

z/OS



DFSMSdfp Checkpoint/Restart

Version 2 Release 1

z/OS



DFSMSdfp Checkpoint/Restart

Version 2 Release 1

Note

Before using this information and the product it supports, read the information in "Notices" on page 95.

This edition applies to Version 2 Release 1 of z/OS (5650-ZOS) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1976, 2013.**

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

Contents

Figures	v
--------------------------	----------

Tables	vii
-------------------------	------------

About This Book.	ix
-----------------------------------	-----------

Required product knowledge	ix
--------------------------------------	----

z/OS information	ix
----------------------------	----

How to send your comments to IBM	xi
---	-----------

If you have a technical problem.	xi
--	----

z/OS Version 2 Release 1 summary of changes	xiii
--	-------------

Chapter 1. Introduction	1
--	----------

Components	1
CHKPT Macro Instruction	1
End-of-Volume Exit Routine	1
Checkpoint at End-of-Volume Facility	1
Restart Definition (RD) Parameter	2
RESTART Parameter	2
SYSCHK DD Statement.	2
SYSCKEOV DD Statement.	2
Dependencies	2
Controlled Use of Checkpoints	2
Specification of Codes Eligible for Restart.	2
Job Journal Requirements	3
User Data Set Security	3
Checkpoint Data Set Security	4

Chapter 2. Coding a Checkpoint	5
---	----------

When to Code a Checkpoint	5
Checkpoint at End-of-Volume.	6
Checkpoint in Exit Routines	7
Coding the DCB for a Checkpoint Data Set	7
Required DCB Parameters	8
DCB Options	8
Notes on the DCB for a Checkpoint Data Set	8
Coding the DD Statement for a Checkpoint Data Set	8
Notes on the DD Statement	9
Using Checkpoint Data Sets	10
Storage Estimates for Checkpoint Data Sets.	10
How Checkpoint Entries Are Written	11
How to Ensure Restart with Sequential Checkpoint Data Sets	13
How Checkpoint Entries Are Identified	14
Coding the CHKPT Macro Instruction	15
CHKPT Return Codes	18
Use of CHKPT with Other Macro Instructions.	18
Canceling a Checkpoint	20
Serially Reusable Resources	20
Requesting Serially Reusable Resources	20
Special Operating System Features.	22
VIO Data Sets	22

Dynamic Allocation and Dynamically Concatenated Data Sets	22
Shared Resources	23
Cross-Memory Support Restrictions	23
Data Spaces	24
Name/Token Services	24

Chapter 3. Requesting Restart	25
--	-----------

Types of Restart	25
Automatic Restart	25
Operator Message Sequence	25
Operator Considerations	26
How to Request Automatic Checkpoint/Restart	27
Automatic Step Restart	28
How Data Sets Being Extended Are Handled during Automatic Step Restart	29
JCL Requirements and Restrictions on Automatic Restart	30
Resource Variations Allowed in Automatic Restart	30
How the System Works at Automatic Restart	30
Deferred Restart	31
Message Sequence	31
Operator Considerations	32
Deferred Step Restart	32
Deferred Checkpoint/Restart	34
How the System Works during a Deferred Checkpoint/Restart.	38
Coding the RD and RESTART Parameters	38
RD Parameter	38
RESTART Parameter	39
SYSCHK DD Statement	40

Chapter 4. Checkpoint/Restart Processing	43
---	-----------

Notes for User Data Sets	43
BDAM Data Sets	43
Extended Format Data Sets	44
Large Format Data Sets	44
Generation Data Sets	44
Graphics Data Sets	45
Multivolume Data Sets	45
Partitioned Data Sets (PDSs).	45
Partitioned Data Sets Extended (PDSEs).	46
SMS-Managed Sequential Data Sets	47
Card Readers.	47
SYSABEND or SYSUDUMP Data Sets	47
SYSIN Data Sets	48
SYSOUT Data Sets	48
TCAM Data Sets.	48
Temporary Data Sets	49
VSAM Data Sets.	49
Work Data Sets	52
Notes for Tape Labels and Files.	52
VSE Tape Files	52

ISO/ANSI Tape Labels	52
Nonstandard Labels	52
Bypass Label Processing (BLP)	52
Automated Tape Library	52
Repositioning and Preserving the Contents of Data Sets	54
Recording Data Set Positioning	54
Buffering of Data Set Records	54
Preserving Data Set Contents	55
Allocating Devices during Deferred Checkpoint/Restart.	55
Handling Checkpoint/Restart Errors	56
I/O errors	56
Internal Errors	56
User Errors	56
Using Checkpoint/Restart with Other Programs	57
Job and Job Step Accounting	57
JES3	58
COBOL	59
DFSORT	59
PL/I.	59
Virtual Fetch	59
Appendix A. Checkpoint/Restart Codes	61
Return Codes.	61

Reason Codes	62
Modules and Return Codes for IHJ Messages	86

Appendix B. End-of-Volume Checkpoint Routine 89

Appendix C. Accessibility 91

Accessibility features	91
Using assistive technologies	91
Keyboard navigation of the user interface	91
Dotted decimal syntax diagrams	91

Notices 95

Policy for unsupported hardware	96
Minimum supported hardware	97

Trademarks. 99

Glossary 101

Index 109

Figures

1. Using One Sequential Checkpoint Data Set to Ensure Restart	13	5. Canceling a Request for Automatic Restart	20
2. Using Two Sequential Checkpoint Data Sets to Ensure Restart	14	6. Requesting a Resource after Restart	21
3. Recording a Checkpoint Identification Assigned by the Control Program	15	7. Requesting Automatic Checkpoint/Restart	28
4. Obtaining Updated TCB Information after Restart	18	8. Requesting Automatic Step Restart.	28
		9. Requesting a Deferred Step Restart.	33
		10. Requesting a Deferred Checkpoint/Restart	35

Tables

1.	Chapter 1. Introduction	1	6.	Chapter 4. Checkpoint/Restart Processing	43
2.	Chapter 2. Coding a Checkpoint	5	7.	Appendix A. Checkpoint/Restart Codes	61
3.	Determining the Size of a Checkpoint Data Set	10	8.	Appendix B. End-of-Volume Checkpoint Routine	89
4.	Cross-Memory Support Restrictions	23			
5.	Chapter 3. Requesting Restart	25			

About This Book

This book is intended to help you understand the checkpoint/restart facility provided by DFSMS. The checkpoint/restart facility lets you:

- Record information about a job at designated checkpoints.
- Restart a job at the beginning of a step or at a checkpoint in a step.

Required product knowledge

To use this book effectively, you should be familiar with job control language (JCL) and data management.

z/OS information

This information explains how z/OS references information in other documents and on the web.

When possible, this information uses cross document links that go directly to the topic in reference using shortened versions of the document title. For complete titles and order numbers of the documents for all products that are part of z/OS[®], see *z/OS Information Roadmap*.

To find the complete z/OS library, including the z/OS Information Center, see z/OS Internet Library (<http://www.ibm.com/systems/z/os/zos/bkserv/>).

How to send your comments to IBM

We appreciate your input on this publication. Feel free to comment on the clarity, accuracy, and completeness of the information or provide any other feedback that you have.

Use one of the following methods to send your comments:

1. Send an email to mhvrcfs@us.ibm.com.
2. Send an email from the "Contact us" web page for z/OS (<http://www.ibm.com/systems/z/os/zos/webqs.html>).
3. Mail the comments to the following address:
IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Building 707
2455 South Road
Poughkeepsie, NY 12601-5400
US
4. Fax the comments to us, as follows:
From the United States and Canada: 1+845+432-9405
From all other countries: Your international access code +1+845+432-9405

Include the following information:

- Your name and address.
- Your email address.
- Your telephone or fax number.
- The publication title and order number:
z/OS V2R1.0 DFSMSdfp Checkpoint/Restart
SC23-6862-00
- The topic and page number that is related to your comment.
- The text of your comment.

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

IBM or any other organizations use the personal information that you supply to contact you only about the issues that you submit.

If you have a technical problem

Do not use the feedback methods that are listed for sending comments. Instead, take one of the following actions:

- Contact your IBM service representative.
- Call IBM technical support.
- Visit the IBM Support Portal at z/OS support page (<http://www.ibm.com/systems/z/support/>).

z/OS Version 2 Release 1 summary of changes

See the following publications for all enhancements to z/OS Version 2 Release 1 (V2R1):

- *z/OS Migration*
- *z/OS Planning for Installation*
- *z/OS Summary of Message and Interface Changes*
- *z/OS Introduction and Release Guide*

Chapter 1. Introduction

Table 1 contains directive information about this chapter's contents:

Table 1. Chapter 1. Introduction

Section
"Components"
"Dependencies" on page 2

Checkpoint/restart is a method of recording information about a job at programmer-designated checkpoints so that the job can be restarted at one of these checkpoints or at the beginning of a job step.

A checkpoint is taken when your program or the system end-of-volume function issues the CHKPT macro instruction. This macro causes the contents of the program's virtual storage area and certain system control information to be written as a series of records in a data set. The restart function retrieves records from the data set if the job terminates abnormally or produces erroneous output, and the job can be restarted. Restart can take place immediately (allowed by the operator at the console) or be deferred until the job is resubmitted. In either case, you can avoid the time-consuming process of rerunning the entire job from the beginning.

Components

Checkpoint/restart consists of the following components.

CHKPT Macro Instruction

Your program or end-of-volume optionally issues the CHKPT macro.

When a CHKPT macro is executed, the contents of the program's virtual storage data area and certain system control information are written in a data set as a series of records. The series of records is called a checkpoint entry, and the data set in which they are written is called a checkpoint data set. The checkpoint entry, which has a unique programmer-specified or system-generated identification called a *checkid*, is retrieved from the data set when restart occurs.

For a detailed explanation about establishing a checkpoint see Chapter 2, "Coding a Checkpoint," on page 5.

End-of-Volume Exit Routine

You can code an end-of-volume exit routine that can issue a CHKPT macro instruction each time the processing of a multivolume physical sequential user data set is continued on another volume. Appendix B, "End-of-Volume Checkpoint Routine," on page 89 contains a brief description of this routine. For more detailed information about coding the EOVS exit for physical sequential data sets see *z/OS DFSMS Using Data Sets*.

Checkpoint at End-of-Volume Facility

A system-supplied routine takes checkpoints at end-of-volume occurrences for multivolume data sets opened with QSAM, BSAM, or EXCP and can be invoked

with a JCL parameter. “Checkpoint at End-of-Volume” on page 6 contains a brief description of this routine. See *z/OS DFSMS Using Data Sets*.

Restart Definition (RD) Parameter

When the RD parameter is used, it is coded in the JOB or EXEC statements and is used to request automatic step restart if job failure occurs, or to suppress, partially or totally, the action of the CHKPT macro instruction. For more information about this parameter see Chapter 3, “Requesting Restart,” on page 25.

RESTART Parameter

The RESTART parameter, coded in the JOB statement, is required when a job is resubmitted for deferred restart. It specifies the step (for deferred step restart) or the step and the checkpoint within that step (for deferred checkpoint/restart) where restart should begin. For more information about this parameter see Chapter 3, “Requesting Restart,” on page 25.

SYSCHK DD Statement

The SYSCHK DD statement is required when a job is resubmitted for deferred checkpoint/restart. It defines the checkpoint data set for the job being restarted. For more information about this statement, see “SYSCHK DD Statement” on page 40.

SYSCKEOV DD Statement

The SYSCKEOV DD statement defines the checkpoint data set that is to contain the checkpoint records generated when a checkpoint is taken at end-of-volume. For more information about this statement see “SYSCKEOV DD Statement” on page 7.

Dependencies

When using checkpoint/restart, consider the following dependencies.

Controlled Use of Checkpoints

Your installation may have special procedures in place to control the use of checkpoints on DASD. Contact your system administrator to determine whether this is the case for your system. For more information about protecting checkpoint data sets on shared DASD see *z/OS DFSMS Using Data Sets*.

Specification of Codes Eligible for Restart

The CKPTREST macro is no longer used to define codes eligible for restart. Instead, use the following technique.

During system installation, your system programmer uses a standard IBM[®] supplied list of system completion codes (codes produced when the system executes abend) to create the table of codes eligible for automatic restart. To change this table, the SYS1.PARMLIB member SCHEDxx RESTART/NORESTART statement type is specified. The RESTART statement type is used to add user-defined completion codes to the table of codes eligible for automatic restart. The NORESTART statement type is used to delete both IBM-supplied and user-defined codes from the table.

Job Journal Requirements

The job journal is a sequential data set that resides on the spool volume of the job entry subsystem. It contains a set of selected job-related control blocks that are critical to automatic restart processing.

The job journal is necessary because the scheduler control blocks are maintained in the scheduler work area (SWA) in pageable storage, rather than in a job queue on external storage. When a job or the system itself fails, the space that contains the SWA and its job control blocks is lost. Because it preserves up-to-date copies of certain critical control blocks, the job journal makes it possible to reconstruct the SWA. SWA control blocks are restored to the state they were in just prior to the failing step for automatic step restart. For automatic checkpoint/restart, SWA control blocks are reconstructed as they appeared at the time of the most recently issued CHKPT.

Job journaling is required before a job is eligible for an automatic restart.

For JES2: Job journaling is provided to a job in JES2 in one of the following ways:

- JOURNAL is specified on the JOB CLASS CHARACTERISTICS statement in the SYS1.PARMLIB member used for initializing JES2 (see *z/OS JES2 Initialization and Tuning Guide*).
- The JCL for the job specifies RD parameter on the JOB statement or the EXEC statement (see *z/OS MVS JCL Reference*).

For JES3: Job journaling is provided to a job in JES3 in one of three ways:

- JOURNAL=YES is specified on the CLASS initialization statement. For more information about the CLASS statement see *z/OS JES3 Initialization and Tuning Guide*.
- JOURNAL=YES is specified on the MAIN JCL statement that overrides the CLASS initialization statement. For more information about the MAIN statement see *z/OS MVS JCL Reference*.
- The JCL has specified either the RESTART or the RD parameter on the JOB statement or the RD parameter on the EXEC statement.

The system creates a job journal to hold restart information for any job specifying job journaling. After a system failure and the system restart of the failing main processor, the jobs in execution that requested job journaling are restarted (warm started) by the system. If a job is eligible for automatic restart, the operator is sent the message IEF225D asking if the job should be restarted. If the job is not eligible for restart or if the operator indicates that the restart should not be attempted, any scratch or VIO data sets the job allocated are deleted, and the job terminates. Jobs that frequently use scratch or VIO data sets should request job journaling.

User Data Set Security

In order to restart a job step that processes RACF[®] protected data sets, you must be authorized to access all RACF[®] protected data sets that were open when the checkpoint was taken. In order to restart, all password-protected VSAM data sets that were open at checkpoint time must have the pointer to their password in the access method control block (ACB) when the checkpoint is taken, or the password must be known to the operator. For RACF[®] protected data sets, the userid specified on the RESTART job statement must be authorized to access the data set.

Checkpoint Data Set Security

An unauthorized program is defined as one that is not authorized by the authorized program facility (APF), not in supervisor state and does not have a system key (keys 0 through 7). An unauthorized program cannot communicate directly with the checkpoint data set. The unauthorized program can take checkpoints and do restarts, but after the operation has begun, a security interface is used by the system to prevent unauthorized alteration of the checkpoint data set.

If an unauthorized program tries to access a checkpoint data set directly, an error message appears, and the job terminates. In addition, the unauthorized program cannot take a checkpoint on a new checkpoint data set if another data control block (DCB) is already open to that data set.

Offline security of the checkpoint data set is ensured by the operator.

Chapter 2. Coding a Checkpoint

The CHKPT macro instruction is coded in your program. When the CHKPT macro executes, job step information (about your program, virtual storage data areas, data set position, and supervisor control) is written as a checkpoint entry in a checkpoint data set. The point at which this information is saved becomes a checkpoint from which a restart may be performed. After the checkpoint entry is written, control returns to your program at the instruction following the CHKPT macro.

This chapter defines a checkpoint and explains how to establish checkpoints to restart job steps. Table 2 contains directive information about this chapter's contents:

Table 2. Chapter 2. Coding a Checkpoint

Section
"When to Code a Checkpoint"
"Coding the DCB for a Checkpoint Data Set" on page 7
"Coding the DD Statement for a Checkpoint Data Set" on page 8
"Using Checkpoint Data Sets" on page 10
"Coding the CHKPT Macro Instruction" on page 15
"Canceling a Checkpoint" on page 20
"Serially Reusable Resources" on page 20
"Special Operating System Features" on page 22

When to Code a Checkpoint

Using the CHKPT macro instruction, you can code a checkpoint within your program or at end-of-volume. The checkpoint routine records information about certain system information, user storage (including user programs in user storage), and the data sets used by the step executing the CHKPT macro. Recorded information about data sets includes the following:

- For all data sets, the checkpoint routine records the information which can be coded on a DD statement, such as the device type and volume serial numbers. The contents of the step's job file control blocks (JFCBs) and job file control block extensions (JFCBEs) are also recorded, as well as the contents of the generation data group name tables (GDGNT).
- For data sets which are (1) open at the checkpoint, (2) are processed on either magnetic tape or direct access devices, (3) and use the BSAM, QSAM, BPAM, VSAM, or EXCP access methods, the checkpoint routine records the information needed to reposition the data set if restart occurs at a checkpoint.

If a checkpoint is taken and an output data set is extended onto a second direct access volume because end-of-volume occurred on the first volume and there was no more space available on the volume, or because the data set contained 16 extents, restart subsequently occurs at that checkpoint and the system does not delete the extension of the data set.

Note that, when a step using the universal character set (UCS) feature is restarted, the system does not determine if the UCS buffer or forms control buffer (FCB) buffer is properly loaded, nor does it alert the operator to the UCS or FCB requirements of the step.

Checkpoint at End-of-Volume

The checkpoint at end-of-volume provides an external checkpoint function with no user program modifications. It is executed immediately prior to the invocation of any checkpoint in a user EOV exit routine. Redundancy occurs if a user exit routine is supplied and checkpoint at end-of-volume is invoked for the same data set's end-of-volume occurrence.

Checkpoint at end-of-volume is executable only if the `CHKPT=EOV` parameter is specified for multivolume data sets or for the second, third, etc., data sets of a sequential concatenation. If the first data set of a concatenation is a multivolume data set, this parameter is also valid on that DD statement.

Checkpoint at end-of-volume issues a `CHKPT` macro. If, at EOV, a return code of 08 is presented, the `CHKPT` macro is executed again. Return code 08 indicates that end-of-extent or end-of-volume occurred for the `SYSCKEOV` DD before completion of the checkpoint entry. The retry of return code 08 uses secondary space, if provided.

For the other unsuccessful return codes (0C, 10, 14, or 18) and for unsuccessful retry of return code 08, the following message is generated:

```
IEC067I CHKPT=EOV FACILITY EXECUTED UNSUCCESSFULLY
```

This message is preceded by a checkpoint/restart error message (prefixed "IHJ") that describes the nature of the problem. For more details about these error messages see *z/OS MVS System Messages, Vol 1 (ABA-AOM)*, *z/OS MVS System Messages, Vol 2 (ARC-ASA)*, *z/OS MVS System Messages, Vol 3 (ASB-BPX)*, *z/OS MVS System Messages, Vol 4 (CBD-DMO)*, *z/OS MVS System Messages, Vol 5 (EDG-GFS)*, *z/OS MVS System Messages, Vol 6 (GOS-IEA)*, *z/OS MVS System Messages, Vol 7 (IEB-IEE)*, *z/OS MVS System Messages, Vol 8 (IEF-IGD)*, *z/OS MVS System Messages, Vol 9 (IGF-IWM)*, and *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

In any event, processing continues, and this message is generated for each unsuccessful invocation of the checkpoint at end-of-volume until step termination occurs. Operator intervention is required to halt further processing prior to step end.

CHKPT at EOV JCL Parameter

The `CHKPT=EOV` parameter on a DD statement requests that a checkpoint be taken for this job step at EOV for the data set whose DD has this parameter. The following restrictions apply to the use of this JCL parameter:

1. The DD must define a QSAM or BSAM sequentially processed data set.
2. The QSAM or BSAM sequentially processed data set must be a multivolume data set or the second, third, etc., set of a concatenated data sets.
3. The DD statement must not define a `SYSOUT` data set. It can be a `DD *`, `SUBSYS`, or `DD DATA` type data set in a concatenation if it is after the first DD.
4. The JCL parameters `DDNAME` and `DYNAM` cannot be specified on the same DD statement with this parameter.
5. The DD must not define a checkpoint data set.

If restriction 1, 2, or 5 is not observed, the following will result:

- No action, no checkpoints are taken
- Processing continues as if the CHKPT=EOV parameter were not specified.

If restriction 3 or 4 is not observed, the following will result:

- JCL error messages appear
- Processing is not initiated.

An example follows:

```
//DD1 DD DSN=DSN1,DISP=OLD,VOL=SER=(TAPE01,TAPE02),  
//      UNIT=TAPE,CHKPT=EOV  
  
//DD2 DD DSN=DSN2,DISP=OLD  
//      DD DSN=DSNX,DISP=OLD,CHKPT=EOV  
//      DD DSN=DSNY,DISP=OLD,CHKPT=EOV  
  
//DD3 DD DSN=DSN3,DISP=NEW,VOL=(,,5),  
//      UNIT=DISK,SPACE=(CYL,(300,300)),CHKPT=EOV
```

- DD1 — Multivolume data set
- DD2 — Concatenated data sets
- DD3 — Multivolume data set on disk

SYSCKEOV DD Statement

The SYSCKEOV DD statement defines the checkpoint data set to contain the checkpoint records generated from the checkpoint at end-of-volume. The same restrictions that apply to other checkpoint DD statements (see “Coding the DD Statement for a Checkpoint Data Set” on page 8) also apply to this DD statement, with the following exceptions:

- DISP=MOD is recommended to reduce risk of loss of checkpoint data in the event of a system failure while taking a checkpoint. This is because restart normally is from the most recent checkpoint in the data set. If you do not code DISP=MOD, each checkpoint overlays the previous one. See “How Checkpoint Entries Are Written” on page 11 and “How to Ensure Restart with Sequential Checkpoint Data Sets” on page 13.
- This DD must define a sequential BSAM data set (BPAM is not supported)
- All the DCB parameters are provided by the checkpoint at EOV routine and should not be coded on the DD statement.

An example follows:

```
//SYSCKEOV DD DSN=CKPTDS,UNIT=TAPE,DISP=MOD
```

Checkpoint in Exit Routines

The CHKPT macro instruction must not be used in an exit routine other than the end-of-volume exit routine. You can take a checkpoint when a BSAM or QSAM data set reaches end-of-volume. For more detailed information about coding the EOV exit for physical sequential data sets, see Appendix B, “End-of-Volume Checkpoint Routine,” on page 89.

Coding the DCB for a Checkpoint Data Set

You must provide either a data control block (DCB) for the checkpoint data set or the name of the DD statement for the checkpoint data set.

Required DCB Parameters

If you provide a DCB, you must code the following parameters:

- DSORG=PS or PO (BSAM or BPAM data set organization)
- MACRF=W (WRITE macro instruction)
- DDNAME= name of DD statement for checkpoint data set.

To control space allocation you can provide DSORG also on the DD statement or in the data class definition. Optional DCB macro and DD statement parameters are in “DCB Options”

DCB Options

You may, optionally, provide the following DCB parameters:

- OPTCD=W (write validity checking)
- NCP=2 (number of channel programs)
- NCP=2 and OPTCD=C (chained scheduling)
- RECFM=U or UT (undefined record format)
- DEVD=DA or TA (direct access or tape device)
- TRTCH=C (data conversion with odd parity; parameter required).

z/OS DFSMS Macro Instructions for Data Sets contains detailed information about coding DCBs.

Notes on the DCB for a Checkpoint Data Set

The following restrictions apply to the DCB for a checkpoint data set:

- The checkpoint routine writes all checkpoint records in both 400- and 4096-byte blocks.
- Requests for two channel programs or chained scheduling apply to the writing of all checkpoint records. Such requests do not apply to the reading of checkpoint records for a restart.
- OPTCD=Q cannot be specified in the DCB, because checkpoints are not character data.
- Following errors that leave the DCB invalid (for example, a D37 abend), the DCB must be closed and opened again, if additional checkpoints will be taken using this DCB.
- Do not code the DCBE parameter on the DCB macro.

Coding the DD Statement for a Checkpoint Data Set

The DD statement for the checkpoint data set must define the data set in the usual way. The publication *z/OS MVS JCL Reference* contains detailed information about coding the DD statement.

“Coding the DCB for a Checkpoint Data Set” on page 7 describes certain DCB parameters that you can code on the DD statement.

The only restrictions on the DD statement are:

- A checkpoint data set cannot be concatenated, because it is an output data set.
- The UNIT parameter must specify a tape or direct access device supported by BSAM or BPAM. The device can be specified by referring to a specific device, a device type, or a group of devices.

If direct access is specified, the device cannot be shared with another processor, unless your installation has special procedures in place to control the use of checkpoints on shared DASD. See *z/OS DFSMS Using Data Sets* for specific details. To avoid allocating to a shared DASD, a special generic device name should be generated at system definition time to include a nonshared DASD of a single device type.

- DEFER should not be coded in the DD statement.
- DISP=SHR cannot be specified.
- Secondary space allocation may be requested by the increment subparameter (see “Notes on the DD Statement”). The increment subparameter must specify at least enough space to contain two complete checkpoint entries. For information about calculating the length of the checkpoint entry see “Storage Estimates for Checkpoint Data Sets” on page 10.
- The LABEL parameter of the DD statement describes the labels of a data set on magnetic tape. For a checkpoint data set, you must specify only IBM standard labels (SL or SUL). Nonstandard labels (NSL), no labels (NL), or International Organization for Standardization/American National Standards Institute (ISO/ANSI) labels (AL or AUL) cannot be specified for a checkpoint data set. If the label type is not specified, the operating system assumes that the data set has IBM standard labels (SL).
- OPTCD=Q cannot be specified as a DCB subparameter.
- CHKPT=EOV is ignored (checkpoint resets the value in the JFCB).
- The RLSE subparameter of the SPACE parameter cannot be used on the DD statement for a checkpoint data set.
- A checkpoint data set cannot be an extended format data set, a PDSE file, or an HFS file.

Notes on the DD Statement

- The initial disposition of the data set (as specified in the DISP operand of the DD statement) is used to position the checkpoint data set each time it is opened. For a discussion of a specific circumstance in which the only checkpoint entry that exists on a data set will be the last checkpoint taken (whether that entry is valid or not), see “How Checkpoint Entries Are Written” on page 11.
- The final and conditional dispositions of the data set have their normal meanings. However, if termination is occurring and an automatic restart at a checkpoint is to occur, the system automatically keeps all data sets that are in use by the job, including the checkpoint data set.
- If you code UNIT or VOL=SER on the checkpoint DD statement, it has no effect unless changed by your installation's ACS routines. If the checkpoint data sets are to be restricted to specific volumes, a storage group should be defined for this purpose.
- Filling the current DASD extent or tape volume is called an end-of-volume (EOV) condition. If EOV is encountered while writing a checkpoint on a direct access or tape volume, a second attempt is made to create the checkpoint entry. If EOV occurs before the second attempt at writing the entry is completed, message IHJ0001 (invalid checkpoint) is issued. Control is returned to you with a X'08' return code in register 15 and a reason code of 027 in register 0.

