that is NOT intended to be used as Programming Interfaces of DFSMS.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States, or other countries, or both:

IBM
DFSMSHsm
DFSMSrmm
IBMLink
MVS
RACF
Resource Link
SecureWay
OS/390
z/OS
z/VM

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

This glossary defines technical terms and abbreviations used in DFSMS documentation. If you do not find the term you are looking for, refer to the index of the appropriate DFSMS manual or view the *Glossary of Computing Terms* located at:

<http://www.ibm.com/ibm/terminology/>

This glossary includes terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036. Definitions are identified by the symbol (A) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published part of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- The *IBM Dictionary of Computing*, New York: McGraw-Hill, 1994.

The following cross-reference is used in this glossary:

See: This refers the reader to (a) a related term, (b) a term that is the expanded form of an abbreviation or acronym, or (c) a synonym or more preferred term.

A

access method services. A multifunction service program that manages both VSAM and non-VSAM data sets and integrated catalog facility catalogs or VSAM catalogs. It defines data sets and allocates space for VSAM data sets, VSAM catalogues, and ICF

catalogs. It converts indexed-sequential (ISAM) data sets to key-sequenced data sets, modifies data set attributes in the catalog, reorganizes data sets, facilitates data portability among operating systems, creates backup copies of data sets and indexes, helps make inaccessible data sets accessible, lists the records of data sets and catalogs, defines and builds alternate indexes, and converts CVOLs and VSAM catalogs to integrated catalog facility catalogs.

access permission. A group of designations that determine who can access a particular AIX or UNIX file and how the user can access the file.

ACS. Automatic class selection.

active control data set (ACDS). A VSAM linear data set that contains an SCDS that has been activated to control the storage management policy for the installation. When activating an SCDS, you determine which ACDS will hold the active configuration (if you have defined more than one ACDS). The ACDS is shared by each system that is using the same SMS configuration to manage storage. See also *source control data set*, *communications data set*.

alias. An alternative name for an ICF user catalog, a non-VSAM file, or a member of a partitioned data set (PDS) or PDSE.

automatic class selection (ACS). A mechanism for assigning Storage Management Subsystem classes and storage groups to data sets.

B

basic partition access method (BPAM). An access method that can be applied to create program libraries in direct access storage for convenient storage and retrieval of programs.

BPAM. See *Basic partitioned access method*.

C

catalog . (1) A directory of files and libraries, with reference to their locations. A catalog may contain other information such as the types of devices in which the files are stored, passwords, blocking factors. (2) A data set that contains extensive information required to locate other data sets, to allocate and deallocate storage space, to verify the access authority of a program or operator, and to accumulate data set usage statistics. (A) (ISO) (3) To enter information about a file or a library into a catalog. (A) (ISO) (4) The collection of all data set indexes that are used by the control program

to locate a volume containing a specific data set. (5) To include the volume identification of a data set in the catalog. (6) See *VSAM master catalog*, *VSAM user catalog*.

CAS. catalog address space.

catalog address space. The area of virtual storage where catalog functions are performed. It contains tables with all user catalog names identified in the master catalog, their aliases, and their associated volume serial numbers. Any changes to the master catalog are automatically reflected in these tables.

coupling facility (CF). The hardware that provides high-speed caching, list processing, and locking functions in a Parallel Sysplex.

D

DASD volume. A DASD space identified by a common label and accessed by a set of related addresses. See also *volume*, *primary storage*, *migration level 1*, *migration level 2*.

data class. A collection of allocation and space attributes, defined by the storage administrator, that are used to create a data set.

device. This term is used interchangeably with unit. For a disk or tape, a unit on which a volume may be mounted. For example, a tape drive is a device; a tape cartridge is a volume. Device also applies to other types of equipment, such as a card reader or a channel-to-channel (CTC) adapter.

Device Support Facilities (ICKDSF). A program used for initialization of DASD volumes and track recovery.

DFSMS. See *Data Facility Storage Management Subsystem*.

DFSMSdfp. A DFSMS functional component or base element of z/OS, that provides functions for storage management, data management, program management, device management, and distributed data access.