The following example shows DD statements for the checkpoint data set:

```
//ddname DD DSN=dsname,UNIT=TAPE,DISP=(MOD,KEEP)
//ddname DD DSN=dsname,UNIT=SYSDA,
// DISP=(NEW,DELETE,KEEP),SPACE=(CYL,(15,17)),
// VOL=SER=CKPTDS
```

Using Checkpoint Data Sets

The checkpoint data set must be empty at the beginning of a job step. For more information, see "How Checkpoint Entries Are Written" on page 11.

Storage Estimates for Checkpoint Data Sets

The following information can be used as a guideline in determining the size of this data set.

Table 3 contains the size and number of records written when a checkpoint is taken. The number of tracks or the amount of tape occupied by the checkpoint data set can be determined by applying the number of records and their sizes against either the track capacities of the direct access device or the recording density and type for the magnetic tape device.

Table 3. Determining the Size of a Checkpoint Data Set

Description of Record	Size (in bytes)	Number Required
Checkpoint header record (CHR)	400	1
DD name table (DDNT)	400	$((D+20)/21)+1)/2$
Data set descriptor record (DSDR)	400	$(A+1)/2$
Problem program image record (PPIR)	4096	B
Supervisor record (SUR)	4096	C
Subsystem checkpoint record (SSCR)		
subsystem data sets	4096	Number of SYSIN or SYSOUT plus number of DDs with SUBSYS=
SPIE	20/4096	E/20
VSAM	4096	1

Records Code:

A = the number of data sets defined in the job step

$B = (\text{Storage allocated} - \text{freespace} + (20 * \text{number of gaps in allocated storage}))/4076$.

Example of calculation of B:

Storage allocated = 9000 to 4FFFFF and 1A000000 to 2FFFFFFF

Freespace=

Address	Size	Number of Gaps
A200 to A3FF	200	1
32900 to 35FFF	3700	1
40000 to 47FFF	8000	1
50000 to 19FFFFFFF	0	1 (not allocated; counts as a gap)

Table 3. Determining the Size of a Checkpoint Data Set (continued)

Description of Record	Size (in bytes)	Number Required
1F900000 to 210F8FFF	17F9000	1
28700000 to 29FFFFFF	1900000	1
=	3104900	6

Therefore:

$$B = ((47000 + 16000000) - 3104900 + (X'14' * 6))$$

/X'FEC'

$$B = 130C0$$

Note: All values in the above example are in hexadecimal.

- C =** A variable number of records for system and protected data management control blocks. This will be approximately three records plus one record for every three DASD data sets and one record for every ten non-DASD data sets.
- D =** Number of JCL-specified data sets dynamically unallocated at checkpoint.
- E =** One for every 90 ESPIEs established before checkpoint.

Checkpoint/restart's dynamic storage requirement varies with the resources used by the caller. The largest impact on the dynamic storage requirement results from the amount of storage allocated to the job and the amount of storage fragmentation. Checkpoint/restart's dynamic storage is allocated out of subpools 229 and 230.

Fragmented Storage

User fragmented storage may result in:

- Reduced performance of a checkpoint
- Insufficient storage for a checkpoint to contain a description of the user's storage (checkpoint is not completed).

How Checkpoint Entries Are Written

Unopened DCB Passed to Checkpoint

- Sequential data sets
 - When a sequential data set's disposition is NEW or OLD, the checkpoint entry is written at the beginning of the data set, even if a checkpoint entry was written on a previous call.
 - When a sequential data set's disposition is MOD, the checkpoint entry is written after the last entry existing in the data set.
- PDSs
 - When the data set is partitioned, each checkpoint entry is a member, and its checkid is its member name. After it writes a checkpoint entry, the checkpoint routine executes the STOW macro to add the checkid of the entry to the directory of the data set.

- If the data set is partitioned, regardless of its disposition, the checkpoint entry is written after the last entry existing in the data set.
- If an identical checkid already exists in the directory, the related address of a member is changed and becomes the address of the new checkpoint entry. The initial disposition specified for the checkpoint data set has no effect on the STOW operation.

Opened DCB Passed to Checkpoint

If your program has a DCB and opens the checkpoint data set for output, the checkpoint routine writes a checkpoint entry at the data set's current position and does not close the data set. A user-opened checkpoint data set does not have to be closed after taking the last checkpoint for the job step. All the checkpoint entries are saved in this case, without specifying DISP=MOD. This provides the ability to request a deferred restart from any of the checkpoints.

- Sequential data sets

For sequential data sets, DISP=OLD, NEW, or MOD has no effect on how the checkpoint entries are written on the checkpoint data set. The first time each checkpoint data set is used, the checkpoint entry is written at the beginning of the file. Each following checkpoint entry is written after the previous entry.

- PDSs

PDSs are treated in the same manner as unopened DCBs passed to checkpoint.

The DCB is left open when checkpoint returns to the user. The open checkpoint data set is closed by step termination.

DDNAME Passed to Checkpoint

If your program provides a DDNAME instead of a DCB, the checkpoint routine opens the checkpoint data set the first time it receives that DDNAME.

- Sequential data sets

For a sequential checkpoint data set, the data set's disposition controls where the first checkpoint is written. If the data set's disposition is NEW or OLD, the first checkpoint entry is written at the beginning of the data set, even if a previous checkpoint call has written an entry. If the data set's disposition is MOD, the first checkpoint entry is written after the last entry existing in the data set. DISP=OLD, NEW, or MOD has no effect on where checkpoint entries other than the first are written. Each following entry is written after the previous entry.

- PDSs

PDSs are treated in the same manner as nonopened DCBs passed to checkpoint.

Note: If your program uses more than one DDNAME on executions of the CHKPT macro, the data set for the previous DDNAME is closed and the data set for the new DDNAME is opened.

General Considerations

If end-of-volume is encountered during writing of a checkpoint entry, a second attempt is made to create the checkpoint entry on another volume if another volume has been allocated. If the checkpoint data set extends during writing of a checkpoint entry, a second attempt is made to create the checkpoint entry by writing over the first attempt and using the new extents. If EOV occurs or the data set extends before the second attempt at writing the entry is completed, message IHJ0001 (invalid checkpoint) is issued. Control is returned to you with a X'08' return code in register 15 and a reason code of 27 in register 0.

The status (opened or closed) and position of a checkpoint data set remain the same at restart as they were after execution of the CHKPT macro instruction that established the checkpoint.

The first time a checkpoint data set is opened, checkpoint routines cause the OPEN macro to request security information for the data set from the operator.

Note: A checkpoint data set must contain only checkpoint entries. If a program attempts to write its own data in a checkpoint data set, message IEC954I will be issued, followed by an abend code of 23F.

How to Ensure Restart with Sequential Checkpoint Data Sets

Abnormal termination or system failure may occur while the new entry is written. To ensure that restart at the most recent valid checkpoint is possible, a new checkpoint entry must not be written over a preceding one. The figures below suggest methods by which you can ensure that restart is possible. All of the methods use sequential checkpoint data sets.

Note: Because of the characteristics of PDSs, checkpoint entries with different member names are not written over each other.

Figure 1 shows the use of one sequential checkpoint data set, one data control block, and one DD statement (CHECKDD) specifying MOD disposition. The user allows the checkpoint routine to open and close the data set each time it writes a checkpoint entry. Checkpoint entries are written sequentially in the data set.

```
Program
.
.
.
CHKPT      CHKDCB
.
.
.
CHKDCB     DCB  DDNAME=CHECKDD,MACRF=W,DSORG=PS
.
.
.
DD Statement

//CHECKDD DD  UNIT=TAPE,DISP=(MOD,KEEP),DSN=WAREHOUS.DB1.CHK
```

Figure 1. Using One Sequential Checkpoint Data Set to Ensure Restart

Alternatively, you can have your program open the data set and leave it open. The disposition of the data set may be NEW, OLD, or MOD. This method uses one sequential checkpoint data set, one DCB, and one DD statement.

Figure 2 on page 14 shows a way to alternate data sets when all checkpoints are taken by one CHKPT macro instruction. The data sets are opened by the control program and identified by two DD statements, CHECKDD1 and CHECKDD2. The data control block (DCB) initially refers to CHECKDD1. Before the second checkpoint, the DCB is changed to refer to CHECKDD2; before the third checkpoint, the DCB is again changed to refer to CHECKDD1; and so forth. One data control block can be used for two data sets that are not open at the same time.

An alternative method of using two sequential data sets is to use two DCBs and two DD statements specifying NEW or OLD dispositions, and to execute alternately two CHKPT macro instructions. Each CHKPT macro instruction then refers to a different data set.

The method illustrated in Figure 1 on page 13 saves all checkpoint entries for possible use in deferred restart. The method illustrated in Figure 2 conserves auxiliary storage. None of the methods requires a particular device type.

```

Program

    LA    2,CHECKDCB          Establish CHECKDCB as
*                               base address for
    USING IHADCB,2           IHADCB
    XC    DCBDDNAM(8),DDHOLD Exchange ddname in
    XC    DDHOLD(8),DCBDDNAM CHECKDCB for ddname
    XC    DCBDDNAM(8),DDHOLD in DDHOLD
    CHKPT CHECKDCB          Open, checkpoint, close
    .
    .
    .
DDHOLD  DC    C'CHECKDD1'
CHECKDCB DCB   DSORG=PS,MACRF=(W),DDNAME=CHECKDD2
        DCBD  DSORG=PS      Define IHADCB (dummy
*                               section that defines
*                               DCBDDNAM)

```

DD Statements

```

//CHECKDD1 DD UNIT=SYSDA,DISP=NEW ...
//CHECKDD2 DD UNIT=SYSDA,DISP=NEW ...

```

Figure 2. Using Two Sequential Checkpoint Data Sets to Ensure Restart

How Checkpoint Entries Are Identified

The number of checkpoint entries, as well as the number of checkpoint data sets used concurrently, is unlimited. In a sequential checkpoint data set, checkids of valid or invalid checkpoint entries in one data set should be unique. In a PDS, checkids of valid entries should be unique.

If you specify checkids instead of having the system generate them, incorrect duplicates might be specified. The system cannot recognize the error of duplicate checkids. When deferred restart at a checkpoint occurs and the checkpoint data set is sequential, the system searches the data set from its beginning for the specified checkpoint entry. The system uses the first entry that it finds which has the specified checkid.

If the data set is partitioned, the system searches the directory to find the location of the specified checkpoint entry. If two or more entries having the same checkid were written in the data set, the most recent of those entries is the one pointed to by the directory, and restart occurs from the most recent entry.

Checkpoint entries have two identifications: primary and secondary.

- The primary identification is the programmer-generated or system-generated checkid specified or requested by the CHKPT macro instruction. The primary identification is used when a search is made for a checkpoint entry.

- The secondary identification is identical to the system-generated checkid that might have been requested by CHKPT. The secondary identification is then used as a base to compute the system-generated checkids of entries written after restart has occurred.

The control program identifies each checkpoint in a message to the operator and user. On request, the control program also makes the identification available to your program. In Figure 3, the CHKPT macro instruction requests the control program to supply an identification ('S' parameter) and place it in the 8-byte field named ID. When the checkpoint is successfully taken, the program prints the identification as part of a message to the programmer.

```

      .
      .
      .
CHKPT  CHKDCB, ID, 'S'    Take checkpoint
LTR    15,15            Was checkpoint taken
BNZ    PHASE2           No, branch to PHASE2
PUT    STEPLOG, MESSAGE Yes, print
*                               checkpoint ID
PHASE2
      .
      .
      .
MESSAGE DS    0F
        DC    H'44'          Record length
        DC    H'0'           Flags
        DC    C'SUCCESSFUL CHKPT AT PHASE2...ID='
ID      DS    CL8
STEPLOG DCB    DSORG=PS,MACRF=(PM),RECFM=V,LRECL=124,DDNAME=LOGDD
CHKDCB  DCB    DSORG=PS,MACRF=(W),RECFM=U,DDNAME=CHKDD

```

Figure 3. Recording a Checkpoint Identification Assigned by the Control Program

Coding the CHKPT Macro Instruction

The CHKPT macro instruction is coded in your program. When the CHKPT macro executes, job step information (about your program, virtual storage data areas, data set position, and supervisor control) is written as a checkpoint entry in a checkpoint data set. The point at which this information is saved becomes a checkpoint from which a restart may be performed. After the checkpoint entry is written, control returns to your program at the instruction following the CHKPT macro.

The CHKPT macro instruction optionally refers to the data control block (DCB) for the checkpoint data set. If the data set is not open, the checkpoint routine opens it and then closes it after writing the checkpoint entry. If the data set is open, the checkpoint routine writes the checkpoint entry, but does not close the data set.

The checkpoint data set must be on one or more magnetic tape volumes or direct access volumes. A checkpoint data set that resides on a magnetic tape must have IBM standard labels. Direct access volumes cannot be shared. For exceptions to this restriction on shared DASD see "Coding the DD Statement for a Checkpoint Data Set" on page 8.)

The standard form of the CHKPT macro instruction is:

[symbol]	CHKPT	{ <i>dcbaddr</i> DDNADDR= <i>ddnaddr</i> CANCEL} [, <i>checkid addr</i> IDADDR= <i>checkid addr</i>] [, <i>checkid length</i> IDLNG= <i>checkid length</i> 'S' IDLNG='S']] [,MF={S (S, <i>addr</i>)} {L (L, <i>name</i>)} {(E, <i>addr</i>) (E, <i>addr1</i> , <i>addr2</i>)}]
----------	-------	--

The options on the CHKPT macro are discussed in alphabetic order below.

CANCEL

Cancels the request for automatic checkpoint/restart. CANCEL can be used only with the standard form (MF= omitted or MF=S or MF=(S,NAME)). An automatic checkpoint/restart can be suppressed by issuing a CHKPT with CANCEL. See "Canceling a Checkpoint" on page 20.

checkid addr

Specifies the address of a programmer-provided field that is to contain a unique, printable identification of the checkpoint entry. This identification is called a checkid. The checkpoint routine writes the checkid as part of the entry and prints it in a message on the operator's console when it finishes writing the entry. The programmer must use the checkid and code it in the JOB statement RESTART parameter to use the corresponding entry to perform a deferred restart at a specific checkpoint. If the *checkid addr* operand is omitted, the *checkid length* or 'S' operand is invalid. Valid characters for a checkid are:

A - Z

0 - 9

\$ #

! *) ; - / , % _ > ? : ' = " and blanks are valid except as first characters or for PDSs.

With the execute form, *checkid addr* is an RX type address or register (2-12).

With the standard or list form, it is an A type address.

checkid length or 'S'

checkid length is the length in bytes of the field that contains the checkid. The maximum length of this field is 16 bytes when the checkpoint data set is physical sequential, 8 bytes when it is partitioned. By coding this operand or by omitting it entirely (in which case a length of 8 bytes is implied), you specify that your program will form an identification and store it in the checkid field before CHKPT is executed. If the *checkid addr* operand is omitted, the *checkid length* operand is invalid.

By coding this operand as 'S', you specify that the checkpoint routine is to generate an identification 8 bytes in length and store it in the checkid field. If the *checkid addr* operand is omitted, this operand is invalid.

If either *checkid length* or 'S' and *checkid addr* and the keyword equivalents are omitted, the checkpoint routine generates an identification. The checkpoint routine then writes this identification in the checkpoint entry and on the operator's console but does not return it to your program.

If you provide the checkpoint identification and the checkpoint data set is sequential, the identification can be any combination of up to 16 alphanumerics, special characters, and blanks. For a PDS, the checkpoint identification must be a valid member name of up to eight alphanumerics. The identification for each checkpoint should be unique. If two identifications differ only by having a different number of trailing blanks, the control program considers them to be

the same. For more information see “How Checkpoint Entries Are Identified” on page 14. When a deferred step restart takes place, the checkpoint identification is reset to 0.

One way to use the *checkid addr* operand, or its keyword equivalent, is to allow your program to select fields in the records of an input data set and use them as checkids. Alternatively, your program can use the *checkid addr* and the ‘S’ operands and include a system-generated checkid in the current record of an output data set.

dcbaddr

is the address of the DCB for the checkpoint data set. With the execute form, *dcbaddr* is an RX type address or register (2-12). With the standard or list form, it is an A type address.

DDNADDR=

Points to an 8-character, fixed-length field that has the name of the DD statement for the checkpoint data set. This keyword may be used instead of the first positional operand DCB. This keyword and the positional operand DCB are mutually exclusive.

With the execute form, *ddnaddr* is an RX type address or register (2-12). With the standard or list form, it is an A type address.

IDADDR=

Points to a variable-length field that contains, or receives, the checkid. (This keyword may be used instead of the second positional operand, *checkid addr*.) This keyword is mutually exclusive with and follows the same rules as the second positional operand.

With the execute form, *checkid addr* is an RX type address or register (2-12). With the standard or list form, it is an A type address.

IDLNG=

Specifies the length of the field that contains, or receives, the checkid. This keyword may be used instead of the third positional operand, *checkid length*. This keyword is mutually exclusive with and follows the same rules as the third positional operand.

Keyword MF= Standard Form (default)

MF=S

Generates an inline parameter list and executable code.

MF=(S,name)

Generates executable code and an inline parameter list with the label name at the beginning of the list. *name* is a symbol.

Keyword MF= List Form

MF=L

Generates a parameter list. *Symbol* is required with this form.

MF=(L,name)

Generates a parameter list with the label name at the beginning of the list. *name* is a symbol.

Keyword MF= Execute Form

The execute form is useful for programs that must be reentrant.

MF=(E,addr)

Generates executable code referencing to the parameter list whose name is *addr*. *addr* can also be expressed as a register, for example, MF=(E,(R5)), if the register is preloaded to point to a parameter list.

MF=(E,addr1,addr2)

Generates executable code that moves a model parameter list from *addr2* to *addr1* and updates the list in *addr1* with the positional or keyword operands specified. *addr1* and/or *addr2* are RX type addresses or may be expressed in register form, for example, MF=(E,(R5),(R6)) if the register or registers are preloaded to point to the respective parameter list.

For more information about the MF= keyword see *z/OS DFSMS Macro Instructions for Data Sets*.

CHKPT Return Codes

The CHKPT macro returns a code in register 15 to indicate whether the CHKPT macro instruction executed successfully. If the value in register 15 does not equal 0, register 0 gives a more detailed code. Appendix A, "Checkpoint/Restart Codes," on page 61, contains a list of these codes and their meanings.

Use of CHKPT with Other Macro Instructions

ATTACH

If ATTACH is issued in the program using CHKPT, all subtasks created must terminate before a CHKPT is issued; the job-step task must be the only task of the step.

EXTRACT

The EXTRACT macro obtains information from the task control block (TCB). TCB information may change when the task terminates and the job step restarts. If the information is needed after restart, the EXTRACT macro instruction should be reissued after the checkpoint is taken, as shown in Figure 4.

```

      .
      .
      .
*     EXTRACT   ANSADDR,FIELDS=(ALL)   Obtain TCB
                                         information
      .
      .
      .
*     CHKPT    CHKPTDCB                Establish
                                         checkpoint
*     CH       15,=H'4'                Is restart in
                                         progress
*     BNE      NRESTART                No, branch to
                                         NRESTART
*     EXTRACT   ANSADDR,FIELDS=(ALL)   Yes, obtain new
                                         information
*
NRESTART
      .
      .
      .
```

Figure 4. Obtaining Updated TCB Information after Restart

FEOV

A user cannot issue an FEOV macro for the checkpoint data set.

IMGLIB

A checkpoint cannot be created while a DCB created by the IMGLIB macro exists in the address space. You can delete the DCB by issuing IMGLIB CLOSE.

PCLINK

If a PCLINK STACK is issued, a PCLINK UNSTACK must also be issued against all existing stacks before a CHKPT can be issued. There must not be any PCLINK STACKS in existence when a CHKPT macro is issued.

SETDEV

If an IBM 3890 Document Processor unit record device is used, the following SETDEV macro instruction must be issued after a successful restart:

```
SETDEV dcbaddr,DEVT=3890,IREC=irecaddr [,PROGRAM=programe]
```

where:

dcbaddr

is the address of the associated DCB.

irecaddr

is the address of an initialization record.

programe

is the name of the stacker control instruction program loaded in SYS1.IMAGELIB. This parameter is optional.

SETPRT

The SETPRT macro in data management selects the universal character set (UCS) buffer for an IBM 3203-5, 3211, 4245, 4248, or 1403 Printer with the Universal Character Set feature and the forms control buffer (FCB) for the 3203-5 or 3211 Printers, 3262 Model 5, or the IBM 3800 Printing Subsystem.

For some printers, such as the 3800, the buffer contents are not saved when a checkpoint is taken. See *z/OS DFSMS Introduction* for a more complete list. To ensure that all lines are printed before taking a checkpoint, you should take one of the following actions:

- Close the printer with the CLOSE macro instruction
- Issue the “Clear Printer” command, perhaps with the SETPRT macro
- Reprint the data after restart.

For the IBM 3800 Printing Subsystem, (Model 1, Model 3, or Model 8 in compatibility mode), the SETPRT macro instruction initially sets or dynamically changes the control information for the printer. For more information about the SETPRT macro see *IBM 3800 Printing Subsystem Programmer's Guide* and *z/OS DFSMS Macro Instructions for Data Sets*.

For the IBM 3800 Printing Subsystem (Model 1, Model 3, or Model 8 in compatibility mode) at checkpoint/restart, the job has the job file control block extension (JFCBE) parameters as modified during the last JFCBE exit. If no exit was requested or if the JFCBE was not flagged as being modified during the exit, the JFCBE reflects the values coded in the restart JCL. For more information about the JFCBE exit see *z/OS DFSMS Using Data Sets*.

STIMER

A time interval established by the STIMER macro must be completed before a CHKPT is issued.

WTOR

The reply to a WTOR macro instruction must be received before a CHKPT is issued.

Canceling a Checkpoint

Under the following conditions, you may not wish to have a job automatically restarted in case of failure.

- If you have not yet taken a checkpoint and the job has already been restarted.
- If you are updating critical data that must be complete and accurate before an automatic restart can be successful.

In these cases, you can take a checkpoint, but avoid automatic restart, by using the CANCEL option on the CHKPT macro. If CHKPT CANCEL is used after a checkpoint is taken, the job is not automatically restarted after failure. For the format of the CHKPT macro see “Coding the CHKPT Macro Instruction” on page 15.

After being restarted, the job step may again terminate abnormally. If the job terminates abnormally, it may be restarted from the same checkpoint, subject to operator authorization. To avoid restarting the job step twice from the same checkpoint, code the sequence shown in Figure 5.

```
      .  
      .  
      .  
CHKPT  CHKPTDCB  Establish checkpoint  
CH      15,=H'4'  Is restart in progress  
BNE     NRESTART No, branch to NRESTART  
CHKPT   CANCEL   Yes, cancel restart request  
NRESTART  
      .  
      .  
      .
```

Figure 5. Canceling a Request for Automatic Restart

After the successful completion of a restart from a checkpoint, the system places a return code of X'04' in register 15 and returns control to your program at the instruction following the CHKPT macro instruction. At this time, a request for another automatic restart at the same checkpoint is normally in effect. In Figure 5, the instruction that follows the CHKPT macro instruction tests the return code to determine whether control has been returned as the result of a restart. If the return code is X'04', a restart has just occurred, and a second CHKPT macro instruction is executed. This macro instruction has a CANCEL operand, which cancels the existing request for an automatic restart. If the job step again terminates abnormally after a restart from the checkpoint, automatic restart can occur only at the beginning of the step. It will not occur at the checkpoint preceding the canceled checkpoint.

Serially Reusable Resources

When a job step terminates, it loses control of serially reusable resources. If the step is restarted, it must request all the resources needed to continue processing.

Requesting Serially Reusable Resources

Explicit use of a serially reusable resource is requested when the user's program issues the ENQ or RESERVE macro instruction. If the program issues the ENQ or RESERVE macro and takes a checkpoint, it must issue the ENQ or RESERVE macro again whenever restart occurs at the checkpoint.

Figure 6 shows a program that requests a serially reusable resource by issuing an ENQ before establishing a checkpoint. After the checkpoint, the program tests for a restart. If a restart has occurred, the program requests the same resource again. The program requests the resource again because the job step terminated, lost control of the resource, and then restarted from the checkpoint.

```

.
.
.
ENQ (QADDR,RADDR)
.
.
.
CHKPT CHKPTDCB
CH 15,=H'4'
BNE NRESTRT
ENQ (QADDR,RADDR)
NRESTRT
.
.
.
DEQ (QADDR,RADDR)
.
.
.

```

Figure 6. Requesting a Resource after Restart

Some serially reusable resources are requested implicitly by issuing data management macro instructions. These resources may be records that you are processing or tracks on a direct access device. To ensure correct processing, you must *not* establish checkpoints during control of these resources:

- If the basic direct access method (BDAM) is used, your program must, before executing the CHKPT macro, execute either the WRITE or the RELEX macro to release a record that is read with exclusive control.
- If BDAM is used to add a record to a data set with variable-length or undefined records, BDAM issues an ENQ macro instruction for the capacity record (R0). Your program must execute the WAIT or CHECK macro to check completion of the write operation before it executes CHKPT. After execution of the WAIT or CHECK macro, the resources are removed from the queue.
- If a VSAM cluster is implicitly specified, the restart program obtains the names of the data set and the index from the catalog and reissues an ENQ macro against each of them. Therefore, no special considerations are required.

Tape Volume DEQ at Demount

Tape volume DEQ at demount is not affected by restarts. If restart is from a checkpoint, all volumes that are dequeued during the processing of the job are enqueued when the restart job is initiated. This includes all volumes referenced in the JFCB at the time of checkpoint. Subsequent restart processing tests each tape data set open at the time of checkpoint to determine if:

- the data set was open for OUTPUT, OUTIN, EXTEND, OUTINX, or INOUT when the checkpoint was taken.
- JFCDQDSP was set to 1 in the JFCB at the time of checkpoint.
- the program restarting is APF authorized.

If these conditions are satisfied, the DEQ at demount is reestablished so that the data set's current volume is dequeued when it is demounted. All volume serials in the JFCB and JFCB extensions prior to the current volume are dequeued during the restart process for that data set.

Because all volumes, including those dequeued through the DEQ at demount, are enqueued at restart, if only for a short time, automatic checkpoint/restart may result in a "waiting for volumes" condition if any of the dequeued volumes are allocated to another job. A deferred checkpoint/restart can be timed to avoid such a delay.

Special Operating System Features

The following sections contain special operating system features.

VIO Data Sets

If a VIO data set is open when the checkpoint is taken, it is supported for automatic restart only.

If restart is deferred, VIO data sets for BSAM or QSAM must be redefined as dummy data sets. If they are not redefined or if an access method other than BSAM or QSAM accesses the VIO data set, the job fails.

Dynamic Allocation and Dynamically Concatenated Data Sets

Checkpoint/restart supports the use of dynamic allocation by program programs. There are, however, several restrictions with a checkpointed job involving the use of dynamically concatenated data sets. See *z/OS MVS Programming: Authorized Assembler Services Guide* for information about the specification of dynamically concatenated data sets.

If you code VOL=REF to refer to a nonspecific tape DD request, then after a restart the two DDs may have different volumes from each other. This occurs when the program writes the first data set before the checkpoint, unallocates that data set and does not open the second data set until after the restart.

If a VIO data set is dynamically allocated prior to a checkpoint and unallocated after checkpoint, the step is ineligible for restart until the next checkpoint is taken.

When the checkpointed job requests dynamic concatenation, the following restrictions apply:

- On a deferred or automatic checkpoint/restart, JCL-defined DD statements that were dynamically concatenated (by SVC 99) with the permanently concatenated attribute will not be properly resolved. This will result in one or more missing DD statements. If the restart attempt fails, message IEF007I, with reason code 34, will be issued.
- On an automatic checkpoint/restart of a job that uses dynamically allocated data sets that had the permanently concatenated attribute, the restart may produce incorrect results. For example, the DDs in the concatenated group may not be matched properly, causing concatenation of the wrong DDs. Restart processing may not detect this type of error. This type of job or step should specify RD=NR in its JOB or EXEC statement. If the job or step fails, it can be resubmitted for a deferred checkpoint/restart.

- On either a deferred or automatic checkpoint/restart of a job which uses dynamically concatenated data sets, unit affinity requests for the concatenated DDs may or may not be accepted on the restart. This is true with or without the permanently concatenated attribute.
- You cannot use the nocapture option of dynamic allocation.

Shared Resources

Checkpoint/restart is supported for the local shared resources (LSR) feature, including use with Hiperspace™ buffers, but repositioning is not allowed. Also, checkpoint is not allowed and cannot be taken when the global shared resources feature is used.

For automatic checkpoint/restart, deleting a data set open at the time the checkpoint was taken or deallocating a SYSOUT data set to the job entry subsystem makes the step ineligible for restart until the next checkpoint is taken.

For automatic step restart, any of the following makes the step ineligible for restart:

- KEEP, CATLG, or UNCATLG a NEW JCL-specified data set
- DELETE JCL-specified data set whose initial status upon entry to the step was OLD
- DELETE an OLD dynamically allocated data set that had volumes specified when allocated
- UNCATLG JCL-specified data set whose initial status upon entry to the step was OLD and did not have volumes specified.

Cross-Memory Support Restrictions

A checkpoint is not allowed when storage access across address space boundaries has been established. When using cross-memory support, the restrictions shown in Table 4 may apply.

Table 4. Cross-Memory Support Restrictions

Macros	Restriction Code	Instructions	Restriction Code
LXRES	2	EPAR	2
ETCRE	2	ESAR	2
ETCON	2	IAC	2
AXRES	2	IVSK	2
AXERT	2	MVCP	3
AXSET	2	MVCS	3
ATSET	2	MVCK	3
PCLINK	1	PC*	4
		PT*	3 and 4
		SAC	1

Table 4. Cross-Memory Support Restrictions (continued)

Macro	Restriction Code	Instructions	Restriction Code
-------	------------------	--------------	------------------

Code Definitions:

1. You must release this resource before checkpoint and reestablish it after checkpoint or restart.
2. You must reestablish this resource after a restart.
3. You must be sure that the keys and ASIDs are correct after a restart.
4. Any attempt to restart from a checkpoint taken in a program call (PC) routine is unpredictable. If a PC routine issues a PCLINK, and then a CHKPT, the error is detected and a checkpoint request is refused. It is impossible to detect a PC routine when the PC routine does not issue a PCLINK and the PC routine is entered without a space switch. If a PC routine issues a checkpoint, that PC routine must ensure that all the requisite values needed for a PT instruction are current before issuing the PT instruction. These values may change after a restart (for example, the ASID might be different).

Data Spaces

A program can use checkpoint/restart while it has one or more entries for a data space on its access list (DU-AL or PASN-AL). In such a case, the CHKPT macro only processes successfully if the program has specified on the ALESERV macro that entries made to the access list for the data space should be ignored for checkpoint/restart processing (CHKPT=IGNORE). A program that specifies the CHKPT=IGNORE option assumes full responsibility for managing the data space storage after a restart operation.

See *z/OS MVS Programming: Assembler Services Reference ABE-HSP* and *z/OS MVS Programming: Assembler Services Guide* for more information about the ALESERV macro.

Name/Token Services

You cannot take a checkpoint if the program has a name/token pair that is not a task-level name/token that was created with the IEANT_CHECKPOINTOK value in the persist parameter. See the IEANTASM macro. You can undefine a name and redefine it after taking a checkpoint or you can add the IEANT_CHECKPOINTOK option. Ensure that the name and the token contain information that will be valid after the restart. For example virtual storage addresses in user subpools will be unchanged but the ASID (address space identifier) probably will be different. The name/token service is described in *z/OS MVS Programming: Assembler Services Reference ABE-HSP*.

Chapter 3. Requesting Restart

Table 5 contains directive information about this chapter's contents:

Table 5. Chapter 3. Requesting Restart

Section
"Automatic Restart"
"Deferred Restart" on page 31
"Coding the RD and RESTART Parameters" on page 38

Types of Restart

The system allows two types of step restart:

- Automatic step restart
- Deferred step restart

The system allows two types of checkpoint/restart:

- Automatic checkpoint/restart
- Deferred checkpoint/restart

Automatic restarts occur in different ways:

- Automatic step restarts at the beginning of a job step that has abended
- Automatic checkpoint/restart restarts at the last checkpoint taken before the job failed

Deferred restart occurs when a job is resubmitted with the RESTART parameter specified on the JOB statement:

- Deferred step restart occurs at the beginning of the job step specified in the RESTART parameter of the JOB statement.
- Deferred checkpoint/restart occurs at the checkpoint specified in the RESTART parameter of the JOB statement.

Automatic Restart

Because an automatic step restart and an automatic checkpoint/restart are similar, they are discussed together in the following sections.

Operator Message Sequence

During processing related to an automatic checkpoint/restart, the system issues the following sequence of messages to the operator:

1. A message is issued each time a checkpoint entry is written. Each message contains the checkpoint identification.
2. If the job step terminates because of an abend condition, an abend message is issued for the job step.
3. If the abend code makes the job step eligible for restart and a job journal is present, an authorization for restart message (IEF225D) is issued that requires a reply.

4. A message is issued, indicating the virtual storage requirements (beginning address and ending address) of the job step to be restarted.
5. Normal mount and restart messages are issued, except when using the Tape Library Dataserver.
6. If password-protected data sets are to be repositioned at restart time, a password message is issued.
7. If additional checkpoint data sets (other than the data set used for restart) are encountered at restart time, checkpoint data set security messages are issued.
8. A successful restart message is issued.

During processing related to an automatic step restart after a job step has terminated abnormally, the system issues the following sequence of messages to the operator:

1. An abend message is issued for the job step.
2. If the abend code makes the job step eligible for restart, and a job journal is present, an authorization message is issued that requires a reply.
3. Normal mount messages are issued.

Note: The abend message, which is issued as:

```
IEF4501 jobname.stepname.procstepname. ABEND code
```

is always displayed if a job step terminates abnormally. In addition, if the job step is being executed and the system fails, this message will be displayed during the next IPL if a system-supported restart is performed. The *code* part of the message has the form Shhh (S followed by a 3-character hexadecimal number) if the system executed the abend macro instruction, or Udddd (U followed by a 4-digit decimal number) if your program executed the abend. It is S2F3 if a system failure occurred. For more information about messages see *z/OS MVS System Messages, Vol 1 (ABA-AOM)*, *z/OS MVS System Messages, Vol 2 (ARC-ASA)*, *z/OS MVS System Messages, Vol 3 (ASB-BPX)*, *z/OS MVS System Messages, Vol 4 (CBD-DMO)*, *z/OS MVS System Messages, Vol 5 (EDG-GFS)*, *z/OS MVS System Messages, Vol 6 (GOS-IEA)*, *z/OS MVS System Messages, Vol 7 (IEB-IEE)*, *z/OS MVS System Messages, Vol 8 (IEF-IGD)*, *z/OS MVS System Messages, Vol 9 (IGF-IWM)*, and *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

Operator Considerations

If a step requests automatic restart and is eligible for restart, the system displays the following message to request authorization for the restart:

```
xxIEF225D SHOULD jobname.stepname.procstepname
[checkid.] RESTART
```

Checkid appears in the message only if restart at a checkpoint is requested. It contains from 1 to 16 characters and identifies the checkpoint entry to be used to perform the restart. The operator must reply to the request for authorization as follows:

```
REPLY xx,{'YES'|'NO'|'HOLD'}
```

YES authorizes the restart, HOLD postpones it, and NO prohibits it.

If it is not clear that the restart should proceed, the operator should reply HOLD to the authorization message. Later, if restart should be allowed, initiate the restart by using the RELEASE command, which will be the same as an initial YES reply. If the decision is to deny the restart authorization, the operator can cancel the job in

the HOLD queue. However, that HOLD, as well as YES, causes special disposition processing to occur during the abnormal termination. This processing keeps all OLD (or MOD) data sets and deletes all NEW data sets if a step restart is requested and keeps all data sets if a checkpoint/restart is requested. If you decide to disallow the restart but want to allow normal disposition (as requested on the job's DD statements) of data sets that were kept, release the job, wait until restart has begun, and then cancel the job.

If a job is canceled in the HOLD queue, any paging space allocated to a VIO data set is not released until the system is reinitialized. To avoid this situation, the operator should release the job, wait until the restart begins, and then cancel the job.

Using the VARY and UNLOAD commands, before the operator replies YES, and after the authorization request, may cause the system's volume and device configuration during a restart execution of the job to be different from what it is during the original execution of the job. The operator may eliminate use of defective volumes and devices.

The ability to use a different volume at restart usually exists only in the case of a NEW data set on a nonspecific volume that is not open at the checkpoint. If a checkpoint/restart is performed, the data set must not be open at the checkpoint. The ability to use a different device does not apply to the device or devices containing the system residence (SYSRES), SPOOL, or PAGE packs. If a checkpoint/restart is to be performed and a data set is open at the checkpoint, the same type of device must be allocated to the data set during both the original and restart executions.

Note: When an initiator selects a job for automatic step restart and the job is reinterpreted, no message is issued to the operator regarding virtual storage requirements, because its execution is not location dependent.

If ADDRSPC=REAL is specified on the JOB or EXEC statement, the restart may be delayed by the system waiting for the allocation of storage. If another job is using the required storage, you do not receive a message; you receive only a delay. Enter a DISPLAY A command to see if a system task or another job is using the storage required by the job with a nonpageable region. You can stop or cancel the conflicting task or job.

How to Request Automatic Checkpoint/Restart

An automatic checkpoint/restart occurs if:

- The job journal option is specified, or
- The step requests restart (RD=R), and
- A successful checkpoint has been taken, and
- The step is eligible for restart because it was terminated by an ABEND macro that returned an eligible completion code (specified by the SYS1.PARMLIB member SCHEDxx RESTART/NORESTART statement type at system initialization), or because system failure occurred, and
- The operator authorizes the restart. This authority enables the operator to control the number of restarts of the same step or from the same checkpoint.

If a step fails and an automatic checkpoint/restart is requested, restart occurs automatically at the last checkpoint taken.

Execution of the CHKPT macro requests this type of restart and establishes the checkpoint. You must provide an ordinary DD statement for the checkpoint data set. Figure 7 illustrates a job requesting automatic restart at a checkpoint. For more information see *z/OS MVS Initialization and Tuning Guide*.

```
//MYJOB   JOB      MSGLEVEL=11,RD=R2
//STEP1   EXEC
          .
          .
//STEP2   EXEC     PGM=MYPROG   MYPROG issues the CHKPT macro
//NAME1   DD       DSN=NAME2   Describes the data set into
//*                               which checkpoint entries
//*                               are to be written
```

Figure 7. Requesting Automatic Checkpoint/Restart

Figure 7 Notes:

- ¹ MSGLEVEL=1 is optional.
- ² RD=R need not be specified if automatic checkpoint/restart is used through job journaling. See “Job Journal Requirements” on page 3.

Automatic Step Restart

An automatic step restart occurs if all of the following are satisfied:

- The step requests restart (RD=R or RD=RNC); and
- The step is eligible for restart because it was terminated by an ABEND macro that returned an eligible completion code (specified by the SYS1.PARMLIB member SCHEDxx RESTART/NORESTART statement type at system initialization), or because system failure occurred; and
- The operator authorizes the restart. This authority enables the operator to control the number of restarts of the same step or from the same checkpoint.

If a step fails and automatic step restart is requested, restart occurs automatically at the beginning of the step that failed.

Automatic step restart may be requested by coding the RD parameter (**RD=R** or **RNC**) on the JOB or EXEC statement in the originally submitted job. Checkpoint processing is suppressed if **RD=RNC**.

Figure 8 illustrates the JCL of a job requesting automatic step restart.

```
//MYJOB   JOB      MSGLEVEL=11,RD=R   Requests automatic
//*                               restart at the
//*                               beginning of any step
//*                               that terminates
//*                               abnormally
//STEP1   EXEC
//STEP2   EXEC     PGM=MYPROG,RD=R2   Requests automatic
//*                               restart of STEP2 if it
//*                               terminates abnormally
//STEP3   EXEC
```

Figure 8. Requesting Automatic Step Restart

Figure 8 Notes:

- ¹ MSGLEVEL=1 is optional.
- ² Note that if RD=R appears on the JOB statement, it is not required on the EXEC statement.

How a Step Restart Is Initiated

A step restart is initiated in the same way as it is during a normal execution. The devices allocated to the restart step can be different (but of the same device type) from the devices allocated originally. If the allocated devices differ, volumes must be moved from one device to another. If automatic volume recognition (AVR) is used and if the devices are available for allocation, devices containing the required volumes are allocated.

Unless the volumes are already mounted, normal mounting messages request the operator to mount the required volumes on the devices after devices are allocated to the restart step. The requested volumes resume processing.

How Data Sets Being Extended Are Handled during Automatic Step Restart

The following sections describe how data sets being extended are handled when automatic step restart is requested.

DISP=MOD

When automatic step restart is requested, the system saves the TTR and track balance of the end of the data set, for each data set with DISP=MOD that is on a direct access volume and used by the step. Saving occurs when each data set is first opened. If restart occurs, the saved TTRs indicate the ends of the data sets when the data sets open again. Thus, if the step writes data in such a data set during the original execution, the step writes over the data during the restart. The action described here does not occur if restart at a checkpoint occurs.

If a data set with DISP=MOD on tape is used in the restart step, the data set is not repositioned at the start of the restart execution. Data written into it during the restart execution follows the data written during the original execution. You may reposition the data set so the data written during the restart execution overlays the data written during the original execution.

OPEN EXTEND or OUTINX

If the data set opens with the OUTINX or EXTEND options, and its DD disposition parameter was

DISP=NEW: there are no restrictions, because the data set reallocates prior to the restart.

DISP=OLD or DISP=SHR: the data set repositions after the last record added before the job terminated. It is possible to add records twice to the data set.

DISP=MOD: the following restrictions apply to a multivolume data set:

- If the volume sequence does not point to the last volume, records overlay or an abend occurs.
- If the volume sequence is not supplied, loading starts on the last volume on which the data set exists. For TAPE data sets only, if the end of data occurs on a volume other than the last allocated volume, processing starts on the wrong volume.

JCL Requirements and Restrictions on Automatic Restart

To allow an automatic step restart or an automatic checkpoint/restart, you must observe the following rules when preparing the job used in the original execution:

1. If a step restart is desired, the RD parameter must be coded to request the restart.
2. If a checkpoint/restart is desired, a DD statement for the checkpoint data set must be included in the step that executes the CHKPT macro instruction.
3. The EXEC statements in the job must have unique names. (Upon restart, the system searches for the named step.)
4. If commands are in the original job, the commands are not reexecuted when restart occurs.
5. The DD statement VOL=REF parameter is ignored if restart is attempted from a checkpoint taken when the DD statements are opened out of order and the referenced DD statement requested nonspecific tape volumes.
6. When using passed data sets, the same temporary data set name should not be reused after the temporary data set has been deleted.
7. If the checkpoint data set is SMS-managed, the automatic class selection (ACS) routine must not override the unit or volume selection.

Resource Variations Allowed in Automatic Restart

The system's device and volume configuration during a restart execution of a job can differ from the status during the original execution of the job.

The ability to use a different volume at restart exists only in the case of a new data set on a nonspecific volume that is not open at the checkpoint. If a checkpoint/restart is to be performed, the data set must not be open at the checkpoint. The ability to use a different device does not apply to the device or devices containing the system residence (SYSRES) volume and the LINKLIB data set. Also, if a checkpoint/restart is to be performed, the same device type must be allocated to the data set during both the original and restart executions.

How the System Works at Automatic Restart

This section describes how the system works when an automatic restart occurs.

How Data Set Disposition Is Determined

When a step requests restart and is eligible for restart, disposition processing of the data sets used by the step or by the job does not occur until the operator replies to the request for authorization. If the operator denies restart, disposition processing occurs normally and programmer-specified final or conditional dispositions are performed. If you indicate that a step is to be executed after abnormal termination (COND=EVEN or COND=ONLY) the step is executed. The next step in a job that will be executed based on abnormal termination (COND=ONLY), will not execute if a restart is successfully performed.

- If step restart is to occur, all data sets with OLD or MOD dispositions in the restart step and all data sets passed around the restart step are kept, even if they are declared temporary. Temporary data sets normally cannot be kept.
- All data sets with NEW dispositions in the restart step are deleted.
- If a checkpoint/restart is to occur, all data used by the job (data sets that were not previously disposed of) is kept.

If the operator performs an operator-deferred restart by replying HOLD to the request for authorization, a CANCEL command for the job can be issued instead of

a RELEASE command. If CANCEL is issued, no further data set disposition processing or step executions occur. The disposition of these data sets remains as it was when the HOLD was issued.

How the Job Is Reinterpreted

After disposition processing is completed, the job is reenqueued by the job entry subsystem and is eligible for selection. After the job is selected, the system begins restart processing by reinterpreting the job. The system uses an internal representation of the original job to perform this function. The job is not read in again.

When the information is merged and step restart occurs, the following changes occur:

- Except for data sets not opened during the original execution that requested nonspecific tape volumes, data sets specified as NEW in the restart step have their dispositions changed to OLD.
- In the case of data sets for which nonspecific volumes are requested in the restart step, the device type and serial numbers of the volumes assigned to the data sets during the original execution are used.
- In the case of multivolume data sets, the volume in use at the checkpoint, and not the first volumes of the data sets, are mounted (if they have not remained mounted) during the restart. For VSAM, however, volume mounting is based on the present catalog information and may result in the mounting of unneeded volumes; for example, the first volume of a sequential data set may be temporarily mounted although it is no longer required at the checkpoint being restarted and may be demounted almost immediately. This situation may be avoided through specification of the parallel mount subparameter of the DD statement UNIT parameter (UNIT=(,P)). For further information, see the *z/OS MVS JCL User's Guide*.
- The system-generated name for any temporary data set uses the time stamp obtained when the job was originally interpreted.

In addition, any modification made to the job's environment by the use of the dynamic allocation facilities prior to the last checkpoint appears as it did at the time of the checkpoint.

Deferred Restart

The following sections describe a deferred restart.

Message Sequence

To perform a deferred checkpoint/restart, you resubmit in an input job stream the job to be restarted. Messages that contain checkpoint entry identifications are displayed on the console during the original execution of the job and thus can be used by the programmer preparing the job for resubmission. When the resubmitted job is restarted, messages appear on the console in the following sequence:

1. A message asking if the checkpoint data set that is being used for restart is a secure volume.
 - For Tape Library Dataserver tape library devices, the tape library configuration data base will be read for confirmation of a secure checkpoint volume.
2. If additional checkpoint data sets, other than the data set used for restart, are encountered at restart time, checkpoint data set security messages are issued.

3. A message indicating the virtual storage requirements of the job.
4. Normal mount messages.
5. If password-protected data sets are to be repositioned at restart time, a password message is issued.
6. A successful restart message.

To perform a deferred step restart, the job to be restarted is resubmitted. Normal mount messages are displayed.

Operator Considerations

When a job is resubmitted to perform a deferred checkpoint/restart (the RESTART parameter is coded on the JOB statement with a checkid operand), the processing is essentially the same as that during an automatic checkpoint/restart after the restart reader has reinterpreted the job. A message is issued to the operator, indicating the virtual storage requirements of the job.

The required virtual storage area can be unavailable for the following reasons:

- The REGION size parameter for the step is larger when the job is resubmitted than in the original execution, and the area required is at different addresses than the area available. The parameter REGION=0K causes this type of restart failure.
- A new IPL is performed and, because of different IPL options specified by the operator, the area required is larger than the area available.

If these conditions exist, a message is displayed, indicating that virtual storage for the job step to be restarted is unavailable.

Deferred Step Restart

You cause a deferred step restart of a job by coding the RESTART parameter without a checkid on the JOB statement and resubmitting the job. The parameter specifies a job step or a step of a cataloged procedure. It restarts the job at the beginning of the specified step. Steps preceding the restart step are interpreted, but not initiated.

The CHKPT macro instruction may or may not be coded in your program. Figure 9 on page 33 illustrates a job originally submitted and the same job resubmitted for step restart. Assume that the results of STEP2 were unsatisfactory because of abnormal termination or incorrect data when the job was executed originally.


```

Original Job
//MYJOB JOB      MSGLEVEL=11      No automatic restart
//*                               requested
//STEP1 EXEC
.
.
.
//STEP2 EXEC      PGM=MYPROG
.
.
.
//STEP3 EXEC

Resubmitted Job
//MYJOB JOB      MSGLEVEL=1,1      Causes restart of job
//                               RESTART=STEP2      at STEP2
//STEP1 EXEC
.
.
.
//STEP2 EXEC      PGM=MYPROG
.
.
.
//STEP3 EXEC

```

Figure 9. Requesting a Deferred Step Restart

Figure 9 Note:

¹ MSGLEVEL=1 is optional.

JCL Requirements and Restrictions on Deferred Step Restart

To perform a deferred step restart, you must provide the data set environment required by the restart job. This can be accomplished by using the conditional disposition subparameter in the appropriate DD statements during the original execution of the job. Conditional dispositions in the original job must be used to:

- Delete all NEW data sets used by the step to be restarted.
- Catalog all data sets passed from steps preceding the restart step, to the restart step, or to steps following the restart step. Abnormal termination of the restart step, when it is originally run, causes cataloging of the passed data sets. The information is available to the following steps when the job is resubmitted.
- Keep all OLD data sets used by the restart step, other than those passed to the step.

If a data set with DISP=MOD on tape is used in the restart step, the data set is not repositioned at the start of the restart execution, and data written into it during the restart execution follows the data written during the original execution. You may want to reposition the data set after restart so the data written during the restart execution overlays the data written during the original execution. Any data sets that are dynamically deallocated have the disposition specified at the time deallocation occurred. Conditional disposition processing is not done during abend.

The following rules apply to the restart job:

1. Code the RESTART parameter on the JOB statement.
2. If data sets are passed from steps preceding the restart step to the restart step or to steps following the restart step, the DD statements receiving the data sets must entirely define the data sets. They must explicitly specify volume serial

number, device type, data set sequence number, and label type, unless this information can be retrieved from the catalog. It is recommended that passed data sets be conditionally cataloged during abnormal termination of the original execution. Label type cannot be retrieved from the catalog.

3. Generation data sets created and cataloged in steps preceding the restart step must not be referred to in the restart step or in steps following the restart step by the relative generation numbers used to create them. They must be referred to by their actual relative generation numbers. For example, a data set created as the +1 data set must be referred to as the 0 data set (assuming that the +2 data set was not also created).
4. The EXEC statement PGM and COND parameters and the DD statement VOL=REF parameter must not be used in the restart step or in steps following the restart step if they contain values of the form *stepname* or *stepname.procstepname*, referring to a step preceding the restart step.

Resource Variations Allowed in Deferred Step Restart

A deferred step restart only allows the restarted execution of a job to begin at other than the first step of the job. Job step initiation and allocation of resources are accomplished normally. The following variations are allowed upon restart:

- Variation of device and volume configuration
- Variation in JCL and data in the resubmitted job.

Deferred Checkpoint/Restart

You can cause a deferred checkpoint/restart of a job by the following procedure:

1. Use the option of coding a special form of the RD parameter (RD=NR) in the original job. This specifies that, if the CHKPT macro instruction executes, a checkpoint entry is written but an automatic checkpoint/restart is not requested. Omission of RD=NR will not preclude a deferred checkpoint/restart.
2. Cause execution of the CHKPT macro instruction, which writes checkpoint entry.
3. Resubmit the job whether or not it terminated abnormally. For example, resubmit it because a volume of one of its input data sets was in error and had caused the corresponding part of an output data set to be in error.
4. Code the RESTART parameter (**RESTART=(stepname,checkid)**) on the JOB statement of the restart job. The parameter specifies both the step to be restarted and the *checkid* that identifies the checkpoint entry to perform the restart.
5. Place a SYSCHK DD statement immediately before the first EXEC statement in the restart job. It specifies the checkpoint data set from which the specified checkpoint entry is read and is additional to any DD statements in the job that define data sets into which checkpoint entries are written. Figure 10 on page 35 illustrates a job when it is originally submitted and when it is resubmitted for a deferred checkpoint/restart. Assume in Figure 10 on page 35 that STEP2, when originally executed, terminates abnormally at some time after CH04 is written. In the resubmitted job, you request that STEP2 be restarted using the checkpoint entry identified as entry CH04.
6. The DD statement VOL=REF parameter is ignored if restart is attempted from a checkpoint taken when the DD statements are opened out of order and the referenced DD statement requested nonspecific tape volumes.

```

Original Job

//MYJOB JOB      RD=NR           Requests that
//*                automatic restart not
//*                occur (optional)
//STEP1 EXEC
.
.
.
//STEP2 EXEC    PGM=MYPROG       MYPROG issues CHKPT
//*                macro
//*
//NAME1 DD      DSN=NAME2,DISP=(,CATLG,KEEP)...
//*                Describes checkpoint
//*                data set
.
.
.
//STEP3 EXEC

Resubmitted Job

//MYJOB JOB      RESTART=(STEP2,CH04) Request restart at
//*                CH04 in STEP2
//SYSCHK DD      DSN=NAME2,DISP=OLD  Describes data set
//*                that contains CH04
//STEP1 EXEC
.
.
.
//STEP2 EXEC    PGM=MYPROG
//NAME1 DD      DSN=NAME2,DISP=(,CATLG,KEEP)...
//*                Describes data set
//*                in which new
//*                checkpoint entries
//*                will be written
.
.
.
//STEP3 EXEC

```

Figure 10. Requesting a Deferred Checkpoint/Restart

JCL Requirements and Restrictions on Deferred Checkpoint/Restart

To perform a deferred checkpoint/restart, you must provide the data set environment required by the restart job, using conditional dispositions during the original execution.