DFSMSdss. A DFSMS functional component or base element of z/OS, used to copy, move, dump, and restore data sets and volumes.

DFSMShsm. A DFSMS functional component or base element of z/OS, used for backing up and recovering data, and managing space on volumes in the storage hierarchy.

DFSMS environment. An environment that helps automate and centralize the management of storage. This is achieved through a combination of hardware, software, and policies. In the DFSMS environment for MVS, this function is provided by DFSMS, DFSORT, and RACF. See also *system-managed storage*.

direct access device space management (DADSM). A collection of subroutines that manages space on disk volumes. The subroutines are: Create, Scratch, Extend, and Partial Release.

dump. A capture of valuable storage information at the time of an error.

E

ECS. Enhanced Catalog Sharing.

extended format. The format of a data set that has a data set name type (DSNTYPE) of EXTENDED. The data set is structured logically the same as a data set that is not in extended format but the physical format is different. Data sets in extended format can be striped or compressed. Data in an extended format VSAM KSDS can be compressed. See also *striped data set*, *compressed format*.

extent. A file extent is a storage area for records allocated to a file by the server. Extents are not formally architected in DDM.

G

GDG. See *generation data group*.

GDS. See *generation data set*.

generation data group (GDG). A collection of historically related non-VSAM data sets that are arranged in chronological order; each data set is a generation data set.

generation data set. One generation of a generation data group.

I

ICKDSF. See *Device Support Facilities program*.

initial program load (IPL). (1) The initialization procedure that causes an operating system to commence operation. (2) The process by which a configuration image is loaded into storage at the beginning of a work day or after a system malfunction. (3) The process of leading system programs and preparing a system to run jobs. (4) Synonymous with system restart, system startup.

Interactive Storage Management Facility (ISMF). The interactive interface of DFSMS/MVS that allows users and storage administrators access to the storage management functions.

Interactive System Productivity Facility (ISPF). An interactive base for ISMF.

IPL. See *initial program load*.

ISMF. See *Interactive Storage Management Facility* .

ISPF. See *Interactive System Productivity Facility*.

J

JCL. See *Job control language*.

Job control language (JCL). A problem-oriented language used to identify the job or describe its requirements to an operating system.

M

master catalog. A key-sequenced data set or file with an index containing extensive data set and volume information that VSAM requires to locate data sets or files, to allocate and deallocate storage space, to verify the authorization of a program or operator to gain access to a data set or file, and to accumulate usage statistics for data sets or files.

multilevel security. A security policy that allows the classification of data and users based on a system of hierarchical security levels (for example: unclassified, secret, top secret) combined with a system of non-hierarchical security categories (for example: Project A, Project B, Project C). In order to access data, a user must have a security level greater than or equal to that of the data, and be authorized to all of the categories assigned to the data.

N

name hiding. Prevents unauthorized users from obtaining names about data sets.

nonVSAM data set. A data set allocated and accessed using one of the following methods: BDAM, BPAM, BISAM, BSAM, QSAM, QISAM.

P

partitioned data set (PDS). A data set on direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

partitioned data set extended (PDSE). A system-managed data set that contains an indexed directory and members that are similar to the directory and members of partitioned data sets. A PDSE can be used instead of a partitioned data set.

PDS. See *Partitioned data set*.

PDSE. See *partitioned data set extended*.

performance. (1) A measurement of the amount of work a product can produce with a given amount of

resources. (2) In a system-managed storage environment, a measurement of effective data processing speed with respect to objectives set by the storage administrator. Performance is largely determined by throughput, response time, and system availability.

pool storage group. A type of storage group that contains system-managed DASD volumes. Pool storage groups allow groups of volumes to be managed as a single entity. See also *storage group*.

R

RACF. See *Resource access control facility*.

Resource Access Control Facility (RACF). An IBM licensed program that is included in z/OS Security Server and is also available as a separate program for the z/OS and VM environments. RACF provides access control by identifying and verifying the users to the system, authorizing access to protected resources, logging detected unauthorized attempts to enter the system, and logging detected accesses to protected resources.