Conditional dispositions should be used to:

- Keep all data sets used by the restart step.
- Catalog all data sets passed from steps preceding the restart step to steps following the restart step. Although the step that terminates abnormally is not using the passed data sets, its termination causes the cataloging of the data sets if the conditional catalog parameter is used in the preceding steps.

Note:

1. Temporary data sets cannot be kept.

2. If the DD statement for any data set at checkpoint time had DISP=(NEW,DELETE,CATLG), you probably want to uncatalog after completion of the restarted job or job step by using access method services (IDCAMS), ISPF or IEHPROGM.

Dynamically deallocated data sets have the disposition processed as specified at the time of deallocation. Conditional disposition processing is not performed during abend.

The following rules must be adhered to when resubmitting a job for a deferred checkpoint/restart:

1. A RESTART parameter with a *checkid* subparameter must be coded on the JOB statement.
2. A SYSCHK DD statement must be placed in the job immediately before the first EXEC statement.
3. The EXEC statements in the job must have unique names. (Upon restart, the system searches for the named restart step.)
4. The JCL statements and data in steps preceding or following the restart step can be different from their original forms. However, all backward references must be resolvable.
5. The restart step must have a DD statement corresponding to each DD statement present in the step in the original, and the names of the statements must be the same as they were originally. However, the restart step can contain, in any position, more DD statements than it contained originally. Except as noted in item 9 below. The total number of volumes specified at restart must equal or exceed the number specified at the checkpoint.
6. If a DD statement in the restart step in the original job defined a data set that was open at the checkpoint to be used, the corresponding statement in the restart job must refer to the same data set, and the data set must be on the same volume, and, in general, have the same extents recorded in its DSCB as it did originally. (See the exceptions in the note that follows.) This means no one can use the DEFRA function of DFSMSdss to consolidate extents. If the data set is multivolume and is processed by the sequential access method (SAM), only the part of the data set on the volume in use at the checkpoint needs to be the same as it was originally.

Note: The extents can differ as follows:

- In the DD statement, you can request that additional space be allocated to the data set when the space currently available is exhausted. If space is allocated after a checkpoint is taken, this space is indicated in the DSCB; after restart from the checkpoint, the space is released, and the DSCB contents are changed to what they were at the checkpoint.
- In the DD statement, you can code RLSE to request that unused space be released when the data set is closed. If the space is released, the DSCB may indicate a reduced extent for the data set when deferred restart at a checkpoint occurs; no space is allocated to replace released space. Space is not released when a data set is open when the step abends, and step termination is followed by automatic restart.

JCL specified data sets that were dynamically unallocated prior to the checkpoint are not allocated at restart.

When there is no need to read or modify a data set after restart, the data set can be replaced by a dummy data set if the original data set was processed by SAM and the job step is not restarted from a checkpoint within the data set's end-of-volume exit routine. A VSAM data set using the ISAM interface to

VSAM cannot be replaced by a dummy data set. Any dummy data set present at the time of the checkpoint must be present as a dummy data set at restart. Allocation is done for each DD statement in the job step where the checkpoint is taken, even if the data set is closed at the time of the checkpoint.

7. The data in the restart step need not be the same as it was originally. If data following a DD * statement was present originally and is entirely omitted in the restart job, the delimiter (/*) statement following the data may also be omitted. The delimiter statement following a DD DATA statement must not be omitted.
8. The VOL parameter of a DD statement must refer to at least those volumes it refers to at checkpoint time. More volumes may be added if desired. The following DD statement parameters cannot be changed: DISP, DCB, LABEL, and UNIT.
9. Except for the requirements stated in rules 4 through 8, the JCL statements and data in the restart step can be different from their original forms. In particular, the DUMMY parameter can be used for any data that was not open at the checkpoint.
10. If data sets are passed from steps preceding the restart step to steps following it, the DD statements receiving the data sets must entirely define them. They must explicitly specify volume serial number, device type, data set sequence number, and label type, unless this information can be retrieved from the catalog. It is recommended that passed data sets be conditionally cataloged during abnormal termination of the original execution. Label type cannot be retrieved from the catalog.
11. The EXEC statement PGM and COND parameters and the DD statement VOL=REF parameter must not be used in the restart step or in steps following the restart step if they contain values of the form *stepname*, or *stepname.procstepname*, referring to a step preceding the restart step.
12. The DD statement VOL=REF parameter is ignored if restart is attempted from a checkpoint taken when the DD statements are opened out of sequence and the referenced DD statement requested nonspecific tape volumes.
13. The volume serial number must be coded on a restart DD for a generation data group if the deferred checkpoint/restart volume serial number list is to be different from the original.

Resource Variations Allowed in a Deferred Checkpoint/Restart

The system's device and volume configuration can differ from the status during the original execution of the job. The allowable differences are those described under "Resource Variations Allowed in Automatic Restart" on page 30.

System Restrictions for Deferred Checkpoint/Restart

The following restrictions apply when submitting a job for deferred checkpoint/restart:

- Jobs specifying nonpageable storage (ADDRSPC=REAL) require an area of storage identical to the storage originally requested at the time the checkpoint is taken and starting at the same address.
- The link pack area modules in use at the time the checkpoint is taken must reside in the same storage locations for the job submitted for deferred checkpoint/restart.
- Any time a change link pack area (CLPA) is required, the locations of modules in the link pack area (LPA) might change without any changes to the LPA modules themselves. This prevents the restart from any checkpoint taken prior to the IPL in which the CLPA is required.

Changes to hardware features and other system reconfigurations may require a CLPA. Review the instructions included with any such change carefully to determine if a CLPA will be required.

- The nucleus must not be changed between checkpoint and restart.
If the required storage is not available, it may be for one of the following reasons:
- The link pack area expands into the required storage
- The system queue area expands into the required storage.

If either of the above conditions occurs, contact your system programmer for a respecification of the system parameter to modify the area size and repeat initial program loading using the new value.

How the System Works during a Deferred Checkpoint/Restart

After the system reads and interprets the restart job, it reads the specified checkpoint entry and merges information from it into the scheduler work area entry for the job. As a result, the work queue entry differs from the entry existing during the original execution, as described earlier. See “How the Job Is Reinterpreted” on page 31.

Next, the system initiates the restart step normally. The system reads the specified checkpoint entry again and functions as in automatic restart. Restart is delayed until the required virtual storage area is available.

When a job restarts correctly, the system operator receives two messages: IHJ006I and IHJ008I. If these messages do not appear for nonpageable storage jobs (ADDRSPC=REAL), enter a DISPLAY A command to see if a system task or another job is using the required storage. The system operator can then stop or cancel the conflicting job.

If there are multiple tape or disk volumes, the system may ask the operator to mount data volumes other than those required at the beginning of the job. In addition, any card input data sets that are used by the failing job step must again be made available to the system.

Coding the RD and RESTART Parameters

The following sections describe how to code the RD and RESTART parameters.

RD Parameter

You code the RD parameter in the JOB or EXEC statement to request automatic step restart if failure occurs, and/or to suppress, partially or totally, the action of the CHKPT macro instruction. If the RD parameter requests an automatic step restart if failure occurs, or if the RD parameter is not coded, the action of CHKPT is normal: CHKPT writes a checkpoint entry and requests a checkpoint/restart to be performed if failure occurs. The RD parameter is ignored for system tasks and generalized start jobs.

The RD parameter, coded on an EXEC statement, applies to the step corresponding to the statement or to all steps of the cataloged procedure referred to by the statement. Coded on a JOB statement, the RD parameter applies to all steps of the corresponding job and overrides an RD parameter coded in any EXEC statement of the job. The syntax of the parameter is:

RD[*procstepname*]={R|NC|NR|RNC}

The possible definitions are:

RD=R

(Restart) Requests an automatic step restart if failure occurs. If the CHKPT macro instruction is executed in the step, the resulting request for an automatic checkpoint/restart overrides the request for an automatic step restart. This parameter is ignored if the job does not contain a job journal. For JES2 and JES3, this parameter forces job journaling.

RD=NC

(No checkpoint) Does not request an automatic step restart. It totally suppresses the action of the CHKPT macro instruction if the macro instruction executes in the step. This allows use of a program containing CHKPT when the checkpoint function is not wanted. It can also suppress the checkpoint at end-of-volume.

RD=NR

(No automatic restart) Does not request an automatic step restart. It suppresses the request for an automatic checkpoint/restart that is made when the CHKPT macro instruction executes in the step. If CHKPT executes, it writes a checkpoint entry normally. A deferred restart can be performed from the checkpoint entry.

RD=RNC

(Restart and no checkpoint) Requests an automatic step restart if failure occurs. It totally suppresses the action of CHKPT if CHKPT executes in the step. It can suppress the checkpoint at end-of-volume. If the job does not contain a job journal, the step is ineligible for automatic restart. For JES2 and JES3, this parameter forces job journaling.

If **RD=value** is coded on an EXEC statement that invokes a cataloged procedure, the parameter applies to all steps of the procedure and overrides all RD parameters present in the EXEC statements of the procedure.

RD.procstepname=value can be coded instead of **RD=value**; it applies to the specified procedure step and overrides the RD parameter that may be coded on the EXEC statement of the procedure step. **RD.procstepname=value** can be coded once for each step of the procedure.

RESTART Parameter

The RESTART parameter requests a deferred restart of a job. Code it in the JOB statement when you resubmit the job. If step restart is to occur, this parameter specifies which step to begin. If the restart is to occur at a checkpoint that was taken during a step, specify both the step and the identification of the particular checkpoint entry. The syntax of the parameter is:

RESTART=(*{stepname | stepname.procstepname | *}* [,*checkid*])

Both operands are used if restart at a checkpoint is to occur. If a step restart is to occur, *checkid* must be omitted; the enclosing parentheses can be omitted.

stepname

name of the step to be started.

stepname.procstepname

if a step of a cataloged procedure is to be restarted, *stepname* is the name of the step executing the procedure and *procstepname* is the name of the step within the procedure.

*

the first step of the job (possibly a step within a cataloged procedure) is to be restarted.

checkid

can contain up to 16 characters in any combination of alphanumeric characters, printable special characters, and blanks. If it contains any special characters or blanks, it must be enclosed in single apostrophes, and apostrophes within it must be represented as double apostrophes.

SYSCHK DD Statement

If you wish to perform a deferred checkpoint restart you must include a SYSCHK DD statement in the resubmitted job. This DD statement specifies the checkpoint data set that contains the checkpoint entry to be used in the restart. The SYSCHK DD statement cannot be included when a deferred step restart is to be performed.

The statement must precede the first EXEC statement in the job that performs a deferred restart at checkpoint. It must follow the JOBLIB DD statement if the JOBLIB DD is present. The desired checkpoint entry must be named by the *checkid* subparameter of the JOB statement RESTART parameter.

The following requirements and restrictions apply to the SYSCHK DD statement:

- The statement must contain or imply DISP=(OLD,KEEP).
- For SMS-managed data sets, UNIT and VOL=SER are not required.
- The statement must define the checkpoint data set. It must specify its name. If it is not cataloged, the statement must also specify the device type or device number and specify the volume serial number.
- If the volume containing the checkpoint data set is mounted on a JES3-managed device, the SYSCHK DD statement must not request deferred mounting.
- The SYSCHK data set cannot be multivolume or concatenated. If the checkpoint data set is multivolume, the SYSCHK DD statement must specify, as the first volume of the data set, the volume and data set name that contain the desired checkpoint entry. The serial number of the volume containing a particular entry appears in the console message that is written when the entry is written.
- If the checkpoint data set is partitioned, the DSNNAME parameter on the SYSCHK DD statement must not contain a member name.
- If a RESTART parameter without the *checkid* subparameter is included in a job, a SYSCHK DD statement must not appear before the first EXEC statement of the job.
- If a RESTART parameter is not included in a job, a SYSCHK DD statement appearing before the first EXEC statement in the job is ignored.
- A SYSCHK DD statement appearing in a step or procedure step of a job is treated as an ordinary DD statement; that is, the name SYSCHK has no special meaning in that case.

An example of a non-SMS SYSCHK DD statement is:

```
//SYSCHK DD DSN=dsnname,DISP=OLD,UNIT=name,  
// VOL=SER=volser
```


For existing SMS-managed data sets, the system ignores the UNIT keyword and any volume serial on the DD statement.

An example of an SMS SYSCHK DD statement is:

```
//SYSCHK DD DSN=dsname,DISP=OLD
```

Chapter 4. Checkpoint/Restart Processing

This chapter describes how the checkpoint/restart routine handles different types of user data sets, and gives additional guidelines on repositioning and preserving the contents of data sets. The end of this chapter includes notes about using checkpoint/restart with various programs and programming languages.

Table 6 contains directive information about this chapter's contents:

Table 6. Chapter 4. Checkpoint/Restart Processing

Section
"Notes for User Data Sets"
"SMS-Managed Sequential Data Sets" on page 47
"VSAM Data Sets" on page 49
"Notes for Tape Labels and Files" on page 52
"Repositioning and Preserving the Contents of Data Sets" on page 54
"Buffering of Data Set Records" on page 54
"Preserving Data Set Contents" on page 55
"Allocating Devices during Deferred Checkpoint/Restart" on page 55
"Handling Checkpoint/Restart Errors" on page 56
"Using Checkpoint/Restart with Other Programs" on page 57

Notes for User Data Sets

The data set types discussed are listed below:

- BDAM (basic direct access method)
- Extended format
- Generation
- Graphics
- Multivolume
- Partitioned data sets (PDSs)
- Partitioned data set extended (PDSEs)
- SMS-managed sequential data sets
- SYSABEND
- SYSIN
- SYSOUT
- TCAM (telecommunications access method)
- Temporary
- VSAM (virtual storage access method)
- Work.

BDAM Data Sets

When the basic direct access method (BDAM) is used to process a data set, processing resumes normally upon restart. However, you must ensure that a

particular block is read or written before a checkpoint is taken. Your program must complete the BDAM I/O operation by executing the CHECK or WAIT macro before it executes the CHKPT macro.

Any change made to a record between a checkpoint and a restart will remain in the data set. You are responsible for backing out any changes that would otherwise produce invalid results. For more information, see "Preserving Data Set Contents" on page 55. If the program does not complete the operation, the block may be read or written either before or after the checkpoint is taken.

For serially reusable resources, your program must also issue either the WRITE or RELEX macro to release a record that is read with exclusive control, before executing the CHKPT macro.

Extended Format Data Sets

A checkpoint data set cannot be an extended format data set.

A checkpoint failure occurs if a checkpoint is requested when an extended format sequential data set is open. You can take a checkpoint while an extended format VSAM data set is open. These extended format VSAM data sets can be striped, compressed format, with extended addressability or any combination thereof.

Large Format Data Sets

Checkpoint/restart will work while a large format data set is open, and you can write a checkpoint to a large format data set.

Generation Data Sets

No automatic cataloging of generation data sets (GDS) takes place. If certain generation data sets of a generation data group are to be cataloged, you must catalog them. The order in which they are cataloged determines the relative generation numbers of the generation data sets for reference by later jobs. The last generation data set becomes the 0 generation, the next-to-last cataloged generation data set becomes the -1 generation, and so on.

You can code the disposition parameter of a generation data set `DISP=(NEW,CATLG,KEEP)` or `DISP=(NEW,CATLG,CATLG)`. Both will allow a restart without any change in the JCL, regardless of how the checkpointed job ended.

If you code `DISP=(NEW,CATLG,KEEP)` and the checkpointed job abnormally terminates, the data set is kept. When the restart job completes successfully, this data set is cataloged. If you code `DISP=(NEW,CATLG,CATLG)` and the checkpointed job abnormally terminates, the data set is cataloged. When the restart job completes, the data set is not cataloged because it already exists on the catalog. However, if the generation data set is to be dynamically allocated, you must code the disposition `DISP=(NEW,CATLG,CATLG)`.

Generation data sets created by one step of a job can be passed to subsequent steps in the same job and can be referenced by the relative generation numbers assigned at the time of creation, whether or not the generations were cataloged.

For deferred step restart, the generation data group name table (GDGNT) is recreated from the catalog. The last generation data set cataloged prior to termination of the job becomes the 0 generation and is used for the base name in

the GDGNT. This may not be the same as the base name when the job was initially run; you must know which generation data sets were cataloged and in what order the data sets were cataloged, and the JCL must be modified accordingly.

For a deferred checkpoint/restart, no modification of the relative generation number in the JCL is necessary.

When a job is started, the base name (0 generation name) of each generation data group is placed in the GDGNT. The GDGNT is never changed unless a new generation level is cataloged. It is saved at a checkpoint and is available for both automatic and deferred checkpoint/restart. The restart takes place without any change in the JCL, whether or not generations previously created by the job were cataloged.

When using generation data groups, the disposition of the data set should be specified correctly to avoid having to change the JCL at restart time. If `DISP=(NEW,CATLG,CATLG)` was specified, then the generation data group (GDG) number is set to (0) at the time of restart and the checkpointed data set is now at `GDG(-1)`, regardless of how the job completed. In order for generation data sets to maintain their respective generation numbers at checkpoint and restart time, the disposition should specify `DISP=(NEW,CATLG,KEEP)`. This will prevent the generation level from changing if the job step fails. If, however, you wish to restart a job that completed without failing, you must change the JCL of the job to be resubmitted to get the correct generation levels.

Note: A job that contains generation data sets referred to by a relative generation number of +1 or greater and with a disposition of OLD, SHR, or MOD is failed by the JES3 interpreter service unless the UNIT parameter is included in the DD statement. With UNIT specified, the JES3 interpreter service permits this data set to be allocated on a deferred basis.

Graphics Data Sets

An open graphics access method (GAM) data control block (DCB) is not supported in Checkpoint/Restart.

Multivolume Data Sets

When multivolume output data sets have nonspecific volume allocation such as `UNIT=SYSDA`, the abnormal disposition must be coded as KEEP to ensure that if the job is restarted, all volume serial numbers are put into the catalog. For example, use `DISP=(NEW,CATLG,KEEP)`, not `DISP=(NEW,CATLG,CATLG)`.

Partitioned Data Sets (PDSs)

If a PDS is compressed after a checkpoint in which the PDS was open, restart at checkpoint should not be requested.

After you take a checkpoint and then open a PDS for adding or replacing members, another checkpoint should be taken before any records are written into the data set. If the second checkpoint is not taken and restart occurs at the first checkpoint, the OPEN routine positions to the current end of the data set instead of to the original end.

If a PDS is open for output (`OPEN,OUTPUT`) at the time of a checkpoint, then the restart function deletes any members added after the checkpoint. If the data set is open with the `UPDAT`, `INOUT`, or `OUTIN` option, then restart does not delete

members unless another DCB is open for output to the data set. For example, the PDS is open for update, and the problem program is updating the third member of a five-member PDS, when checkpoint 1 is taken. If a member is updated as a newly rewritten member after checkpoint 1, and the program is restarted from that point, rewriting the updated member will cause problems because the member already exists.

Adding Members

When a PDS is updated, be careful to preserve the contents of the directory. The directory consists of entries that point to each member of the data set.

When a member is added to a PDS, an entry is also added to the directory. If one member is added, the STOW macro can be used to create the entry, or the member name can be specified in the DD statement. In the latter case, the control program creates the directory entry when the data set is closed or when the job step terminates. Programs create members by issuing a STOW macro after writing each member. You might choose to take a checkpoint after each STOW macro or after several STOW macros.

When one or more members are added to a PDS, a checkpoint should be taken immediately after opening the data set. After taking the checkpoint, the new member can be written and its entry added to the directory. If the step is restarted from a checkpoint, the restart function repositions the data set. This means restart deletes the member so your program can recreate it.

Updating Members

To update a member of a PDS, updated records can either be written back to their original locations, or the entire member (in updated form) may be rewritten as a new member of the data set. In the latter case, the directory entry must be updated to point to the rewritten member.

If a checkpoint is taken before rewriting an entire member, one must also be taken immediately after updating the directory, because the control program deletes the updated directory entry if it repositions the data set for restart from the original checkpoint. Because no entry then points to the original member, the postrestart processing will be unsuccessful.

Deleting Members

The checkpoint restart function deletes members that were added after the checkpoint. Because this action may delete members written by another job (another job may have been executed between the original and restart executions of the subject job), restart at a checkpoint should not be requested.

If a PDS is open for output at checkpoint time, any new members added later will be deleted during a restart from this checkpoint. **No members will be deleted when a PDS is open for update.** Be cautious when restarting from checkpoints having PDSs open for update.

Caution: When a program copies a PDS, the members in the copy generally reside in different locations within the data set. If you take a checkpoint while a PDS is open and then copy members, perhaps with IEBCOPY or DFSMSdss[™], you cannot do a deferred checkpoint restart.

Partitioned Data Sets Extended (PDSEs)

A failure occurs if a checkpoint is requested when a DCB is open to a PDSE, or if a PDS was opened at checkpoint time but was changed to a PDSE by restart time. A

checkpointed data set cannot be a PDSE. For more information on PDSEs see *z/OS DFSMS Using the New Functions* and *z/OS DFSMS Using Data Sets*.

Note: A PDSE can be used as JOBLIB/STEPLIB.

SMS-Managed Sequential Data Sets

When an application program takes a checkpoint, the system records information about the status of that program in a checkpoint data set. This information includes the location on disk or tape where the application is currently reading or writing each open data set. If a data set that is open at the time of the checkpoint is moved to another location before the restart, you cannot restart the application from the checkpoint because the location-dependent information recorded by checkpoint/restart is no longer valid.

There are several system functions (for example, DFSMSHsm™ or DFSMSdss™) that might automatically move a data set without the owner specifically requesting it. To ensure that all checkpointed data sets remain available for restart, the checkpoint function sets the unmovable attribute for each SMS-managed sequential data set that is open during the checkpoint. An exception is the data set containing the actual recorded checkpoint information (the checkpoint data set), which does not require the unmovable attribute.

You can move checkpointed data sets when you no longer need them to perform a restart. DFSMSHsm™ and DFSMSdss™ FORCECP(days) let you use operations such as MIGRATE, COPY or DEFrag to move an SMS-managed sequential data set based on a number of days since the last access. DFSMSHsm™ recall, and DFSMSdss™ restore and copy, are operations that turn off the unmovable attribute for the target data set.

Card Readers

When BSAM, QSAM, or EXCP is used to read a data set from a card reader, your program can reposition the data set upon restart. If you provide a repositioning routine, instruct the operator to position the data set to the beginning if a restart becomes necessary. The program might be designed to operate as follows:

- The program saves the first record read from the data set and keeps a count of the number of records read before each checkpoint.
- After a restart, the repositioning routine reads a record from the data set and compares it with the first record read before abnormal termination. If the records are identical, the data set has been positioned to the beginning. The routine then repositions it by reading (without otherwise processing) the number of records read before the checkpoint.

SYSABEND or SYSUDUMP Data Sets

Whether or not checkpoint/restart is used, abnormal termination causes the system to write a SYSABEND or SYSUDUMP data set if you provide a SYSABEND or SYSUDUMP DD statement. The system uses its own data control block to write the data set, and it opens the data set during abnormal termination processing. You can code or omit the SYSOUT parameter on the SYSABEND or SYSUDUMP DD statement.

When the SYSOUT parameter is coded and automatic restart occurs after abnormal termination, the SYSABEND or SYSUDUMP data set is not printed for step restart. Because the SYSABEND or SYSUDUMP data set is created by the job step, it is deleted during restart.

In all other cases, the SYSABEND or SYSUDUMP data set is printed, whether or not the restart is successful. If a second abnormal termination occurs, a second SYSABEND or SYSUDUMP data set is written. The same rules that apply to the first data sets also apply to the second data set and to all subsequent data sets.

SYSIN Data Sets

When restart at a checkpoint occurs, a SYSIN data set (data following a DD * or DD DATA statement) is repositioned. Unit record data sets are never repositioned. The checkpoint routine waits until all requested input/output operations are complete, then requests that the job entry subsystem save positioning information. When automatic restart occurs, the system keeps the SYSIN data of the job restarting. During the restart execution, the job can read data from the SYSIN data sets as it could during the original execution.

To perform deferred restart, you include any necessary SYSIN data in the resubmitted job.

Information about altering SYSIN data in a restart job is given under “JCL Requirements and Restrictions on Automatic Restart” on page 30.

SYSOUT Data Sets

The following describes how SYSOUT data sets (data sets having the SYSOUT parameter coded on their DD statements) are handled during various types of restart.

Automatic Restart

When your program writes SYSOUT data, the system stores it in spool data sets. During an automatic step restart, JES3 releases the output spool space that the original execution used. After automatic step restart JES allocates new spool spaces.

If a checkpoint/restart occurs, the data sets used during the original execution are kept. If a SYSOUT data set was open when the last checkpoint was taken, it is repositioned to its position at the time the checkpoint was taken.

The checkpoint routine waits until all requested input/output operations are complete, then requests that the job entry subsystem save positioning information. Data written during the restart execution overlays only the data written between the time the last checkpoint was taken and the time the job step terminated abnormally. If a SYSOUT data set is closed at the checkpoint, the data set is not repositioned. If the restart step opens the same data set again, the data written during the restart follows the data originally written. The data set has an implied disposition of MOD.

Deferred Checkpoint/Restart

When a checkpoint/restart occurs, and a SYSOUT data set is open at the checkpoint, the data set written into during the restart is a new data set different from the data set used originally.

TCAM Data Sets

A successful restart of a telecommunications access method (TCAM) data set depends on the following conditions:

- The message control program (MCP) region must be active and have enough virtual storage to build the required control blocks.

- The QNAME parameter in the DD statement of the checkpoint job must be available in the terminal table of the MCP region.

Temporary Data Sets

A temporary data set is a data set where you do not supply a name or one where you code a name that starts with an ampersand. You might choose to allocate space for a temporary data set in a step before a program writes in it. You might do this to simplify the JCL. The system suppresses checkpoints and automatic restarts for any job step that uses a temporary data set that is allocated in a prior job step.

VSAM Data Sets

For VSAM data sets, you are responsible for handling checkpoint/restart problems that arise because of changes in the data. For example, consider a program that updates records in a data set by adding a number to a value already existing in some field within the record. If the program terminates and is restarted, you must ensure that the records processed between the checkpoint and the termination are not processed again after the restart.

During checkpoint/restart processing, no user ACB exits are taken. For checkpoint processing, the ACB exception exit is taken; for restart processing, the ACB exception exit is not taken.