S

SCDS. See *Source control data set*.

sequential data set. A data set whose records are organized on the basis of their successive physical positions, such as on magnetic tape. Contrast with *direct data set*.

SMS. Storage Management Subsystem.

SMS class. A list of attributes that SMS applies to data sets having similar allocation (data class), performance (storage class), or backup and retention (management class) needs.

SMS-managed data set. A data set that has been assigned a storage class.

source control data set (SCDS). A VSAM linear data set containing an SMS configuration. The SMS configuration in an SCDS can be changed and validated using ISMF. See also *active control data set*, *communications data set*.

storage administrator. A person in the data processing center who is responsible for defining, implementing, and maintaining storage management policies.

storage class. A collection of storage attributes that identify performance goals and availability requirements, defined by the storage administrator, used to select a device that can meet those goals and requirements.

storage group. A collection of storage volumes and attributes, defined by the storage administrator. The collections can be a group of DASD volumes or tape volumes, or a group of DASD, optical, or tape volumes treated as a single object storage hierarchy. See also *VIO storage group, pool storage group, tape storage group, object storage group, object backup storage group, dummy storage group.*

storage management. The activities of data set allocation, placement, monitoring, migration, backup, recall, recovery, and deletion. These can be done either manually or by using automated processes. The Storage Management Subsystem automates these processes for you, while optimizing storage resources. See also *Storage Management Subsystem.*

Storage Management Subsystem (SMS). A DFSMS facility used to automate and centralize the management of storage. Using SMS, a storage administrator describes data allocation characteristics, performance and availability goals, backup and retention requirements, and storage requirements to the system through data class, storage class, management class, storage group, and ACS routine definitions.

stripe. In DFSMS, the portion of a striped data set, such as an extended format data set, that resides on one volume. The records in that portion are not always logically consecutive. The system distributes records among the stripes such that the volumes can be read from or written to simultaneously to gain better performance. Whether it is striped is not apparent to the application program.

sysplex. A set of z/OS systems communicating and cooperating with each other through certain multisystem hardware components and software services to process customer workloads.

system-managed storage. Storage managed by the Storage Management Subsystem. SMS attempts to deliver required services for availability, performance, and space to applications. See also *system-managed storage environment.*

system programmer. A programmer who plans, generates, maintains, extends, and controls the use of an operating system and applications with the aim of improving overall productivity of an installation.

T

threshold. A storage group attribute that controls the space usage on DASD volumes, as a percentage of occupied tracks versus total tracks. The *low migration threshold* is used during primary space management and interval migration to determine when to stop processing data. The *high allocation threshold* is used to determine candidate volumes for new data set allocations. Volumes with occupancy lower than the

high threshold are selected over volumes that meet or exceed the high threshold value.

U

UNIX. A highly portable operating system originally developed by Bell Laboratories that features multiprogramming in a multi-user environment. UNIX is implemented in the C language. UNIX was originally developed for use on minicomputers but has been adapted on mainframes and microcomputers. It is especially suitable for multiprocessor, graphics, and vector-processing systems.

V

VLF. Virtual lookaside facility

virtual storage access method (VSAM). An access method for direct or sequential processing of fixed and variable-length records on direct access devices. The records in a VSAM data set or file can be organized in logical sequence by a key field (key sequence), in the physical sequence in which they are written on the data set or file (entry-sequence), or by the relative-record number.

volume. The storage space on DASD, tape, or optical devices, which is identified by a volume label. See also *DASD volume, optical volume, tape volume.*

VSAM. See *virtual storage access method.*

Z

z/OS. z/OS is a network computing-ready, integrated operating system consisting of more than 50 base elements and integrated optional features delivered as a configured, tested system.

z/OS Network File System. A base element of z/OS, that allows remote access to z/OS host processor data from workstations, personal computers, or any other system on a TCP/IP network that is using client software for the Network File System protocol.