VSAM functions you can use with checkpoint and restart include:

- VSAM Hiperbatch™ environment
- VSAM LSR environment using Hiperspace™ buffering.
- Extended format data sets, including compressed format data sets and those with extended addressability

A checkpoint cannot be taken if VSAM data sets are open in multiple address spaces and if your program uses the CBIC (control blocks in common) option.

When the checkpoint routine records positioning for a VSAM data set, all outstanding I/O requests for the data and index are completed before the contents of your address space are saved. If an error occurs while these I/O requests are processed, the checkpoint procedure stops and a code of X'0C' is returned in register 15. You can handle the error condition and reissue the CHKPT macro.

You cannot take a checkpoint if a data set with the RLS (record level sharing) option is open.

Catalog Updates

The checkpoint program issues a VSAM CLOSE macro with TYPE=T to update the catalog. It records information about VSAM data sets in a checkpoint data set. If a failure occurs, the latest checkpoint is used to reconstruct the situation that prevailed when the checkpoint was taken.

When using VSAM data sets, a checkpoint is required after any change in the catalog information. A restart will fail if a difference in the checkpoint information and current catalog information is detected.

Note: To determine if the catalog information has changed, examine the contents of the RPL feedback word after each VSAM request. The SHOWCB macro is one way to do this (see *z/OS DFSMS Macro Instructions for Data Sets* for more information).

An RPL feedback word of X'00000004', X'0000002B', or X'0000002C' (see *z/OS DFSMSdfp Diagnosis* for more information) means that EOVS has been called and the catalog information may have been updated. Immediately issue a checkpoint to keep the checkpoint data set current.

Use of Hiperbatch

Hiperbatch is a processing option which provides better performance for jobs accessing VSAM data sets by creating a copy of the data set in expanded storage. Checkpoint/restart processing is supported for VSAM data sets using hiperbatch. No external changes are required when coding the CHKPT macro or when specifying restart processing for VSAM data sets using hiperbatch.

The AMP Parameter

The AMP parameter has a subparameter for specifying checkpoint/restart options that handle two special situations in restarting a processing program:

- Modifications to the data set other than records added sequentially to the end of an entry-sequenced data set. The restart program cannot restore a data set to its checkpoint status if there have been internal modifications to it since the checkpoint, so the restart program does not attempt restart processing.
- Addition of records to the end of a data set by way of a job step other than the job step that issued the checkpoint. Any records added to the end of an entry-sequenced data set are erased in restoring the data set to its checkpoint status.

The AMP options for checkpoint/restart are:

- Let restart take its normal action for both situations
- Override either one or the other of the two actions
- Override both actions.

If you override the check for internal modification, the processing program is restarted, even though the data set being processed cannot be restored. If you override erasing data at the end of a data set and the catalog has been updated, the processing program is not restarted unless you also override the check for modification. For more detail on multiple ACBs open against the same data set, see *z/OS DFSMS Macro Instructions for Data Sets*.

To prevent data from being erased or to allow restart with modified data sets, the CROPS subparameter of the AMP parameter must be coded in the DD statement for the cluster or data set.

When more than one ACB is opened for output for the same data set using DDN sharing and the same DD statement, or using DSN sharing, the first ACB to be opened will establish the CROPS parameter for the data set. If this is not done, the results are unpredictable.

Repositioning is mandatory for all VSAM data sets open for create mode processing, except for relative record data sets processed in direct mode. If AMP='CROPS=NRE' or AMP='CROPS=NRC', no checkpoint is taken. If a checkpoint is attempted, message IHJ000I is issued, and a return code of X'08' is returned in register 15. A reason code of 41 in register 0 is also returned to you.

If a VSAM data set supported for repositioning extends to a new volume after the checkpoint, VSAM restart cannot reposition the data set. A restart from that checkpoint is not successful unless the no-reposition option is taken by specifying AMP='CROPS=NRE' or AMP='CROPS=NRC'.

VSAM Data Set Types

VSAM has key-sequenced, entry-sequenced, and relative record data sets. The main difference among the three is the order in which the data records are loaded into the data sets.

Key-Sequenced Data Sets: Records are loaded into a key-sequenced data set (KSDS) in key sequence, the order defined by the collating sequence of the contents of the key field in each of the records. Each record has a unique value in the key field, such as employee number or invoice number. VSAM uses an index and optional free space to insert a new record into the data set in key sequence.

A key-sequenced data set is prepared for restart by the restoration of any statistical information (such as number of records inserted) to its checkpoint status.

For a VSAM KSDS, restart does not erase any data except in create mode. It does, however, detect modification of the data set by either the checkpointed program or another program that used the data set between the checkpoint and the restart. If the data set is modified, the restart is terminated, unless you override the testing of the data set by using the AMP='CROPS=NCK' subparameter in the DD statement for the data set.

Entry-Sequenced Data Sets: Records are loaded into an *entry-sequenced data set* regardless of the contents of the records. The sequence is determined by the order in which the records are received to be placed in the data set. This is what determines the physical arrangement. New records are stored at the end of the data set.

A checkpoint cannot be taken if a VSAM entry-sequenced data set is open for output with an immediate-upgrade path (or alternate index) open over it, unless the no-reposition option, AMP='CROPS=NRE' or AMP='CROPS=NRC', is specified. VSAM immediate-upgrade data sets are key-sequenced data sets, and repositioning is not supported for them.

For a VSAM entry-sequenced output data set, all data added after the last checkpoint was taken is physically erased by restart unless the AMP='CROPS=NRE' or NRC subparameter is specified in the DD statement for the data set. If data is erased by restart, the catalog record for the data set is updated to reflect the current end of data, and the data set statistics are adjusted to reflect the new status.

During restart, entry-sequenced output data sets are restored by the elimination of all records added at the end since the checkpoint.

Relative Record Data Sets: Records are loaded into a *relative record data set* in relative record number sequence. The data set is a string of fixed-length slots, identified by a relative record number. When a record is inserted, you can assign the relative record number or allow VSAM to assign the record the next available number in sequence. No index is used.

A checkpoint cannot be taken for a VSAM relative record data set (RRDS) in load mode if direct processing was or is performed.

If a checkpoint is taken with a relative record data set open for load mode nondirect processing and direct processing is performed after the checkpoint (thus affecting data that was loaded before the checkpoint was taken), no attempt is made by restart to reset the data that existed at checkpoint time. After restart, the

data set resets to checkpoint status but still contains the results of any direct processing performed on that part of the data set that existed at checkpoint time.

Work Data Sets

Many programs use “work” data sets, which are alternately written and read, rewritten and reread. If a work data set is used, a checkpoint should be taken each time you have finished reading the data set and before rewriting it.

For example, a program may perform the following sequence of operations to produce different versions of data set A:

1. Write and then read back A1.
2. Write and then read back A2.
3. Write and then read back A3.

A checkpoint should be taken at the beginning of operations 2 and 3 before any rewriting of data set A takes place. If, for example, the job step is abnormally terminated while operation 2 is in progress, the job step can be restarted from the checkpoint taken at the beginning of operation 2. At this checkpoint, there is no need for the data in version A1.

Notes for Tape Labels and Files

This section provides information on tape labels and files in relation to checkpoints.

VSE Tape Files

Checkpoints can be taken with VSE tape files opened with the bypass leading tapemark option LABEL=(LTM) and/or the bypass embedded VSE checkpoint records option DCB=(OPTCD=H) specified. However, a checkpoint must not be taken when an opened data set resides on a VSE 7-track tape, is written in translate mode, and contains embedded checkpoint records.

ISO/ANSI Tape Labels

A tape data set with ISO/ANSI labels or with OPTCD=Q cannot be used as a checkpoint data set.

Nonstandard Labels

You cannot restart from a checkpoint that was taken while a tape with nonstandard tape labels is open.

Bypass Label Processing (BLP)

If you are using a multivolume tape data set that is described by concatenated DD statements with the bypass label processing (BLP) option, the deferred option should be coded in the UNIT parameter of those DD statements. This prevents a misleading scheduler premount message that asks for the wrong volume to be mounted at restart time.

Automated Tape Library

During tape installation, checkpoint procedures must be modified when checkpoint tape volumes reside in a tape library automated by a Tape Library Dataserver. A checkpoint indicator will be recorded in the volume's tape configuration data base record, in addition to its file label. The checkpoint volume's tape configuration data

base record indicator will be added and/or deleted whenever the volume changes status. Instead of requesting checkpoint volume confirmation from the operator, the volume's tape configuration data base record will be read. For more information see *z/OS DFSMS Access Method Services Commands*.

When a cartridge is ejected, checkpoint status will be displayed to aid the tape library personnel in executing installation checkpoint volume procedures.

File Protection

The tape operator can file-protect a tape cartridge by manually setting a switch on the cartridge. If it is not in a tape library, a console message might direct the tape operator to set the switch.

The Tape Library Dataserver gripper is unable to manipulate the file protect switch on the cartridge. Automation makes it impossible for a tape operator to change the switch setting of the cartridge. Thus, the read/write capability of the tape cartridge is determined by the switch setting of the tape cartridge at the time it is placed in the library.

When needed, private cartridges can be used within a Tape Library Dataserver and permanently file-protected. This means that they are read-only and any attempt to write will result in a permanent error.

Cartridges can be used within a Tape Library Dataserver and not file-protected by the cartridge switch, but when accessed for input, the logical file protect feature is used. When the OPEN function enables it with a Perform Subsystem Function (PSF) order, this feature causes rejection of all attempts to write to the associated tape drive, for as long as the current volume is mounted.

The following messages will not be issued for tape devices. Instead, the following actions will be taken:

Number	Condition	Action
IEC068A	A tape volume is opened for INPUT, is not file protected and the user is not authorized to write on the volume.	An attempt will be made to logically protect the volume. If that fails, the job will be terminated.
IEC400A	A tape is to be mounted.	This message will be written to the job messages instead of being routed to the consoles.
IEC401A	A tape volume is file protected and opened for OUTPUT or OUTIN.	An attempt will be made to logically unprotect the volume. If that attempt fails or the tape is physically protected, the job will be terminated.
IEC402D	A tape volume is file protected and opened for INOUT.	An attempt will be made to logically unprotect the volume. If that attempt fails, or the tape is physically protected, the job will be terminated.

Note: If you are not using 3495 tape devices see *z/OS MVS System Messages, Vol 1 (ABA-AOM)*, *z/OS MVS System Messages, Vol 2 (ARC-ASA)*, *z/OS MVS System Messages, Vol 3 (ASB-BPX)*, *z/OS MVS System Messages, Vol 4 (CBD-DMO)*, *z/OS MVS System Messages, Vol 5 (EDG-GFS)*, *z/OS MVS System Messages, Vol 6 (GOS-IEA)*, *z/OS MVS System Messages, Vol 7 (IEB-IEE)*, *z/OS MVS System Messages*,

Vol 8 (IEF-IGD), z/OS MVS System Messages, Vol 9 (IGF-IWM), and z/OS MVS System Messages, Vol 10 (IXC-IZP) for actions to take.

Tape Label Verification

If a tape label conflict is encountered while verifying the tape volumes in the Tape Library Dataserver, checkpoint/restart will demount the volume and fail the restart.

For more information about tape labels see *z/OS DFSMS Using Magnetic Tapes*.

Repositioning and Preserving the Contents of Data Sets

The control program repositions data sets, but does not preserve their contents. After taking a checkpoint, you must ensure that data set contents are not changed in such a manner as to make successful postrestart processing impossible.

Recording Data Set Positioning

The checkpoint routine records positioning information for user data sets as follows:

- Data sets are repositioned at restart only if they were open when the checkpoint was taken. The OPEN routine positions all data sets opened after the checkpoint was taken.
- Unit record data sets (printer, punch, or card reader) are not repositioned at restart.
- If you use EXCP to process a reel tape data set open at a checkpoint, ensure that the block count in the data set's data control block is correct. You can do that by coding REPOS=Y in the DCB macro. If the block count is incorrect, the data set may be incorrectly positioned by restart.

If input/output operations were requested, but not begun (for example, if a READ macro instruction was executed, but the related channel program was not started), the checkpoint routine then stops any processing associated with the I/O request, records the positioning information, and reestablishes I/O operations.

If I/O operations have already begun, the checkpoint routine waits until they are complete before recording positioning information.

Buffering of Data Set Records

When QSAM is used to process a data set, an indeterminate number of virtual-storage buffers may contain data when a checkpoint is taken. If restart at a checkpoint occurs, the system's action depends on whether a card reader or another type of device is used to process the data set:

- **Card reader used (QSAM only).** Upon restart, existing buffer contents are released. Restarting from a checkpoint causes buffers to be reprimed by reading records from the current data set into them.
- **Another device used.** Upon restart, the buffer contents are restored to virtual storage, and processing continues normally. It is not possible to predict the time (either before or after the checkpoint) at which a given record will be transferred between a buffer and the recording medium. If your program failed due to a bad record read just before or after taking a checkpoint, then the same failure probably will occur from that checkpoint. This might happen even if you repair the data set after the checkpoint.

If you close a QSAM output data set immediately after restarting the program at a checkpoint, be aware that the data set may not be restored to the same condition it was in when the checkpoint was originally taken.

Preserving Data Set Contents

The system does not save and restore the contents of data sets. You must ensure that input data sets and system data sets contain all necessary data when restart occurs. If a data set on a direct access volume is open at the checkpoint, the data set's label (the DSCB in the VTOC) must have the same location and reflect the same extents upon restart as it did when the checkpoint was taken.

If your program reads records from a data set, updates them, and writes them back to their original locations, it may be useless to take a checkpoint before completing this processing. If a checkpoint is taken earlier, postrestart processing is unsuccessful under these circumstances:

- Your program updates a record before abnormal termination and repeats the update after restart, and
- The updated record contents depend on the original contents.

For example, suppose that the purpose of the update is to exchange the positions of two fields in each record. If the record is updated twice, the fields are returned to their original positions, and the results are invalid. In a different application, an update might place a value in a record field, regardless of the field's original contents. You could then restart the step at a checkpoint taken before or during the update procedure, because an updated record would not be changed if updated again after restart.

When data set records are processed in an update-in-place manner (records are read, changed, and written back into their original location in the data set), bad data can be prevented if records updated after the last checkpoint are restored to their original state or if your program keeps track of the records that are updated and avoids updating them again during restart.

Allocating Devices during Deferred Checkpoint/Restart

When you do a deferred restart from a checkpoint, the type of device allocated for the data set depends on the specification in the UNIT parameter of the DD statement. In addition to assuring the same device type for a checkpoint and restart, the system also attempts to allocate a device with the same optional features present at the time the checkpoint was taken.

- If a device number was specified (for example, UNIT=190), then a device of that type is allocated. It may or may not be the device requested.
- If a device name was specified, then a device of that type is allocated.
- If a user-defined name for a single type of device was specified (for example, UNIT=DISK1), then a device of the defined type is allocated.
- If a name for a mixed group of devices was specified (for example, UNIT=SYSDA), then a device of the same type as that used when the checkpoint was taken is allocated.

However, if the mixed group includes devices with varying optional characteristics (for example with and without RPS or DASD shared and not shared between processors), a device with the same optional characteristics is not guaranteed. To insure that the device allocated by checkpoint/restart has the

same optional characteristics as the device used when the checkpoint was taken, define some generics at system definition time that include only a single group with the same optional characteristics.

In jobs that can restart, you should avoid using generic unitnames or esoteric unitnames (for example, UNIT=TAPE) that contain more than one device type. Allocation failure may result during restart if too few units of a specific device type are available. For example, if UNIT=TAPE includes more than one type of an IBM 3400 Tape Drive (that is 3400-3, 3400-4, and 3400-6), allocation failure may occur if the device type available at restart time is not the same device type as allocated in the original run of the job. To avoid this allocation failure, define some esoteric unitnames at system initialization that include only a single device type.

Handling Checkpoint/Restart Errors

This sections provides information on checkpoint/restart errors.

I/O errors

The checkpoint routine issues return code 0C if it encounters a permanent I/O error when:

- Completing an outstanding VSAM I/O request
- Quiescing queued access method I/O operations
 - An exception occurs when QSAM is used and the skip or accept option is specified in the EROPT operand of the data set's data control block. In this case, code 00 is returned.
- Writing the checkpoint data set.

When an access method other than QSAM is used, your program can ensure that input/output operations are complete before it executes the CHKPT macro instruction; it can thereby avoid having read or written an erroneous record while quiescing.

If a permanent error occurs when the system reads a checkpoint data set to perform a restart, the restart step is terminated abnormally with the system completion code 13F. Further restart of the step is not attempted.

Internal Errors

On return to the caller, register 15 will contain a return code, and register 0 will contain a reason code describing the error. For more information see Appendix A, "Checkpoint/Restart Codes," on page 61.

Depending on the type of error, you may want to consult a SYS1.LOGREC or a SYS1.DUMPxx data set.

User Errors

Checkpoint will make a number of checks to discover user errors. If errors are discovered *before* the first record of the checkpoint entry is written, the checkpoint request is refused; if errors are discovered *after* the first record of the checkpoint entry is written, the checkpoint entry is invalid and checkpoint gives a non-zero return and reason codes.

Using Checkpoint/Restart with Other Programs

This section provides information about using checkpoint/restart with other programs.

Job and Job Step Accounting

The system accumulates processor time used for each job step and job. To access these time values, an installation can provide an accounting routine to receive control at step initiation, step termination, and job termination. Accounting routines are discussed in *z/OS MVS Programming: Authorized Assembler Services Guide*. The relationships between checkpoint/restart and the step time and job time values available to the accounting routine are listed below:

- At termination (either normal or abnormal) of an original execution, the step and job times accumulated are available to the accounting routine.
- At initiation of the restart step during an automatic restart, the step and job times accumulated for the original execution are again available to the accounting routine.
- At initiation of the restart step during a deferred restart, the step and job times are zero.
- At termination of a restart step and at all subsequent times when the accounting routine is given control during the restart execution, the step and job times reflect only the time used during the restart execution.
- Checkpoint/restart restores the TCBUSER field to its value at the time of the checkpoint.

For example, in an original execution, Step A uses 2 minutes of processor time, and Step B uses 3 minutes of processor time and abnormally terminates. At step termination, the step time is 3 minutes and the job time is 5.

- If automatic restart is performed for Step B, a step time of 3 minutes and a job time of 5 are again available to the accounting routine at the reinitiation of Step B.
- If Step B then uses 4 minutes of processor time and terminates, a step time of 4 and a job time of 4 are available to the accounting routine at step termination.

The two values available at the time the restart step is initiated are provided for information purposes only. They are not reflected in the step and job running times presented at termination time of the restarted job. You will not be charged twice for the time accumulated up to the abend.

Another point to consider in your accounting routine is the effect of a restart on the step sequence number available to the accounting routine. The following list indicates the sequence number presented to the accounting routine under the various restart conditions:

Condition	Step Sequence Value for Step n
Original execution	n
Automatic step restart	1
Automatic checkpoint/restart	n+1
Deferred step restart	1
Deferred checkpoint/restart	2

Whenever an automatic restart is performed, the step sequence value accurately reflects the position of the step in the job.

In the case of a deferred restart, the restarting step is the first step of the restart job.

Job Step Time Limit

The EXEC statement TIME parameter specifies a limit on the processor time used by the related step. With any kind of restart, the entire value of the limit specified for the job step applies to the restart step. For a deferred restart, you may specify a limit different from the limit originally specified.

If the processor time used by a step exceeds the specified limit while a checkpoint entry is written, the entry is invalid and an abnormal termination occurs. You cannot do an automatic restart in this case but you can do a deferred checkpoint restart from an earlier checkpoint.

Completion of Step or Job Termination at System Restart

If a step or a job is terminating when system failure occurs, the termination is completed during the system restart that the operator may perform after the failure. This occurs whether or not the step or job uses checkpoint/restart, although the job journal option must be in use.

If other than the last step of a job is terminating when the failure occurs, the termination completes during the system restart, and the next step of the job initiates.

If the last step of a job terminates, or, if the job terminates, all necessary terminations are completed.

If a job requests an automatic restart, then abnormally terminates and system failure occurs before the restart termination processing is complete, the processing completes during the system restart.

JES3

If a job cannot complete executing because of system failure and the job requested job journaling, the job is restarted (warm started) by the system. If the job is eligible for automatic restart, the operator is sent message IEF225D asking if the job should be restarted. If the job is not eligible for automatic restart or if the operator indicated that restart should not be attempted, any temporary or VIO data sets the job allocated are deleted and the job is processed based upon the FAILURE option specified in the MAIN JCL statement. If FAILURE=RESTART is specified, JES3 automatically reschedules the job for execution from the beginning of the first step.

If automatic restart is requested for a job, the job's data sets are set up by JES3 for the first job step and not for the step being restarted. Canceling a job in a dependent network prevents successor jobs from executing if they are dependent upon successful completion of the canceled job. Any operator commands in the input stream of the job step being restarted are not executed. Restart of JES3 controlled jobs may be accompanied by message IAT2006 and/or IAT2575. For responses to these messages see *z/OS JES3 Messages*.

Note: A job that contains references to non-SMS-managed generation data sets that are referenced by a relative generation number of +1 or greater and with a disposition of OLD, SHR, or MOD is failed by the JES3 interpreter service unless the UNIT parameter is included on the DD statement. With UNIT specified, the JES3 interpreter service permits this data set to be allocated on a deferred basis.

COBOL

The COBOL RERUN clause can be used to provide the COBOL user with linkage to checkpoint/restart. Cautions and restrictions on the use of checkpoint/restart also apply to the use of the RERUN clause. For further information see the appropriate COBOL reference manual.

DFSORT™

When performing a sort with the IBM DFSORT™ program, you can, by including the CKPT parameter in sort control statements, cause checkpoint entries to be written and an automatic checkpoint/restart to be requested. The job control language can be used to request automatic or deferred step restarts or a deferred restart at a checkpoint.

PL/I

The PL/I user can use automatic and deferred step restart and can use COMMAND PQUIT to take checkpoints and get automatic and deferred checkpoint restarts. To cause a checkpoint entry to be written and request an automatic checkpoint/restart, the user codes CALL PLICKPT.

Each checkpoint entry in the checkpoint data set is identified by a system-generated checkid. A system message (including the checkid) at the console notifies the operator that a checkpoint entry is written.

The organization of the checkpoint data set is always physical sequential, and the data set may be written on magnetic tape or a direct access volume. Partitioned organization cannot be used.

A DD statement must be present in the job stream to define the checkpoint data set. The DISP parameter in this DD statement is used to specify whether single or multiple checkpoint entries are to be written. DISP=(NEW,KEEP) specifies a single checkpoint entry; DISP=(MOD,KEEP) specifies multiple checkpoint entries.

If CALL PLIREST is used in PL/I programs, the SCHEDxx PARMLIB member must specify 4092 as an eligible user completion code. For more information see "Specification of Codes Eligible for Restart" on page 2.

Virtual Fetch

Modules managed by virtual fetch cannot issue a checkpoint. In addition, job steps that call virtual fetch cannot issue a checkpoint.

Appendix A. Checkpoint/Restart Codes

Table 7 contains directive information about this chapter's contents:

Table 7. Appendix A. Checkpoint/Restart Codes

Section

"Return Codes"

"Reason Codes" on page 62

"Modules and Return Codes for IHJ Messages" on page 86
--

This appendix lists the return codes and reason codes associated with checkpoint/restart functions and describes the status conditions and error messages that may be encountered during normal operation. For more detailed information see the IHJxxxx messages in *z/OS MVS System Messages, Vol 1 (ABA-AOM)*, *z/OS MVS System Messages, Vol 2 (ARC-ASA)*, *z/OS MVS System Messages, Vol 3 (ASB-BPX)*, *z/OS MVS System Messages, Vol 4 (CBD-DMO)*, *z/OS MVS System Messages, Vol 5 (EDG-GFS)*, *z/OS MVS System Messages, Vol 6 (GOS-IEA)*, *z/OS MVS System Messages, Vol 7 (IEB-IEE)*, *z/OS MVS System Messages, Vol 8 (IEF-IGD)*, *z/OS MVS System Messages, Vol 9 (IGF-IWM)*, and *z/OS MVS System Messages, Vol 10 (IXC-IZP)*.

Return Codes

The return codes shown below call attention to various successful or unsuccessful operating conditions.

Return Code

(Hex)

Meaning

00	Successful completion. Code 00 is also returned if the RD parameter was coded as RD=NC or RD=RNC to totally suppress the function of CHKPT.
----	--

04	Automatic checkpoint restart has occurred at the checkpoint taken by the CHKPT macro instruction during the original execution of the job. Another restart from the same checkpoint can still occur. If a deferred restart was performed and RD=NC, NR, or RNC was specified in the resubmitted job, a request for another restart is not in effect.
----	---

The following return codes have reason codes in register 0.

08	Unsuccessful completion. A checkpoint entry was not written, and a restart from this checkpoint was not requested. A request for an automatic restart from a previous checkpoint remains in effect.
----	--

0C

System-related or I/O-related problem:

- An error occurred during the handling of a system request, such as ESTAE, SETLOCK, or PURGE; or
- an I/O error occurred in processing the CHKPT request; or
- A VSAM error was detected while preparing a VSAM data set for the CHKPT request.

This might occur for an attempted checkpoint or restart. If it is a restart, the checkpoint received return code 0, 10, or 18.

Return Code (Hex)	Meaning
10	<p>Successful completion with possible error condition. The task has control, by means of an explicit or implied use of the ENQ macro instruction, of a serially reusable resource; if the task terminates abnormally, it will not have control of the resource when the job step is restarted. Your program must, therefore, restore the enqueues.</p> <p>Additional information regarding explicit and implicit use of the ENQ macro instruction will be found under "Requesting Serially Reusable Resources" on page 20.</p> <p>Be sure that the dynamic link pack area is not altered until you no longer intend to attempt a restart from the identified checkpoint.</p>
14	<p>Unsuccessful completion. Internal error detected.</p> <p>This might occur for an attempted checkpoint or restart. If it is a restart, the checkpoint received return code 0, 10, or 18.</p>
18	<p>Error encountered restoring purged I/O. The checkpoint was successful, but because of the error that was detected, restart may not be possible. In addition, the job now in progress may fail because of I/O errors.</p>

When one of the errors indicated by code 08, 0C, 10, 14, or 18 occurs, the system prints an error message at the operator's console. The message contains a reason code that further identifies the error. The operator should report the contents of the message to the programmer.

Reason Codes

Some reason codes may appear in messages associated with abends 13F, 23F, 33F, 43F, and 53F. Abend codes are indicated by an 'S' in the following table. For more detailed information about the IHJxxxx reason codes see *z/OS MVS System Messages, Vol 1 (ABA-AOM)*, *z/OS MVS System Messages, Vol 2 (ARC-ASA)*, *z/OS MVS System Messages, Vol 3 (ASB-BPX)*, *z/OS MVS System Messages, Vol 4 (CBD-DMO)*, *z/OS MVS System Messages, Vol 5 (EDG-GFS)*, *z/OS MVS System Messages, Vol 6 (GOS-IEA)*, *z/OS MVS System Messages, Vol 7 (IEB-IEE)*, *z/OS MVS System Messages, Vol 8 (IEF-IGD)*, *z/OS MVS System Messages, Vol 9 (IGF-IWM)*, and *z/OS MVS System Messages, Vol 10 (IXC-IZP)*. Also see *z/OS MVS System Codes* for more information about system abend codes.

Some reason code programmer responses are "This is probably an error in checkpoint/restart record management". For each of them, the value and the fact that it represents an error in system logic is part of the programming interface. The nature of the error is not part of the interface.

Reason Code	Return or Abend Code (hex)	Meaning
000	10	<p>Possible enqueue error.</p> <p>Programmer Response: Ensure the program reestablished the enqueues upon restart.</p>
001	08	<p>A CHECKPOINT parameter list error was encountered.</p> <p>Programmer Response: Correct the parameter list.</p>

Reason Code	Return or Abend Code (hex)	Meaning
002	08	The DD statement for the checkpoint data set is missing. Programmer Response: Ensure that the DD statement is correct and is included.
003	08, S13F	Insufficient storage for a checkpoint or restart. Programmer Response: Use a larger storage area. If the IEALIMIT routine has been modified, it may be necessary to alter it to request less storage.
004	0C	Open failed for checkpoint data set. Programmer Response: Locate the problem and correct it.
005	08	The checkpoint data set key length was not equal to zero. Programmer Response: Ensure that the checkpoint data set is not a keyed data set.
006	08	The checkpoint data set record format was not U. Programmer Response: Ensure that the record format is U.
007	08	The data control block (DCB) for the checkpoint data set was opened for other than BSAM or BPAM processing. Programmer Response: Ensure that only BSAM or BPAM is used.
008	08	The time interval specified in a STIMER macro had not elapsed. Programmer Response: Cancel the time interval before issuing the CHKPT macro instruction, and restore the interval after a restart.
009	08	The CHKPT macro instruction was issued in an exit routine other than the end-of-volume exit routine, or an RB of an unacceptable type was found on the RB chain. Programmer Response: Ensure that the CHKPT macro instruction is not issued in an exit routine other than the end-of-volume exit routine.
010	08	An open graphic data control block (DCB) has been detected, but it is not supported in Checkpoint/Restart. Programmer Response: Ensure that the CHKPT macro instruction is not issued while graphic DCBs are open.
011	08	The current task is a subtask or is not represented by the first job step TCB. Programmer Response: Ensure that the CHKPT macro instruction is not issued when multiple tasks (created by the ATTACH macro instruction) exist.
012	08	The current task has a subtask. Programmer Response: Ensure that the CHKPT macro instruction is not issued when multiple tasks (created by the ATTACH macro instruction) exist.

Reason Code	Return or Abend Code (hex)	Meaning
013	08	<p>At the time the checkpoint was attempted, the job had issued the PCLINK macro instruction with the STACK option, but had not issued a corresponding PCLINK macro instruction with the UNSTACK parameter. (That is, the PCLINK stack was not empty.)</p> <p>Programmer Response: Ensure that the PCLINK stack is empty.</p>
014	08	<p>A WTOR macro instruction was issued prior to the CHKPT macro, and a reply had not been received.</p> <p>Programmer Response: Issue a WAIT macro instruction after the WTOR macro instruction, or a DOM macro instruction for any outstanding WTORS before the CHKPT macro instruction.</p>
015	08	<p>An incorrect checkpoint identification length or format was encountered.</p> <p>Programmer Response: When supplying a checkpoint identification, ensure that the length and format are correct.</p>
016	08	<p>The checkpoint data set device type was not magnetic tape or direct access.</p> <p>Programmer Response: Take checkpoints only to tape or DASD.</p>
017	10	<p>Insufficient storage to check enqueues—possible error.</p> <p>Programmer Response: Provide a larger storage area.</p>
018	10	<p>GQSCAN found an abnormal condition.</p> <p>Programmer Response: See the information provided by the GQSCAN service to locate and correct the problem.</p>
019	08	<p>Either MACRF=W was not specified for the checkpoint data control block (DCB), or the data control block was opened by the user's program, but was not opened for output.</p> <p>Programmer Response: Ensure that the DCB for the checkpoint data set has MACRF=W and, if opened, is opened for output.</p>
020	0C	<p>I/O error during open of CHKPT data set.</p> <p>Programmer Response: Locate the cause of the I/O error and correct the problem.</p>
021	08	<p>LABEL was coded in the DD statement for the checkpoint data set with a value other than SL or SUL. Other values for the LABEL parameter must not be used in a checkpoint data set DD statement.</p> <p>Programmer Response: A tape data set to be used for creating a checkpoint data set must have standard labels. Correct the JCL and/or tape label.</p>

Reason Code	Return or Abend Code (hex)	Meaning
022	0C	<p>Error occurred reading or writing a SWA control block.</p> <p>Programmer Response: This is probably an error in the SWA manager. Contact your software support personnel.</p>
023	0C, S13F	<p>An uncorrectable input/output error occurred while writing the checkpoint data set.</p> <p>Programmer Response: Locate the cause of the I/O error and correct it.</p>
024	S13F	<p>A data set has more than five volumes, but there is no JFCBX.</p> <p>Programmer Response: This is probably an error in the SWA manager. Contact your software support personnel.</p>
025	08	<p>A CHKPT request was made by a task that was open to a PDSE, or a PDSE was specified as a checkpoint data set.</p> <p>Programmer Response: Close PDSEs open at the time of checkpoint.</p>
026	0C	<p>Execution of a STOW macro instruction was unsuccessful; a permanent input/output error occurred.</p> <p>Programmer Response: Locate the cause of the I/O error and correct it.</p>
027	08	<p>End-of-volume occurred while writing a checkpoint data set entry on a direct access volume, and secondary allocation was requested and allocated more than once; or end-of-volume occurred twice while writing a checkpoint data set entry on tape.</p> <p>Programmer Response: Ensure that all volumes (or DASD extents) have enough room for more than one checkpoint entry.</p>
028	0C, S33F, 10	<p>I/O error occurred during SYNCDEV to a user data set. The user's data set may not be available for a restart.</p> <p>Programmer Response: Locate the cause of the I/O error and correct it.</p>
029	08	<p>The OPTCD=Q subparameter of the DCB parameter was coded. The value coded is not acceptable for a checkpoint data set. ASCII translation on checkpoint data set is not allowed.</p> <p>Programmer Response: Correct the indicated error and resubmit the job.</p>
030	0C	<p>An uncorrectable I/O error occurred while quiescing pending I/O requests. The I/O error occurred with BSAM or with QSAM EROPT not equal to ACC.</p> <p>Programmer Response: Code EROPT=ACC (on your DD statement or DCB) for QSAM data sets, or locate the cause of the I/O error and correct it.</p>

Reason Code	Return or Abend Code (hex)	Meaning
031	S13F	An uncorrectable I/O error occurred while reading from the checkpoint data set. Programmer Response: Locate the cause of the I/O error and correct it. If all restrictions are met for a deferred restart, you can submit a deferred restart. Either restart at an earlier checkpoint or repeat the original execution using a different volume.
033	08	Cannot reposition to a tape data set or record because the block count in the DCB is negative. This may be caused by the label type being NL or BLP and the data set being open for RDBACK.
034	S13F	A DD statement was missing for the restart step. Programmer Response: Supply the missing DD statement.
035	S13F	Wrong length record detected when reading checkpoint data set. Programmer Response: Locate and correct the problem. Be alert for a possible security violation.
036	0C	Checkpoint received a nonzero return code from the MVS [™] supervisor concerning the number of SSCR records it requires. Programmer Response: See “Modules and Return Codes for IHJ Messages” on page 86 for IHJ Messages.
037	0C	Checkpoint received a nonzero return code after requesting that the MVS supervisor fill in the SSCR records with checkpoint data. Programmer Response: See “Modules and Return Codes for IHJ Messages” on page 86 for IHJ Messages.
038	S13F	Restart received a nonzero return code when it passed the SSCR records containing MVS supervisors' checkpoint data. Programmer Response: See “Modules and Return Codes for IHJ Messages” on page 86 for IHJ Messages.
039	08	A PDS was converted to a PDSE between the time a checkpoint was taken and the time an attempt was made to restart at the checkpoint.
040	S13F	An uncorrectable I/O error occurred while reading standard volume labels. Programmer Response: Locate the cause of the I/O error and correct it.
041	08	CHKPT was issued with a VSAM data set open for create mode processing, and no reposition was specified. Programmer Response Ensure that a checkpoint is not issued when this condition exists, or correct this condition.
042	0C	An error occurred during repositioning for a VSAM data set. Programmer Response: Locate the error and correct it.

Reason Code	Return or Abend Code (hex)	Meaning
043	08	CHKPT was issued with an entry-sequenced data set open for output, an upgrade set open over it, and reposition required. NRE or NRC must be specified in the DD statement. Programmer Response: Locate the error and correct it.
044	08	CHKPT was issued for a relative-record data set open for create mode direct processing. Programmer Response: Ensure that a checkpoint is not issued when this condition exists, or correct this condition.
045	08	CHKPT was issued with a VSAM data set open with the GSR option. Programmer Response: Ensure that a checkpoint is not issued when this condition exists, or correct this condition.
046	08	A checkpoint is not allowed if any of the VSAM data sets in the address space are utilizing the control block update facility processing, or if any of the VSAM data sets in the address space are associated with an alternate index that was using the enhanced share option 4 processing. Programmer Response: Change the data set to be DISP=OLD or close the data set prior to issuing the CHKPT.
047	08	CHKPT was attempted with a VSAM data set that was using the UBF or CBIC option. Programmer Response: Ensure that a checkpoint is not issued when this condition exists, or correct this condition.
048	08	VSAM—media manager was in use. Programmer Response: Remove the CHKPT macro instruction.
049	08	VSAM RLS was in use. Programmer Response: Remove the CHKPT macro instruction.
050	S13F	A volume serial number at restart is not the same as it was when the checkpoint was taken. Programmer Response: Ensure that the volumes used at checkpoint are the same at restart and that units are available to mount them.
051	08	A SPIE or ESPIE exit was active when CHKPT request was made. Programmer Response: Do not issue a CHKPT request from a SPIE or ESPIE exit.

Reason Code	Return or Abend Code (hex)	Meaning
052	S13F	<p>A module that was loaded in the link pack area (LPA) or the nucleus and was being used by the current job when the checkpoint was taken, is not in the LPA or nucleus at restart.</p> <p>Programmer Response: Ensure that the nucleus and LPA are the same at restart as when the checkpoint was taken.</p>
053	S13F	<p>A module has a different entry point address in the link pack area (LPA) or nucleus at restart than it did when the checkpoint was taken.</p> <p>Programmer Response: Ensure that the nucleus and LPA are the same at restart as when the checkpoint was taken.</p>
054	08	<p>A primary address space ID or secondary address space ID is not the same as the home address space ID. Cross-memory addressing was active when a checkpoint was requested.</p> <p>Programmer Response: Do not issue the CHKPT macro when in cross memory mode.</p>
055	0C	<p>Unacceptable return code received from system routine (VSMREGN) used to obtain starting address and size of private area regions.</p> <p>Programmer Response: Information about this problem should have been provided by the VSMREGN service of VSM (virtual storage manager).</p>
056	0C, 14	<p>Unacceptable return code received from system routine (VSMLIST) used to obtain ranges of virtual storage allocated.</p> <p>Programmer Response: Information about this problem should have been provided by the VSMLIST service of VSM (virtual storage manager).</p>
057	08	<p>Checkpoint DD statement is concatenated.</p> <p>Programmer Response: Checkpoint data sets must not be concatenated.</p>
058	14, S13F	<p>Internal parameter list is incorrect.</p> <p>Programmer Response: Use the diagnostic information provided by checkpoint/restart to locate and correct the problem.</p>
059	14, S13F	<p>Internal unexpected return code was received.</p> <p>Programmer Response: Use the diagnostic information provided by checkpoint/restart to locate and correct the problem.</p>
060	14	<p>Scheduler checkpoint routine discovered error in SWA control block.</p> <p>Programmer Response: This is probably an error in the SWA manager. Contact your software support personnel.</p>

Reason Code	Return or Abend Code (hex)	Meaning
061	08	<p>Caller of CHKPT was in secondary addressing mode.</p> <p>Programmer Response: Do not issue the CHKPT macro when in cross memory mode.</p>
062	0C, 14	<p>Checkpoint or Restart detected an internal error while processing checkpoint records.</p> <p>Programmer Response: This is probably an error in the SWA manager. Contact your software support personnel.</p>
063	S13F	<p>I/O error while repositioning a tape to a record within a data set.</p> <p>Programmer Response: Either restart at an earlier checkpoint or repeat the original execution, using a different volume.</p>
064	10	<p>DLPA module detected on load list.</p> <p>Programmer Response: Be sure that the dynamic link pack area is not altered until you no longer intend to attempt a restart from the identified checkpoint.</p>
065	0C	<p>Checkpoint received a nonzero return code from the MVS supervisor due to internal errors.</p> <p>Programmer Response: This is probably an error in the checkpoint/restart MVS supervisor routine. Contact your software support personnel.</p>
066	0C	<p>Restart received a nonzero return code from the MVS Supervisor due to internal errors.</p> <p>Programmer Response: This is probably an error in the checkpoint/restart MVS supervisor routine. Contact your software support personnel.</p>
067	0C	<p>Checkpoint received a nonzero return code from the MVS supervisor because a service routine failed.</p> <p>Programmer Response: This is probably an error in the checkpoint/restart MVS supervisor routine. Contact your software support personnel.</p>
068	0C	<p>Restart received a nonzero return code from the MVS supervisor because a service routine failed.</p> <p>Programmer Response: This is probably an error in the checkpoint/restart MVS supervisor routine. Contact your software support personnel.</p>
069	08	<p>Too much VSMLIST data—user's storage is fragmented beyond checkpoint's capabilities.</p> <p>Programmer Response: Reduce the region size, change the program to reduce storage fragmentation, or increase the size of extended private storage.</p>

Reason Code	Return or Abend Code (hex)	Meaning
070	0C	Checkpoint received a nonzero return code from the MVS Contents Supervisor, because a provider routine return code was not zero. Programmer Response: This is probably an error in the checkpoint/restart MVS supervisor routine. Contact your software support personnel.
071	0C	Restart received a nonzero return code from the MVS contents supervisor, because a provider routine return code was not zero. Programmer Response: This is probably an error in the checkpoint/restart MVS supervisor routine. Contact your software support personnel.
072	0C	Checkpoint received a nonzero return code from the MVS supervisor because a provider routine return code was not zero. Programmer Response: This is probably an error in the checkpoint/restart MVS supervisor routine. Contact your software support personnel.
073	0C	Restart received a nonzero code from the MVS supervisor because a provider routine return code was not zero. Programmer Response: This is probably an error in the checkpoint/restart MVS supervisor routine. Contact your software support personnel.
074	S13F	An unacceptable I/O error occurred while reading a data set control block (DSCB). This reason code can also be issued if a data set was being scratched or deleted. Programmer Response: If it is during restart and the restart was deferred, the restart may be attempted again. If a preceding input/output error message for the restarted job identified the device in error, locate and correct the error. Ensure that no data set has been scratched or deleted.
075	08	STOW encountered a full directory on the checkpoint data set. Programmer Response: Increase the size of the directory for the checkpoint data set.
076	S13F	A direct access data set (other than a spooled data set) being processed for input did not occupy the same extent as it did originally; that is, discrepancies were found in comparing the space allocations described in the data extent block (DEB) and the data set control block (DSCB). Programmer Response: Ensure that no input data set (other than a spooled data set) has been expanded or rewritten in its original volume.

Reason Code	Return or Abend Code (hex)	Meaning
078	S13F	Restart was attempted on a different level of the system than when the checkpoint was taken. Programmer Response: Restart the job in the same processing environment as the one in which it was checkpointed.
079	S13F	DUMMY was specified for an open data set; and either the data set was not being processed by the basic or queued sequential access methods, or the checkpoint at which restart was to occur was established in an end-of-volume exit routine for the data set. Programmer Response: Ensure that DUMMY is specified only for data sets being processed by the queued or basic sequential access methods. Also ensure that restart is not to occur at a checkpoint that was established in an end-of-volume exit routine for a data set that has been made DUMMY.
080	S13F	A subsystem data set was made DUMMY at restart. This usually refers to a spooled data set. Programmer Response: Do not make subsystem data sets DUMMY.
081	S13F	The checkpoint job or step specified ADDRSPC=REAL(V=R) and the restarting job or step did not, or vice versa. Programmer Response: Ensure that the V=R status of both the checkpoint and restarting job or step are the same. That is, if ADDRSPC=REAL was specified at checkpoint time, it can only be specified (and must be specified) at restart time.
082	S13F	The TCAM control program was not active at restart, but a TCAM DCB was open at checkpoint. Programmer Response: Ensure that the TCAM control program region is active and is large enough to add control blocks.
083	S13F	The value of the QNAME keyword parameter on the DD statement is not defined in the TCAM table. Programmer Response: Correct the DD statement or define the QNAME value in the TCAM table.
084	S13F	The process entry defined by the QNAME is already in use by another application program. Programmer Response: Restart the job when the specified QNAME is not being used.
085	S13F	A GETMAIN failed during TCAM restart. Programmer Response: Provide more storage.
086	S13F, 08	An abend occurred during a checkpoint or restart. Programmer Response: Use the diagnostic information provided to correct the error.

Reason Code	Return or Abend Code (hex)	Meaning
087	S13F	Record order on checkpoint data set incorrect and was detected by restart. Programmer Response: Use the diagnostic information provided to correct the error. Be alert for a possible security violation.
088	S13F	Storage not allocated as expected and was detected by restart. Programmer Response: Use the diagnostic information provided to correct the error.
090	14	Checkpoint or restart detected an error internal to checkpoint/restart. Programmer Response: Use the diagnostic information provided to correct the error.
091	S13F, 14	Error building a message. Checkpoint/restart detected an error while building a message for the operator. Programmer Response: Use the diagnostic information provided to correct the error.
092	S13F	An error occurred during basic partitioned access method (BPAM) processing. For example, an error was detected after the convert routine, or after reading the directory blocks. Programmer Response: Locate the problem and correct it.
093	14	An internal parameter list in an MVS supervisor call is incorrect. Programmer Response: This is probably an error in the checkpoint/restart MVS supervisor routine. Contact your software support personnel.
094	S13F	Nucleus routines or tables have moved or been deleted since checkpoint. LPA module entry points, which reside in the nucleus table, have been altered. This may result from a change link pack area (CLPA). Programmer Response: Resubmit the job on the same system with the same nucleus and LPA on which the checkpoint was taken.
095	14, 08	A control block chaining error occurred.
096	S13F	An error occurred during use of the STOW macro instruction to delete a member from the directory of a PDS. Programmer Response: Locate the cause of the error and correct it.
097	S13F	An unacceptable I/O error occurred in the installation's nonstandard label (NSL) routine (NSLREPOS). Programmer Response: Locate the cause of the I/O error and correct it.

Reason Code	Return or Abend Code (hex)	Meaning
098	S13F	<p>Insufficient number of devices were available for restart.</p> <p>Programmer Response: Ensure that the UNIT parameter of the DD statement accurately reflects the device requirements for the data set.</p>
099	S13F	<p>A data set association block (DSAB) is missing at restart. This probably means a DD statement was omitted.</p> <p>Programmer Response: Probable system error. Locate and correct the problem.</p>
100	S13F	<p>MSS error occurred during mount processing. This is a probable system error.</p>
101	S13F	<p>An uncorrectable I/O error occurred while reading a volume label on a direct access device.</p> <p>Programmer Response: Locate the cause of the I/O error and correct it.</p>
103	S13F	<p>A subsystem checkpoint record (SSCR) was written, but the corresponding DCB is not in the open data set table at restart.</p> <p>Programmer Response: Probable error in checkpoint/restart. Contact your software support personnel.</p>
104	S13F	<p>An I/O error occurred while repositioning to the correct data set on tape.</p> <p>Programmer Response: Locate the cause of the I/O error and correct it.</p>
105	S43F	<p>Library Automation services detected an error while processing a tape. See message text following message IHJ014I for detailed information about the error.</p> <p>Programmer Response: Probable system error. Contact your software support personnel.</p>
106	S53F	<p>A Tape Library Dataserver tape volume was mounted, and the internal label serial number did not match the volume serial number of the tape requested.</p> <p>Programmer Response: After verifying that the correct tape is available, resubmit the job.</p>
107	S53F	<p>The file protect status for a Tape Library Dataserver tape volume was incorrect. Restart attempted to reset it using PSF, but encountered an error.</p> <p>Programmer Response: Insure that the tape volume is correctly file protected, then resubmit the job.</p>

Reason Code	Return or Abend Code (hex)	Meaning
108	04	Restart was unable to update statistics (i.e., last mount date and last write date) for a Tape Library Dataserver tape volume, because SMS was not active. Programmer Response: This is a warning indicating that checkpoint/restart was unable to update the tape configuration data base (TCDB) tape volume record; however, the restart has completed successfully.
109	04	Restart was unable to update statistics (i.e., last mount date and last write date) for a Tape Library Dataserver tape volume. An error return code was returned from the SSI call. Programmer Response: This is a warning indicating that checkpoint/restart was unable to update the tape configuration data base (TCDB) tape volume record; however, the restart has completed successfully.
110	04	Restart was unable to update statistics (i.e., last mount date and last write date) for a Tape Library Dataserver tape volume. An error code was returned by SMS. Programmer Response: This is a warning indicating that checkpoint/restart was unable to update the tape configuration data base (TCDB) tape volume record; however, the restart has completed successfully.
112	0C	Checkpoint was unable to load the MVS CSV (Contents Supervisor) routine. Programmer Response: This is probably an error with the CSV routine. Contact your software support personnel.
113	0C	Restart was unable to load the MVS CSV (Contents Supervisor) routine. Programmer Response: This is probably an error with the CSV routine. Contact your software support personnel.
114	08	UNIX [®] system services is active in the current address space. Programmer Response: Checkpoint is not supported for HFS files. Contact your software support personnel.
117	08	The system cannot take the checkpoint because the checkpoint data set was dynamically allocated with the XTIO option, with the DSAB-above-the-line option or without UCB capture. Programmer Response: Remove the CHKPT macro instruction or remove the XTIO, NOCAPTURE and DSAB-above-the-line options from the checkpoint data set allocation.

Reason Code	Return or Abend Code (hex)	Meaning
124	08	<p>A CHKPT request was made by a task that was open to a checkpoint data set with a DCBE.</p> <p>Programmer Response: Ensure that checkpoint is not issued while this condition exists. DCBE's are not supported by checkpoint.</p>
125	08	<p>A CHKPT request was made by a task that was open to a sequential extended format data set, or an extended format data set was specified as a checkpoint data set.</p> <p>Programmer Response: Ensure that checkpoint is not issued while this condition exists, or correct this condition.</p>
126	08	<p>A PDS was converted to an extended format data set between the time a checkpoint was taken and the time an attempt was made to restart at the checkpoint.</p> <p>Programmer Response: Ensure that the restart is not issued while this condition exists, or correct this condition.</p>
181	S13F	<p>An error occurred during VSAM restart preformat routine. <i>yyy</i> is the value from the RPL error code.</p> <p>Programmer Response: The error is for VSAM only. For more information see the RPL error codes in <i>z/OS DFSMS Macro Instructions for Data Sets</i> for VSAM data sets.</p>
182	S13F	<p>An error occurred during VSAM restart verify routine. <i>yyy</i> is the value from the RPL error code.</p> <p>Programmer Response: The error is for VSAM only. For more information see the RPL error codes in <i>z/OS DFSMS Macro Instructions for Data Sets</i> for VSAM data sets.</p>
183	S13F	<p>An error occurred during VSAM restart put routine. <i>www</i> is the value from the RPL error code.</p> <p>Programmer Response: The error is for VSAM only. For more information see the RPL error codes in <i>z/OS DFSMS Macro Instructions for Data Sets</i> for VSAM data sets.</p>
184	S13F	<p>An error occurred during VSAM restart index put routine. <i>www</i> is the value from the RPL error code.</p> <p>Programmer Response: The error is for VSAM only. For more information see the RPL error codes in <i>z/OS DFSMS Macro Instructions for Data Sets</i> for VSAM data sets.</p>
185	S13F	<p>VSAM open error during restart. <i>www</i> is the value from the ACB error code.</p> <p>Programmer Response: See message IEC161I or IEC252I.</p>
186	0C	<p>VSAM close error. <i>www</i> is the value from the ACB error code.</p> <p>Programmer Response: The message is for VSAM only. See message IEC161I or IEC252I.</p>

Reason Code	Return or Abend Code (hex)	Meaning
190	S13F	<p>An error occurred while attempting to obtain catalog information for the cluster identified by the ddname in message IHJ0091.</p> <p>Programmer Response: The error is for VSAM only. For more information, see messages IEC331I and IDA3009I, if issued.</p>
191	S13F	<p>An error occurred while attempting to mount volumes for the ddname identified by message IHJ0091.</p> <p>Programmer Response: Increase the number of units allocated for the ddname.</p> <p>The error is for VSAM only. For more information, see message IEC161I or IEC252I, if issued.</p>
193	S13F	<p>Checkpoint was taken during VSAM create mode, but, at restart time, the data set was no longer in create mode.</p> <p>Programmer Response: A restart is not allowed. Take the checkpoint after the data set is loaded.</p>
194	S13F	<p>The catalog entries for a VSAM upgrade path have been altered since checkpoint.</p> <p>Programmer Response: If the upgrade path can be restored to its condition at checkpoint, the restart may be resubmitted.</p>
195	S13F, 0C	<p>Insufficient storage is available for a VSAM restart.</p> <p>Programmer Response: Increase the available storage.</p>
196	S13F	<p>The catalog for one or more VSAM data sets has been updated, indicating that another job successfully processed the data set after the checkpoint was taken. Restart cannot restore the data set to its checkpoint status.</p> <p>Programmer Response: Override the condition, and allow restart by coding AMP='CROPS=NCK' in the DD statement for the data set. Be aware, however, that changes might have been made to the data between the time the checkpoint was taken and when restart is attempted. A data set that has been extended to span volumes since the checkpoint was taken cannot be restarted.</p> <p>The error is for VSAM only. For more information, see messages IEC161I or IEC252I, if issued.</p>
197	S13F	<p>A restart is not allowed for control block update facility (CBUF) processing.</p> <p>Programmer Response: Do not issue the CHKPT macro instruction while using CBUF processing.</p>