Index

A

accessibility 41
ACDS 29
APAR OW52855 27

B

BPAM concatenation limit 5
BPAM enhancements
 description 5
 enabling with application
 programs 7
 setting SEARCH and READ
 authority 6
 verifying SEARCH and READ
 authority 6

C

catalog enhancements
 description 9
 disabling JOBCAT and STEPCAT DD
 statements 10
 disabling the SYMREC function 13
 enabling JOBCAT and STEPCAT DD
 statements 10
 enabling the SYMREC function 13
 ending a service task abnormally 12
 setting the catalog usage
 threshold 12
 specifying AUTOADD through
 SYS1.NUCLEUS(SYSCATxx) 11
 specifying AUTOADD through
 SYS1.PARMLIB(LOADxx) 11

commands

ALTER_RENAME 32, 33
ALTER_ROLLIN 32, 33
BUILD IX 27
CVAFDIR 25
CVAFFILT 25
CVAFSEQ 25
DADSM obtain 25
DELETE_NONVSAM 32, 33
LISTCAT 25
LISTVTOC 25
MODIFY CATALOG,
 END(xx),[REDRIVE | NOREDRIVE],[FORCE] 12
MODIFY
 CATALOG,ABEND(xx),[FORCE] 13
MODIFY
 CATALOG,DISABLE(JOBSTPCAT) 11
MODIFY
 CATALOG,DISABLE(SYMREC) 13
MODIFY
 CATALOG,ENABLE(JOBSTPCAT) 10
MODIFY
 CATALOG,ENABLE(SYMREC) 13
MODIFY
 CATALOG,NOTIFYEXTENT(xxx) 12
S IFGEDI 17, 18

commands (continued)

SET SMS=SAVESCDS 29
SET SMSxx 29
SETSMS GDS_RECLAIM(YES) 33
SMS GDS_RECLAIM(NO) 31

D

disability 41
documents, licensed vi

E

enhanced data integrity function
 description 15
 enabling 16
 excluding applications from data
 integrity processing 18
 modifying 17, 18

F

file sequence numbers up to 65 535
 description 21
 modifying applications that use
 IECDUCBT 24
 modifying applications that use
 IEFJFCBN 24
 specifying for existing data sets 23
 specifying for new data sets 22

G

GDS reclaim processing 29, 31

H

high-allocation-threshold 29

K

keyboard 41

L

licensed documents vi
LOADxx member of SYS1.PARMLIB
 example 11
LookAt message retrieval tool vi

M

macros
 IECDUCBT 24
 IEFJFCBN 24
 RDJFCB 23
message retrieval tool, LookAt vi

MODE(DISABLE)
 example 18
MODE(ENFORCE)
 example 16
MODE(FORCE)
 example 17
MODE(WARN) 16
 example 16, 17
multi-tiered storage group option 33

N

name-hiding
 description 25
 implementing 26

P

physical sequential data set striping
 enhancement
 See striping enhancement

R

re-indexing online volumes 27
ICKDSF requirements 27
run-time procedure requirements 27
software requirements 27

S

SCDS 29
sequential data set striping enhancement
 See striping enhancement
shortcut keys 41
SMS enhancements 29
 disabling GDS reclaim processing 31
 enabling GDS reclaim processing 33
 using the multi-tiered storage group
 option 33
SMS EOVS processing 30
SMS volume selection of storage
 groups 30
storage management subsystem
 enhancements
 See SMS enhancements
striping 37
striping enhancement
 description 37
 using 37
sustained data rate 37
SYS1.PARMLIB members
 IFGPSEDI 16, 17, 18
 IGDSMSxx 29
 LOADxx 11
SYSCATxx member of SYS1.NUCLEUS
 example 11

T

tapes

- bypass label processing (BLP) 21
- standard label (SL) tapes 21
- unlabeled (NL) tapes 21

U

UNIX files

- list command 6
- security 5
- types 5

V

virtual storage access method (VSAM)

enhancements

- See* VSAM enhancements

VSAM enhancements

- description 39
- implementing the message improvements 39

VSAM messages

- catalog locate request fails 39
- format example 39
- sharing option violation 39

Readers' Comments — We'd Like to Hear from You

z/OS

DFSMS: Using DFSMSdfp in the z/OS V1R5 Environment

Publication No. SC26-7473-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY 12601-5400



Fold and Tape

Please do not staple

Fold and Tape



Program Number: 5694-A01

Printed in USA

SC26-7473-00