Reason Code	Return or Abend Code (hex)	Meaning
198	S13F	<p>A VSAM data set extended to a new volume after the checkpoint was taken.</p> <p>Programmer Response: The checkpoint cannot be restarted. For a deferred restart, resubmit the job selecting a checkpoint taken after the data set extended to the new volume.</p> <p>The error is for VSAM only. For more information, see message IEC161I or IEC252I, if issued.</p>
199	S13F	<p>An error occurred during BLDVRP processing in VSAM restart. <i>www</i> is the return code from the BLDVRP request.</p> <p>Programmer Response: The error is for VSAM only. For more information see <i>z/OS DFSMS Macro Instructions for Data Sets</i>.</p>
200	0C	<p>A PURGE I/O request failed.</p> <p>Programmer Response: Locate and correct the error.</p>
201	S13F	<p>Page Fix (PGSER) failed. This is an uncorrectable error.</p> <p>Programmer Response: Probable system error. Contact your software support personnel.</p>
202	0C, S13F	<p>SETLOCK failed. This is an uncorrectable error.</p> <p>Programmer Response: Probable system error. Contact your software support personnel.</p>
203	S13F	<p>An invalid attempt has been made to invoke RESTART SVC routines.</p> <p>Programmer Response: Only the scheduler may invoke RESTART. Remove either the JCL or program logic that attempts to execute RESTART.</p>
204	0C	<p>VIO data sets could not be checkpointed for a subsequent automatic restart. This checkpoint is unavailable for automatic restarts; however, deferred restarts are possible if the VIO data sets are made DUMMY.</p> <p>Programmer Response: An error occurred while checkpointing VIO data sets to the job journal. Deferred restarts may be accomplished on the checkpoint entry; however, the VIO data sets must be made DUMMY. Locate and correct the error.</p>
205	0C	<p>VIO data sets could not be checkpointed for a subsequent automatic restart. Any checkpoints taken during this step are ineligible for automatic restarts; however, deferred restarts are possible if the VIO data sets are made DUMMY. An automatic restart cannot occur until a new and successful checkpoint is taken.</p> <p>Programmer Response: An error occurred while checkpointing VIO data sets to the job journal. Deferred restarts may be accomplished on the checkpoint entry; however, the VIO data sets must be made DUMMY. Locate and correct the error.</p>

Reason Code	Return or Abend Code (hex)	Meaning
206	08	<p>A VTAM[®] ACB was open when a CHKPT was attempted.</p> <p>Programmer Response: Ensure that all VTAM[®] ACBs are closed prior to issuing CHKPT.</p>
207	0C, S13F	<p>Failure in the subsystem interface when checkpointing or restarting subsystem data sets.</p> <p>Programmer Response: Probable subsystem error. Contact your software support personnel.</p>
208	08	<p>Erroneous data exists in tape for the new checkpoint data set presented by the user. The integrity of this data set is uncertain and therefore unusable.</p> <p>Programmer Response: If the new checkpoint data set is a sequential data set, the failure occurred because the data set was open for MOD and residual data was found. Either remove the MOD disposition, reinitialize the data set to an empty status, or reallocate the data set. If the data set is a PDS, the failure occurred as a result of residual members in the directory. Either scratch the members or reallocate the data set.</p>
209	08	<p>Another user DCB is concurrently open to the checkpoint data set presented by the user.</p> <p>Programmer Response: Remove the code that violates this restriction, or change the DD statement for the checkpoint data set.</p>
210	08	<p>The checkpoint data set presented by the user was defined with DISP=SHR.</p> <p>Programmer Response: Change the JCL to reflect a disposition other than SHR for the checkpoint data set.</p>
211	08	<p>An attempt to create a new checkpoint data set has failed because it resides on an insecure volume.</p> <p>Programmer Response: The candidate volume for a checkpoint data set was rejected by the operator as not being secure. Contact installation personnel and verify the eligibility of the volume for creating checkpoint data sets.</p>
213	08	<p>The checkpoint data set offered by the user is a subsystem data set. Subsystem data sets cannot be used as checkpoint data sets.</p> <p>Programmer Response: Correct the JCL and resubmit the job.</p>
214	08	<p>The new checkpoint data set offered by the user exists on a shared direct access device. Concurrent opens cannot be controlled; therefore, the data set is unacceptable.</p> <p>Programmer Response: Contact operations personnel to ensure that this does not occur. See also "Coding the DD Statement for a Checkpoint Data Set" on page 8.</p>

Reason Code	Return or Abend Code (hex)	Meaning
215	18	An error occurred while restoring the purged I/O operations. A checkpoint entry was successfully written; however, continuation of the user's program may produce errors. Programmer Response: Locate and correct the error.
216	S13F	Residual subpool allocations in the user's region prevented a successful reallocation of the user's checkpoint-time region image. Programmer Response: Locate and correct the error. Possible system problem.
219	S13F	There is a password error for a password-protected tape data set. Programmer Response: Obtain the proper password for the password-protected data set.
220	S13F	The data set control block (DSCB) address of a data set to be repositioned has changed since the checkpoint was taken. Programmer Response: If possible, return the volume to the condition it was in at checkpoint.
221	S13F	The wrong password was given for a password-protected data set that was to be repositioned. Programmer Response: Obtain the proper password for the password-protected data set.
222	S13F	The tape data set name, recorded in header label 1, has changed since the checkpoint was taken. Programmer Response: If possible, return the volume to the condition it was in at checkpoint.
224	08	An open data set has a type of ACB that checkpoint does not support. Checkpoint supports VSAM ACB's and an ACB that the system builds when the user opens a DCB for a subsystem data set, such as spooled data sets. Programmer Response: Ensure that all data sets of this type are closed prior to issuing the checkpoint.
240	0C, S13F	A recovery (ESTAE) environment could not be established. Programmer Response: Probable system error. Contact your software support personnel.
241	0C, S13F	An indeterminate error occurred when processing VSAM data sets. Programmer Response: Locate and correct the error using the diagnostic information provided.
242	0C, S13F	A machine check occurred while processing VSAM data sets. Programmer Response: If the machine check was transitory, resubmit the job. Otherwise, have the problem fixed before resubmitting the job.

Reason Code	Return or Abend Code (hex)	Meaning
243	S13F	An invalid SSCR record was encountered in the checkpoint data set entry when repositioning VSAM data sets. This is a probable system logic error. Programmer Response: Resubmit the job, selecting another checkpoint entry. Be alert for a possible security violation.
250	08	SYS1.IMAGELIB data set open with the IMGLIB macro when checkpoint taken. Programmer Response: Ensure that all IMAGELIB DCBs are closed prior to issuing CHKPT. See also “IMGLIB” on page 18.
251	S23F	A data set that was not a checkpoint data set at checkpoint time was found to be open as a secure checkpoint data set at restart time. Programmer Response: Ensure the proper data set volume was mounted at restart time. Also, verify the status of the subject volume by contacting the operations personnel. Be alert for a possible security violation.
252	S23F	In a job using more than one checkpoint data set, one of the checkpoint data sets (not the one used for restart) was deemed not secure. Programmer Response: Ensure the proper data set volume was mounted at restart time. Also verify the status of the subject volume by contacting the operations personnel. Be alert for a possible security violation.
255	S13F, S23F	Security violation—user not authorized to access a RACF [®] protected data set at restart. Programmer Response: See IHJ101I or IHJ102I for information. Be alert for a possible security violation.
513	14	An internal error was detected for a bad parameter list identifier. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
514	14	An internal error was detected for an incorrect parameter list level. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
515	14	An internal parameter list is too short. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
516	14	An internal storage list has been invalidated. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.

Reason Code	Return or Abend Code (hex)	Meaning
517	14	More than one internal requester flag was selected. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
518	14	A reserved field in an internal parameter list is not zeros. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
519	14	An invalid request code was passed in an internal parameter list for record management processing. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
520	14	A write request was made with an invalid record address. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
521	14	A write request was made with an invalid record length. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
522	14	No record header address was provided in a write request and the record ID is for a record type that requires a separate record header. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
523	14	No record header length was provided in a write request when a record header was provided. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
524	14	The caller of an internal request is an input function and requested an output operation or vice versa. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
526	14	A DCB address was not passed on an OPEN for output. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.

Reason Code	Return or Abend Code (hex)	Meaning
527	14	No length was specified on a write request for a record appendage. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
528	14	Error positioning to the checkpoint data set. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
529	14	An OPEN for input or output must precede any other operation requests. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
530	20	An invalid record type was found in the checkpoint data set. Programmer Response: This is probably an error in checkpoint/restart record management. Be alert for a possible security violation. Contact your software support personnel.
531	20	An unrecognizable record code was found in the checkpoint data set. Programmer Response: This is probably an error in checkpoint/restart record management. Be alert for a possible security violation. Contact your software support personnel.
533	12	A checkpoint/restart GETMAIN request received a nonzero return code. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
537	12	An attempt to open the checkpoint data set failed. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
538	20	A buffer with data was found when none was expected. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
539	20	A buffer that should have been full was found to be only partially filled. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.

Reason Code	Return or Abend Code (hex)	Meaning
540	14	The storage list pointer in an internal parameter list is zero. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
541	14	An attempt to do a DEB validity check failed. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
542	14	The record code passed for a write request is an end-of-block code. This code can not be passed as a record code. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
543	14	A request to write a record included a record too big for a buffer and of a type that can not be split. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
544	08	A checkpoint ID was specified that is not a member of the checkpoint PDS. Programmer Response: This is probably an error in checkpoint/restart record management. Possible user error. Contact your software support personnel.
545	08	An OPEN for input was attempted to a data set that is not a secure checkpoint data set. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
546	08	A requested checkpoint ID could not be found in the checkpoint data set. Programmer Response: This is probably an error in checkpoint/restart record management. Possible user error. Contact your software support personnel.
547	20	A buffer that should have been empty was marked full. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
548	08	A volume switch occurred while writing a checkpoint entry in the checkpoint data set. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.

Reason Code	Return or Abend Code (hex)	Meaning
549	14	A request to write an SSCR record was passed with an incorrect header length. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
550	14	A request to write an SSCR record was passed with nonzero values in reserved fields. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
551	14	A request to write an SSCR record was passed with its data length set to zero. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
552	14	A request to write an SSCR record was passed with the DCB address set to zero. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
553	14	A request to write a record was passed with an unknown record code. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
554	14	The DCB passed for OPEN for output was not a checkpoint protected DCB. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
555	08	The ddname passed in an open for input request is not allocated. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
556	12	On a call to SWA to read a JFCB, a nonzero return code was received. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
557	14	No ddname was passed on an OPEN for input request by scheduler restart. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.

Reason Code	Return or Abend Code (hex)	Meaning
558	14	No checkpoint ID address was passed on an OPEN for input request by scheduler restart. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
559	14	A zero length was passed for the checkpoint ID by scheduler restart. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
560	14	No checkpoint header record address was passed on an OPEN for input request by scheduler restart. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
561	14	A zero length was passed for the checkpoint header record by scheduler restart. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
562	12	An error occurred while calculating block size for the checkpoint data set. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
563	12	An optimum block size could not be calculated for the checkpoint data set. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
564	08	An attempt was made by scheduler restart to open a checkpoint data set that was created on a release of DFP or MVS/SP that did not have complete scheduler checkpoint support. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
565	12	A nonzero return code was returned from an ESTAE service request. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.

Reason Code	Return or Abend Code (hex)	Meaning
566	08	The extents of a DASD checkpoint data set have been extended. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.
567	14	A record split was requested for a record type that cannot be split. Programmer Response: This is probably an error in checkpoint/restart record management. Contact your software support personnel.

Modules and Return Codes for IHJ Messages

When Checkpoint/Restart issues messages IHJ001I and IHJ007I with reason codes 036, 037, or 038, variables **m** and **rc** might appear. **m** designates a module id, which corresponds with a module name; **rc** designates a return code issued by module **m**. The return codes are in decimal format.

Module ID	Return Code	Meaning
01		Module IEAVCKRS—Supervisor Control Checkpoint/Restart Router.
	08	The number of requested SSCRs for the second checkpoint invocation does not match the number requested in the first checkpoint invocation.
02		Module IEAVSPI—SPIE/ESPIE Checkpoint/Restart Routine.
	08	ESPIE exits were created while the task was running with a PSW key other than the assigned key.
	255	IEAVSPI requested a negative number of SSCRs.
03		Module IEAVCRVF—Vector Checkpoint/Restart Routine.
	04	The pointer to the subject task's VSSA is invalid.
	08	The vector SLIH determined that there are no eligible vector processors with an attached Vector Facility (VF) online.
	12	The vector SLIH could not obtain a VSSA for the task.
	16	The VF section size recorded during the checkpoint is not equal to the VF section size during the restart.
	20	In the time interval between checkpoint and the restart, the mapping of the VSSA changed.
	24	There are no processors with an attached VF in the complex.
	28	An unexpected error occurred in IEAVCRVF.
04	32	VF status has been lost because of a machine check.
	255	IEAVCRVF requested a negative number of SSCRs.
		Module IEAVLSCR—Linkage Stack Checkpoint/Restart Routine.
05	08	There are entries on the task's linkage stack.
		Module IEAVALCR—Access List Checkpoint/Restart Routine.

Module ID	Return Code	Meaning
	08	The task has entries in use on either the dispatchable unit access list (DU-AL) or on the primary address space access list (PASN-AL). That is, the task has established addressability to another address/data space via an access list.
	0C	An unexpected error occurred during IEAVALCR processing.
06		Module IEANTCKR-Name/Token checkpoint/restart extension routine.
	08	Address space owns name/token pairs that were not created with the IEANT_CHECKPOINTOK value in the persist parameter. Remove the definition before taking the checkpoint or add the IEANT_CHECKPOINTOK option.

Appendix B. End-of-Volume Checkpoint Routine

Table 8 contains directive information about this chapter's contents:

Table 8. Appendix B. End-of-Volume Checkpoint Routine

Section
Appendix B, "End-of-Volume Checkpoint Routine"

In a data control block exit list, you can specify the address of a routine that receives control when end-of-volume is reached in processing a physical sequential data set (BSAM or QSAM). For information about forming an exit routine and coding the EOVS exit list for physical sequential data sets, see *z/OS DFSMS Using Data Sets*. The routine is entered after a new volume is mounted and all necessary label processing is completed. If the volume is magnetic tape, the tape is positioned after the tape mark that precedes the beginning of the data. The end-of-volume exit routine may take a checkpoint by issuing the CHKPT macro instruction. The job step can be restarted from this checkpoint. When the job step is restarted, the volume is mounted and positioned as upon entry to the routine.

The end-of-volume exit routine returns control the same way as any other data control block exit routine. Restart is impossible if changes are made to the link pack area or nucleus after the checkpoint is taken. When the step is restarted, the virtual storage addresses of end-of-volume modules must be the same as when the checkpoint was taken.

Appendix C. Accessibility

Accessible publications for this product are offered through the z/OS Information Center, which is available at www.ibm.com/systems/z/os/zos/bkserv/.

If you experience difficulty with the accessibility of any z/OS information, please send a detailed message to mhvrcfs@us.ibm.com or to the following mailing address:

IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Building 707
2455 South Road
Poughkeepsie, NY 12601-5400
USA

Accessibility features

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 such products 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 Vol 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.

Dotted decimal syntax diagrams

Syntax diagrams are provided in dotted decimal format for users accessing the z/OS Information Center using a screen reader. In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), they can appear on the same line, because they can be considered as a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that your screen reader is set to read out punctuation. All the syntax elements that have the same dotted decimal number (for example, all the syntax elements that have the number 3.1) are mutually

exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, you know that your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, it is preceded by the backslash (\) character. The * symbol can be used next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is given the format 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol giving information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, this indicates a reference that is defined elsewhere. The string following the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you should refer to separate syntax fragment OP1.

The following words and symbols are used next to the dotted decimal numbers:

- ? means an optional syntax element. A dotted decimal number followed by the ? symbol indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element, (for example 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that syntax elements NOTIFY and UPDATE are optional; that is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.
- ! means a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicates that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the same dotted decimal number can specify a ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In this example, if you include the FILE keyword but do not specify an option, default option KEEP will be applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1!

(KEEP), and 2.1.1 (DELETE), the default option KEEP only applies to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.

- * means a syntax element that can be repeated 0 or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1* data area, you know that you can include one data area, more than one data area, or no data area. If you hear the lines 3*, 3 HOST, and 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

Note:

1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
 2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you could write HOST STATE, but you could not write HOST HOST.
 3. The * symbol is equivalent to a loop-back line in a railroad syntax diagram.
- + means a syntax element that must be included one or more times. A dotted decimal number followed by the + symbol indicates that this syntax element must be included one or more times; that is, it must be included at least once and can be repeated. For example, if you hear the line 6.1+ data area, you must include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE, you know that you must include HOST, STATE, or both. Similar to the * symbol, the + symbol can only repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loop-back line in a railroad syntax diagram.

Notices

This information was developed for products and services offered in the U.S.A. or elsewhere.

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 character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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:

Site Counsel
IBM Corporation
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.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

COPYRIGHT LICENSE:

This information might contain sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Policy for unsupported hardware

Various z/OS elements, such as DFSMS, HCD, JES2, JES3, and MVS, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted

for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

Minimum supported hardware

The minimum supported hardware for z/OS releases identified in z/OS announcements can subsequently change when service for particular servers or devices is withdrawn. Likewise, the levels of other software products supported on a particular release of z/OS are subject to the service support lifecycle of those products. Therefore, z/OS and its product publications (for example, panels, samples, messages, and product documentation) can include references to hardware and software that is no longer supported.

- For information about software support lifecycle, see: IBM Lifecycle Support for z/OS (<http://www.ibm.com/software/support/systemsz/lifecycle/>)
- For information about currently-supported IBM hardware, contact your IBM representative.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml (<http://www.ibm.com/legal/copytrade.shtml>).

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.

ABSTR

Absolute track (value of SPACE).

ACB Access method control block.

ACS See *automatic class selection*.

ADSP See *automatic data set protection*.

AL American National Standard Labels.

allocation

Generically, the entire process of obtaining a volume and unit of external storage, and setting aside space on that storage for a data set. Also refers to the process of logically connecting a program and a data set. The program is running or is about to run.

AMODE

Addressing mode (24, 31, ANY).

ANSI American National Standards Institute.

AOR Application owning region.

APF Authorized program facility.

application

The use to which an access method is put or the end result that it serves; contrasted to the internal operation of the access method.

ASCII American National Standard Code for Information Interchange.

AUL American National Standard user labels (value of LABEL).

ATL See *automated tape library*.

automated tape library data server

A device consisting of robotic components, cartridge storage areas, tape subsystems, and controlling hardware and software, together with the set of tape volumes that reside in the library and can be mounted on the library tape drives. Contrast with *manual tape library*. See also *tape library*.

automatic class selection (ACS) routine

A procedural set of ACS language statements. Based on a set of input variables, the ACS language statements generate the name of a predefined SMS class, or a list of names of predefined storage groups, for a data set.

automatic data set protection (ADSP)

In z/OS, a user attribute that causes all permanent data sets created by the user to be automatically defined to RACF with a discrete RACF profile.

- AVGREC** Average record scale (JCL keyword).
- basic format**
The format of a data set that has a data set name type (DSNTYPE) of BASIC. A basic format data set is a sequential data set that is specified to be neither large format nor extended format. The size of a basic format data set cannot exceed 65 535 tracks on each volume.
- BCDIC** Binary coded decimal interchange code.
- BCS** Basic catalog structure.
- BDAM** Basic direct access method.
- BDW** Block descriptor word.
- BFALN** Buffer alignment (parameter of DCB).
- BFTEK** Buffer technique (parameter of DCB).
- BISAM** Basic index sequential access method.
- BLP** Bypass label processing.
- BLKSIZE** Block size (parameter of DCB).
- BLT** Block locator token.
- BPAM** Basic partitioned access method.
- BPI** Bytes per inch.
- BSAM** Basic sequential access method.
- BSM** Backspace past tape mark and forward space over tape mark (parameter of CNTRL).
- BSP** Backspace one block (macro).
- BSR** Backspace over a specified number of blocks or records (parameter of CNTRL).
- BUFC** Buffer control block.
- BUFCB** Buffer pool control block (parameter of DCB).
- BUFL** Buffer length (parameter of DCB).
- BUFNO** Buffer number (parameter of DCB).
- BUFOFF** Buffer offset (length of ASCII block prefix by which the buffer is offset; parameter of DCB).
- CA** Control area.
- catalog**
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. See *master catalog*.
- CBIC** Control blocks in common.
- CBUF** Control block update facility.
- CCHHR** Cylinder/head/record address.
- CCW** Channel command word.
- CF** Coupling facility.
- CI** Control interval.
- CIDF** Control interval definition field.
- CKDS** Cryptographic key data set.
- CLPA** Change link pack area.
- CONTIG** Contiguous space allocation (value of SPACE).
- CSA** Common storage area.
- CSW** Channel status word.
- CYLOFL** Number of tracks for cylinder overflow records (parameter of DCB).
- DA** Direct access (value of DEVD or DSORG).
- DADSM** Direct access device space management.
- DASD** Direct access storage device.
- DATACLAS** Data class (JCL keyword).
- data control block (DCB)**
A control block used by access method routines (other than VSAM or VTAM) in storing and retrieving data.
- Data Facility Storage Management Subsystem (DFSMS)**
An operating environment that helps automate and centralize the management

- of storage. To manage storage, SMS provides the storage administrator with control over data class, storage class, management class, storage group, and automatic class selection routine definitions.
- Data Facility Storage Management Subsystem data facility product (DFSMSdfp)**
A DFSMS functional component and a base element of z/OS that provides functions for storage management, data management, program management, device management, and distributed data access.
- data security**
Prevention of access to or use of data or programs without authorization. As used in this publication, the safety of data from unauthorized use, theft, or purposeful destruction.
- DAU** Direct access unmovable data set (value of DSORG).
- DBCS** Double-byte character set.
- DCB** See *data control block* (control block name, macro, or parameter on DD statement).
- DCBD**
Data control block dummy section.
- DCBE** Data control block extension.
- DD** Data definition.
- DDM** Distributed Data Management.
- DEB** Data extent block.
- DECB** Data event control block.
- DEN** Magnetic tape density (parameter of DCB).
- DES** Data encryption standard.
- DEV D**
Device-dependent (parameter of DCB).
- 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.
- DISP** Disposition (JCL DD parameter).
- DIV** Data-in-virtual.
- DLF** Data lookaside facility.
- DSCB** Data set control block.
- DSORG**
Data set organization (parameter of DCB and DD and in a data class definition).
- EBCDIC**
Extended binary coded decimal interchange code.
- ECB** Event control block.
- ECSA** Extended common service area.
- ENABLE**
Relationship that allows a system to allocate and access data sets in a VIO storage group, a pool storage group, or individual volumes within a pool storage group.
- entry** A collection of information about a cataloged object in a master or user catalog.
- entry-sequenced data set**
A data set whose records are loaded without respect to their contents, and whose RBAs cannot change. Records are retrieved and stored by addressed access, and new records are added at the end of the data set.
- EOB** End-of-block.
- EOD** End-of-data.
- EODAD**
End-of-data set exit routine address (parameter of DCB, DCBE, and EXLST).
- EOV** End-of-volume.
- ESDS** See *entry-sequenced data set*.
- ESETL**
End-of-sequential retrieval (QISAM macro).
- ESPIE** Extended specify program interruption exits.
- exclusive control**
Preventing multiple WRITE-add requests from updating the same dummy record or writing over the same available space on a track. When specified by the user, exclusive control requests that the system prevent the data block about to be read from being modified by other requests; it

is specified in a READ macro and released in a WRITE or RELEX macro. When a WRITE-add request is about to be processed, the system automatically gets exclusive control of either the data set or the track.

EXCP Execute channel program.

EXLST
Exit list (parameter of DCB and VSAM macros).

EXPDT
Expiration date for a data set (JCL keyword).

export To create a backup or portable copy of a VSAM cluster, alternate index, or integrated catalog facility user catalog.

FCB See *forms control buffer*.

FEOV Force end-of-volume (macro).

file system

In the z/OS UNIX environment, the collection of files and file management structures on a physical or logical mass storage device, such as a diskette or minidisk. See also *HFS data set*.

FOR File owning region.

forms control buffer (FCB)
A buffer for storing codes for control of certain printers such as the 3205-5, 3211, 3262 Model 5, or IBM 3800 Printing Subsystem.

FSM Forward space past tape mark and backspace over tape mark (parameter of CNTRL).

FSR Forward space over a specified number of blocks or records (parameter of CNTRL).

GCR Group coded recording (tape recording).

GDG See *generation data group*.

GDGNT
Generation data group name table.

GDS Generation data set.

generation data group (GDG)
A collection of data sets with the same base name, such as PAYROLL, that are kept in chronological order. Each data set is called a generation data set.

GL GET macro, locate mode (value of MACRF).

global resource serialization (GRS)
A component of z/OS used for serializing use of system resources and for converting hardware reserves on DASD volumes to data set enqueues.

GM GET macro, move mode (value of MACRF).

GRS See *global resource serialization*.

GSR Global shared resources.

GTF Generalized trace facility.

HCD Hardware configuration dialog.

HFS See *hierarchical file system*.

hierarchical file system (HFS) data set
A data set that contains a POSIX-compliant file system, which is a collection of files and directories organized in a hierarchical structure, that can be accessed using z/OS UNIX System Services. See also *file system*.

IBG Interblock gap.

ICFC Integrated catalog facility catalog.

ICI Improved control interval access.

import

To restore a VSAM cluster, alternate index, or integrated catalog facility catalog from a portable data set created by the EXPORT command.

IPL Initial program load.

INOUT
Input then output (parameter of OPEN).

I/O Input/Output.

IOB Input/output block.

IRG Interrecord gap.

IS Indexed sequential (value of DSORG).

ISAM Indexed sequential access method.

ISMF Interactive Storage Management Facility.

ISO International Organization for Standardization.

ISU Indexed sequential unmovable (value of DSORG).

JCL Job control language.

JES Job entry subsystem.

JFCB Job file control block.

JFCBE Job file control block extension.

KEYLEN

Key length (JCL and DCB keyword).

key-sequenced data set (KSDS)

A VSAM data set whose records are loaded in ascending key sequence and controlled by an index. Records are retrieved and stored by keyed access or by addressed access, and new records are inserted in key sequence because of free space allocated in the data set. Relative byte addresses of records can change because of control interval or control area splits.

KSDS See *key-sequenced data set*.

large format

The format of a data set that has a data set name type (DSNTYPE) of LARGE. A large format data set has the same characteristics as a sequential (non-extended format) data set, but its size on each volume can exceed 65 535 tracks. There is no minimum size requirement for a large format data set.

link-pack area (LPA)

An area of virtual storage that contains re-entrant routines that are loaded at IPL (initial program load) time and can be used concurrently by all tasks in the system.

LPA See *Link-pack area*.

LRECL

Logical record length (JCL keyword).

LRI Logical record interface.

LSR Local shared resources.

MACRF

Macro instruction form (parameter of DCB and ACB).

master catalog

A catalog that contains extensive data set and volume information that VSAM requires to locate data sets, to allocate and deallocate storage space, to verify the authorization of a program or operator to gain access to a data set, and accumulate usage statistics for data sets.

MBBCHHR

Module number, bin number, cylinder number, head number, record number.

MCP Message control program.

MGMTCLAS

Management class (JCL keyword).

MLA Multilevel alias.

MOD Modify data set (value of DISP).

MSHI Main storage for highest-level index (parameter of DCB).

MSWA

Main storage for work area (parameter of DCB).

NCI Normal control interval.

NCP Number of channel programs (parameter of DCB).

NOPWREAD

No password required to read a data set (value of LABEL).

NRZI Nonreturn-to-zero-inverted.

NSL Nonstandard label (value of LABEL).

NSR Nonshared resources.

NTM Number of tracks in cylinder index for each entry in lowest level of master index (parameter of DCB).

NUB No user buffering.

NUP No update.

OPTCD

Optional services code (parameter of DCB).

OUTIN

Output then input (parameter of OPEN).

OUTINX

Output at end of data set (to extend) then input (parameter of OPEN).

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 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.

PDAB Parallel data access block.

PDS See *partitioned data set*.

- PDSE** See *partitioned data set extended*.
- PE** Phase encoding (tape recording mode).
- performance subsystem function (PSF)**
A command to transfer data from the control program to the subsystem to control certain conditions or functions within a control unit.
- PL** PUT macro, locate mode (value of MACRF).
- PM** PUT macro, move mode (value of MACRF).
- PO** Partitioned organization (value of DSORG).
- POSIX**
Portable operating system interface for computer environments.
- POU** Partitioned organization unmovable (value of DSORG).
- PRTSP**
Printer line spacing (parameter of DCB).
- PS** Physical sequential (value of DSORG).
- PSF** Printing Subsystem Facility or Perform subsystem function
- PSU** Physical sequential unmovable (value of DSORG).
- PSW** Program status word.
- QISAM**
Queued indexed sequential access method.
- QSAM**
Queued sequential access method.
- R0** Record zero.
- RACF** Resource Access Control Facility.
- RACF authorization**
The facility for checking a user's level of access to a resource against the user's desired level of access.

The result of that check.
- RBA** Relative byte address.
- RDBACK**
Read backward (parameter of OPEN).
- RDF** See *record definition field*.
- RDW** Record descriptor word.
- RECFM**
Record format (JCL keyword and DCB macro parameter).
- record definition field (RDF)**
A field stored as part of a stored record segment; it contains the control information required to manage stored record segments within a control interval.
- REFDD**
Refer to previous DD statement (JCL keyword).
- relative record data set (RRDS)**
A VSAM data set whose records are loaded into fixed-length slots.
- RETPD**
Retention period (JCL keyword).
- RKP** Relative key position (parameter of DCB).
- RLS** Record-level sharing
- RLSE** Release unused space (DD statement).
- RMODE**
Residence mode.
- RPL** Request parameter list.
- RPS** Rotational position sensing.
- RRDS** Variable-length or fixed-length relative record data set. Also see *relative record data set*.
- RRN** Relative record number.
- SAA** Systems Application Architecture®.
- SBCS** Single-byte character set.
- SDW** Segment descriptor word.
- SEOF** Software end-of-file.
- SER** Volume serial number (value of VOLUME).
- SETL** Set lower limit of sequential retrieval (QISAM macro).
- SF** Sequential forward (parameter of READ or WRITE).
- shared resources**
A set of functions that permit the sharing of a pool of I/O-related control blocks, channel programs, and buffers among several VSAM data sets open at the same time.
- SI** Shift in.

simple name

The rightmost component of a qualified name. For example, APPLE is the simple name in TREE.FRUIT.APPLE.

- SK** Skip to a printer channel (parameter of CNTRL).
- SL** IBM standard labels (value of LABEL).
- SMB** System-Managed Buffering
- SMS** Storage Management Subsystem, and system-managed storage.
- SMSI** Size of main-storage area for highest-level index (parameter of DCB).
- SMSW**
Size of main-storage work area (parameter of DCB).
- SO** Shift out.
- source catalog**
An existing catalog that may be exported into a target catalog.
- SP** Space lines on a printer (parameter of CNTRL).
- SRB** Service request block.
- SS** Select stacker on card reader (parameter of CNTRL).
- SSCR** Subsystem checkpoint record.
- SSI** Subsystem interface.
- STORCLAS**
Storage class (JCL keyword).
- SUL** IBM standard and user labels (value of LABEL).
- SWA** Scheduler work area.
- SYNAD**
Synchronous error routine address (parameter of DCB, DCBE, and EXLST).
- T** Track overflow option (value of RECFM); user-totaling(value of OPTCD).
- task control block (TCB)**
Holds control information related to a task.
- TCB** See *task control block*, and trusted computing base.
- TCDB** Tape configuration data base.
- TIOT** Task input/output table.
- TMP** Terminal monitor program.

TRC Table reference character.

TRTCH

Track recording technique (parameter of DCB and of DD statement).

TTR Track record address.

UCB See *unit control block*.

UCS See *universal character set*.

UHL User header label.

unit control block (UCB)

A control block in storage that describes the characteristics of a particular I/O device on the operating system.

universal character set (UCS)

A printer feature that permits the use of a variety of character arrays.

UPD Update.

USAR User security authorization record.

USVR User security verification routine.

UTL User trailer label.

VIO Virtual input output.

VIO storage group

A type of storage group that allocates data sets to paging storage, which simulates the activity of a DASD volume. VIO storage groups do not contain any actual DASD volumes.

VSAM

Virtual storage access method.

VSAM RLS

VSAM record-level sharing.

VSI VSAM shared information.

VTOC Volume table of contents.

VVDS VSAM volume data set.

XLRI Extended logical record interface.

Index

Numerics

13F completion code 56
3400 Tape Drive 56
3800 Printing Subsystem 19
3890 Document Processor 19
4092 completion code 59
7-track tape 8, 52

A

ABEND
 13F 56, 62
 23F 13, 62
 33F 62
 43F 62
 53F 62
 D37 8
abend codes 26
ABEND macro 26, 27, 28
abnormal termination
 ABEND macro 27
 automatic checkpoint/restart 27
 BSAM (basic sequential access method) 47
 CANCEL operand 15, 20
 checkpoint entry 13, 58
 CKPTREST macro 2
 COND=EVEN, COND=ONLY 30
 data set (job step) 52
 data set (passed) 33, 35, 37
 data set (SYSOUT) 48
 data set (update) 55
 DD statement 9
 HOLD queue 26
 job (recording information) 1
 job step 25
 restart step 33
 SYSABEND, SYSUDUMP 48
ACB (access method control block) 3, 49
 VSAM (virtual storage access method) 44
accessibility 91
 contact IBM 91
 features 91
ACS (automatic class selection) 30
ACS routine 9
address (checkid) 15
ALESERV macro 24
allocation
 avoiding shared DASD 8
 dynamic
 checkpoint 31
 checkpoint/restart 22
 concatenated data sets 22
 control information change 19
 generation data set 44
 nocapture option 23, 74
 space 36
 storage requirement 11
 VIO data set 22

allocation (*continued*)
 space 3, 8, 27, 33, 70
 restart 49
 secondary 8, 9
AMP parameter 50
APF (authorized program facility) 4
assistive technologies 91
ATL (automated tape library)
 Tape Library Dataserver 52
ATTACH macro 18, 63
automatic checkpoint/restart
 canceling 20
 data set disposition 9, 30
 description 25
 DFSORT 59
 ensuring restart 13
 identification (primary, secondary) 15
 JES3 58
 job journaling 3
 job reinterpretation 31
 requesting (conditions) 27
 requesting (example) 28
 system description 30
 temporary data set 49
 type 25
automatic restart
 ADDRSPC=REAL 27
 JCL (requirements, restrictions) 30
 operator options 26
 RD parameter 38
 resource variations 30
 step termination 36
 system 30
 VIO data set 22
 virtual storage 38
 volume 30
automatic step restart
 abnormal termination 26
 data set disposition 29
 eligible codes 2
 establishing checkpoints 6
 EXTRACT macro 18
 initiating, AVR (automatic volume recognition) 29
 job journaling 3
 MOD data set 27, 29
 PL/I 59
 RD parameter 2, 28
 restrictions 33
 SETPRT macro 19
 type 25
AVR (automatic volume recognition) 29

B

BDAM (basic direct access method)
 data sets 43
 serially reusable resources 20
BLDVRP processing
 VSAM restart 77

block count
 DCB (negative) 66
 nonstandard tape labels 52
 reel tape
 EXCP 54
BLP (bypass label processing) 52
BPAM (basic partitioned access method)
 checkpoint considerations 5
 DCB parameter 8
 SYSCKEOV DD statement 7
 UNIT parameter 8
BSAM (basic sequential access method)
 DCB parameter 8
 DD statement 7
 EOV (end-of-volume) 1, 89
 repositioning data set 47
 UNIT parameter 8
 using 5
 VIO data set 22
bypass label processing (BLP) 52, 66

C

CALL PLICKPT 59
CANCEL command 30
CANCEL operand 15, 20
capacity
 checkpoint data set 10
card readers 47
cataloged procedure
 data sets 35
 EXEC statement 39
 step 32
cataloged procedure step 38
CBIC option 67
CBUF (control block update facility)
 option 76
chained scheduling
 DCB option 8
channel programs
 DCB option 8
CHECK macro 21, 44
checkid (checkpoint entry identification)
 message at restart 26
 not specified 40
 PDS (partitioned data set) 11
checkpoint canceling 20
checkpoint data set
 DCB (data control block) 8
 DD statement 8
 disposition 13
 extending 12
 opening, ensuring restart 13
 sequential 14
 space allocation 8
 storage estimates 10
 SYSCHK DD statement 2
 VIO data set 22
checkpoint entry 29
checkpoint entry, writing 11
checkpoint EOV 6

CHECKPOINT parameter 62
 checkpoint routine 5, 48
 checkpoint, taking 46
 checkpoint/restart
 data spaces 24
 dependencies 2
 description 1
 ensuring restart 13
 EOV (end-of-volume) 6
 job accounting 57
 messages 12
 checkpoint/restart, automatic
 accounting 57
 checkpoint/restart, deferred
 job accounting 57
 CHKPT macro 1, 5, 6, 7, 11, 15, 18, 20,
 30, 32, 34, 38, 44, 49, 50, 56, 61, 63, 64,
 67, 68, 69, 74, 76, 89
 CHKPT=EOV parameter 6
 CKPT parameter 59
 CKPTREST macro 2
 CLASS initialization statement 3
 CLOSE macro 19, 49
 CLPA (change link pack area) 37
 nucleus table 72
 COBOL RERUN clause 59
 command
 CANCEL 30
 DISPLAY 27, 38
 RELEASE 26, 30
 completion code (13F) 56
 completion code (4092) 59
 completion codes
 returned in register 61, 86
 concatenated data set 6
 COND parameter 34
 conditional disposition 33, 35
 COPY operation 47
 CROPS parameter 50
 cross memory support
 restrictions 23
 cross-memory
 addressing 68
 mode 69
 CSV (Contents Supervisor) 74

D

DASD (direct access storage device) 55
 nonshared, shared (space allocation) 8
 data conversion (with odd parity) 8
 data set
 checkpoint security 4
 DD statement (checkpoint) 30
 dummy 36
 generation 33
 open (RDBACK) 66
 passed 30
 password protection 79
 positioning 54
 preallocated, temporary 49
 repositioning, direct access 29
 tape (reposition) 33
 TCAM (telecommunications access
 method) 48
 temporary 30
 VIO (virtual I/O) 27, 58

data set (*continued*)
 VSAM (virtual storage access
 method) 49, 50
 data set (entry-sequenced) 51
 data set (key-sequenced) 51
 data set (password protection) 3
 data set (relative record) 51
 data set buffers 54
 data set disposition 9, 29, 30, 32, 34
 data set move 47
 data set types (VSAM) 51
 data set update
 abnormal termination 55
 data spaces 24
 DCB (data control block) 4, 8, 11, 19
 exit list 89
 invalid, ABEND 8
 opened, passed to checkpoint 12
 restrictions 8
 DCB macro 54
 DCB parameter 65
 DCBE (data control block extension) 8,
 75
 DD statement 5, 9, 13, 34
 checkpoint data set 28
 DATA delimiter 37
 requesting automatic
 checkpoint/restart 28
 DDNADDR parameter 17
 DDNAME
 passed to checkpoint 12
 DEB (data extend block) 70, 83
 DEFER parameter 9
 deferred
 DD statement 32
 step restart 33
 step restart (requesting example) 32
 step restart (requirements,
 restrictions) 33
 step restart (resource variations) 34
 step restart (SYSCHK DD
 statement) 40
 step restart (type) 25
 deferred checkpoint/restart
 ADDRSPC=REAL 37, 38
 checkpoint entries 12, 14
 description 25
 dynamic concatenation 22
 JCL (requirements, restrictions) 35
 job step time limit 58
 messages 31
 nucleus 38
 operator considerations 32
 PL/I 59
 requesting 2, 34
 requesting (example) 35
 resource variations allowed 37
 submitting jobs 36
 SYSCHK DD statement 2, 34, 40
 SYSIN data set 48
 system operation 38
 system restrictions 37
 type 25
 VIO data set 22
 volume serial number 37
 DEFrag operation 36, 47
 delimiter (/*) statement 37

demount (tape volume DEQ) 21
 DEQ macro 21, 22
 device
 allocation (failure) 56
 step restart, deferred 34
 device allocation
 direct access 48
 restart 27, 30
 restart, virtual storage 29
 UNIT parameter 55
 device change
 restart 30
 DFSMSdss 36, 46, 47
 DFSMSHsm 47
 DFSORT program 59
 direct access
 device processing 5
 EOV (end-of-volume) 9, 10
 volume 15
 VSAM data set 51
 directory (preserving contents) 46
 DISP operand 9
 DISP parameter 59
 DISPLAY command 27, 38
 disposition, conditional 33, 35
 DOM macro 64
 DSAB (data set association block)
 DD statement (restart) 73
 DSCB (data set control block) 36, 55, 70
 duplicate 29
 DSNNAME parameter 40
 dummy data set 22, 36
 DUMMY parameter 37
 duplicate DSCB (data set control
 block) 29
 dynamic
 allocation
 checkpoint 31
 checkpoint/restart 22
 data sets 22
 generation data set 44
 nocapture option 23, 74
 space 36
 VIO data set 22
 concatenated data sets 22
 control information change
 SETPRT macro 19
 link pack area 62, 69
 storage requirement
 checkpoint/restart 11
 unallocation
 checkpoint 11
 conditional disposition 36
 data sets 33
 JCL specified data sets 36

E

ENQ macro 20, 21, 62
 EOV (end-of-volume) 1, 2, 6, 9, 12, 36,
 39
 checkpoint routine, macro 89
 exit routine (CHKPT) 7
 SYSCKEOV DD statement 7
 EROPT (error options) 56, 65
 error options (EROPT) 56
 errors (input/output) 56

ESDS (entry-sequenced data set) 51
 ESETL macro 21
 ESPIE macro 67, 86
 ESTAE
 recovery environment 79
 service request 85
 system request 61
 EXCP access method 54
 EXCP macro 5
 EXEC statement 28, 30, 33, 34, 36, 40
 automatic checkpoint/restart 27
 operator considerations 27
 TIME parameter 58
 extended format 43
 extended format data set 9, 44
 extents, data set 36
 external checkpoint function 6
 EXTRACT macro 18

F

FAILURE option 58
 FCB (forms control buffer) 6, 19
 FEOV macro 18
 file protection 32, 53
 FORCECP(days)
 SMS-managed sequential data set 47
 fragmented storage 11

G

GAM (graphics access method) 45
 GDG (generation data group) 45
 GDGNT (generation data group name table) 5, 44
 GDS (generation data set) 33, 45
 cataloging 44
 JES3 restriction 58
 GETMAIN failure 71
 graphic data control block 63
 GSR (global shared resources) 23

H

HFS (hierarchical file system)
 extended format data set 9
 files 74
 hiperbatch 50
 QSAM (queued sequential access method) 86
 VSAM (virtual storage access method) 49
 hiperspace buffer 23
 HOLD reply 31
 home address space ID 68

I

IAT2006 58
 IAT2575 58
 IDA3009I 76
 IDADDR parameter 17
 identification
 primary, secondary 14
 IDLNG parameter 17

IEBCOPY restart restriction 46
 IEC068A 53
 IEC161I 75, 76, 77
 IEC252I 75, 76
 IEC252I 77
 IEC331I 76
 IEC400A 53
 IEC401A 53
 IEC402D 53
 IEC954I 13
 IEF007I 22
 IEF225D 3, 25, 58
 IHJ000I 9, 12
 IHJ000I 50
 IHJ001I 86
 IHJ006I 38
 IHJ007I 86
 IHJ008I 38
 IHJ009I 76
 IHJ014I 73
 IHJ101I 80
 IHJ102I 80
 IMGLIB macro 18, 80
 increment subparameter 9
 INOUT option 45
 input work queues 38
 interface to VSAM (ISAM data set) 36
 IPL (initial program load) 37
 IPL (initial program load) 26, 32, 37
 ISAM (indexed sequential access method) 36
 ISO/ANSI (International Organization for Standardization/American National Standards Institute) 9
 ISO/ANSI standard labels 52

J

JCL (job control language)
 requirements, restrictions 30, 35
 JES2
 job journaling 3, 39
 JES3
 GDS (generation data set) 49
 GDS restriction 58
 job journaling 3, 39, 58
 output spool space 48
 JFCB (job file control block) 5, 21, 84
 JFCBE (job file control block extension) 5, 19
 job accounting 57
 job canceling 27
 job failure 26
 after automatic restart 28
 automatic restart 38
 causes 29
 during checkpoint entry 13
 job journaling 58
 saving checkpoint data (MOD) 7
 job journaling 3, 25, 26
 system restart 58
 job reinterpretation 31
 JOB statement 28, 32
 automatic checkpoint/restart 28
 job step
 abnormal termination 25, 26
 JOBLIB DD statement 40

K

keyboard
 navigation 91
 PF keys 91
 shortcut keys 91
 KSDS (key-sequenced data set) 51

L

label
 bypass label processing (BLP) 66
 ISO/ANSI (AL or AUL) 9
 ISO/ANSI tape (data set) 52
 LTM (leading tapemark option) 52
 no label (NL) 9, 66
 nonstandard (NSL) 9
 nonstandard label (NSL) 72
 nonstandard tape 52
 serial numbers 73
 standard 15
 standard (checkpoint data set) 8
 standard (SL, SUL) 9
 LABEL parameter 9, 64
 large format data set 44
 library
 automated tape 52
 link library 30
 linkage stack 86
 LINKLIB (data set) 30
 LPA (link pack area) 37
 list, DCB exit 89
 nucleus routines/tables 72
 LSR (local shared resources) 23
 LTM (leading tapemark option)
 LABEL=(,LTM) 52

M

macro
 ABEND 26, 27, 28
 ALESERV 24
 ATTACH 18, 63
 CHECK 21, 44
 CHKPT 1, 5, 11, 15, 18, 20, 30, 32, 34,
 38, 44, 49, 50, 56, 61, 63, 64, 67, 68,
 69, 74, 76, 89
 CKPTREST 2
 CLOSE 19, 49
 DCB 54
 DEQ 21, 22
 DOM 64
 ENQ 20, 21, 62
 ESETL 21
 ESPIE 67, 86
 EXCP 5
 EXTRACT 18
 FEOV 18
 IMGLIB 18, 80
 OPEN 9
 PCLINK 19, 24, 64
 READ 54
 RELEX 21, 44
 RESERVE 20
 SETDEV 19
 SETL 21
 SETPRT 19

- macro (*continued*)
 - STIMER 19, 63
 - STOW 11, 12, 46, 65, 72
 - WAIT 21, 44, 64
 - WRITE 8, 21, 44
 - WTOR 19, 64
- magnetic tape
 - checkpoint entry 12
 - data set organization 59
 - devices 56
 - processing (checkpoint entry) 5, 9
 - restriction 40
 - standard label 8
 - volume 15
 - volumes 9, 22, 34, 41
- MCP (message control program) 48
- member
 - checkpoint entry 11
 - deleting 46
- message
 - abend (data set security) 26
 - authorization for restart 26
 - checkpoint 9
 - checkpoint restart 12
 - checkpoint return codes
 - associated with IHJ Messages 86
 - control program 48
 - EOV (end-of-volume) 6
 - error 61, 86
 - IAT2006 58
 - IAT2575 58
 - IDA3009I 76
 - IEC068A 53
 - IEC161I 75, 76, 77
 - IEC252I 75, 76
 - IEC252I 77
 - IEC331I 76
 - IEC400A 53
 - IEC401A 53
 - IEC402D 53
 - IEC954I 13
 - IEF007I 22
 - IEF225D 3, 25, 58
 - IHJ0001 9, 12
 - IHJ000I 50
 - IHJ001I 86
 - IHJ006I 38
 - IHJ007I 86
 - IHJ008I 38
 - IHJ009I 76
 - IHJ014I 73
 - IHJ101I 80
 - IHJ102I 80
- mount 32
- replies 26, 27
- restart sequence 25
- sequence 31
- MIGRATE operation 47
- MOD data set 30, 33
 - automatic restart 27
 - automatic step restart 29
 - checkpoint entries 11
 - example 13
 - PL/I programs 59
 - restart, repositioning 29
 - SYSCKEOV DD statement 7

- modules, return codes
 - associated with IHJ Messages 86
- MSS error
 - system 73
- multivolume data set 6, 40, 41
 - JCL parameter 1
 - nonspecific volume allocation 45
 - physical sequential 1
- MVS/SP
 - scheduler checkpoint support 85

N

- name/token
 - restriction 86
 - pair 24
- navigation
 - keyboard 91
- no-reposition option 50
- nocapture option 23, 74
- nonpageable region 27
- nonpageable virtual storage 37
- nonspecific type volume 22, 27, 31
- nonstandard labels 52
- NORESTART statement type 2
- Notices 95
- NRC subparameter 51
- NSLREPOS
 - nonstandard label (NSL) routine 72
- nucleus
 - deferred checkpoint/restart 38

O

- OPEN macro 9
- OPEN routine 54
- operator considerations 26, 32
- operator message
 - automatic checkpoint/restart 25
- OPTCD parameter 8
- OUTIN option 45

P

- PAGE pack 27
- parameter
 - AMP 50
 - CHECKPOINT 62
 - CKPT 59
 - COND 34
 - CROPS 50
 - DCB 65
 - DEFER 9
 - DISP 59
 - DSNAME 40
 - DUMMY 37
 - LABEL 9, 64
 - PGM 33, 34
 - QNAME 49, 71
 - RD 2, 3, 28, 38, 39, 61
 - RESTART 2, 3, 16, 32, 33, 34, 36, 39
 - RLSE 9
 - SPACE 9
 - SYSOUT 47, 48
 - TIME 58
 - UNIT 8, 45, 52, 55, 58, 73

- parameter (*continued*)
 - UNSTACK 64
 - VOL 22, 34, 37
- PASN-AL (primary address space access list) 87
- password
 - protected 25
 - protection 79
 - restart (message) 26, 32
- PC routine, restriction 24
- PCLINK macro 19, 24, 64
- PDS (partitioned data set) 43
 - checkid (checkpoint entry identification) 14
 - checkid length 16
 - considerations (checkpoint/restart) 45
 - directory (members) 46
 - open for update 46
 - restriction (SYSCHK DD statement) 40
 - writing checkid 11
- PDSE (partitioned data set extended) 43
 - checkpoint 47
 - extended format data set 9
- PGM parameter 33
- PGSER (Page Fix) 77
- PL/I program 59
- PLIREST 59
- PQUIT command 59
- preallocated data set
 - temporary 49
- primary
 - address space ID 68
 - identification 14
- primary address space access list (PASN-AL) 87
- printer
 - setup (universal character set) 19
- printer repositioning
 - at restart 54
- processor time 57
- program control assigned
 - checkpoint identification 15
- programmer-specified
 - checkid (checkpoint entry identification) 1, 16
- PSF (file protect status)
 - Tape Library Dataserver 73
- PSF (performance subsystem function) 53
- PSW key 86

Q

- QNAME parameter 49, 71
- QSAM (queued sequential access method)
 - card reader 54
 - checkpoint/restart 5
 - CHKPT JCL parameter 6
 - data set 55
 - EOV (end-of-volume) 89
 - I/O errors 56
 - VIO data set 22
- queue, input work 31

- quick reference
 - checkpoint/restart codes 61
 - checkpoint/restart processing 43
 - coding a checkpoint 5
 - end-of-volume checkpoint routine 89
 - introduction 1
 - requesting restart 25

R

- RACF (Resource Access Control Facility) 3
- RD parameter 2, 3, 28, 38, 39, 61
- RDBACK (read backward)
 - BLP (bypass label processing) 66
 - data set open 66
- READ macro 54
- reason codes
 - checkpoint/restart 62
- record catalog 51
- record update
 - after checkpoint 55
- recording data set positioning 54
- reduced performance 11
- REGION size parameter 32
- region, nonpageable 27
- register contents
 - EOV (end-of-volume) 89
 - return/reason codes 61
- RELEASE command 30
 - initiate restart 26
- RELEX macro 21, 44
- repositioning
 - data set (VSAM) 50
 - data sets (VSAM) 49
 - deferred step restart 33
 - password protected 25
- repositioning card reader 54
- repositioning restriction 23
- RERUN clause 59
- RESERVE macro 20
- resource variations 30, 37
 - deferred step restart 34
- restart
 - EOV (end-of-volume) 5
 - responses 26
 - unauthorized programs 4
 - VSAM data sets 50
- restart card reader 54
- restart job eligibility 2
- restart message 3
- RESTART parameter 2, 3, 16, 25, 32, 33, 34, 36, 39
- restart request 2
- RESTART statement
 - JES3 58
 - RACF protected data sets 3
 - SCHEDxx
 - PLIREST 59
 - SVC 77
 - SYS1.PARMLIB
 - SCHEDxx 2
 - SYS1.PARMLIB member
 - SCHEDxx 27, 28
- restart types 25
- restart, automatic
 - data set disposition 9

- restart, automatic (*continued*)
 - description 25
 - job step accounting 57
- restart, deferred
 - messages 31
- restart, ensuring 13
- RESTART/NORESTART 2, 27, 28
- return codes
 - add, delete 2
 - checkpoint/restart 61, 62, 86
 - CHKPT 18
- RLS (record level sharing) 49, 67
- RLSE parameter 9, 36
- routine (checkpoint) 5, 9
- RPL feedback word 49, 75
- RPS (rotational position sensing) 55
- RRDS (relative record data set) 51

S

- secondary
 - address space ID 68
 - addressing mode 69
 - allocation 65
 - identification
 - checkid (checkpoint entry identification) 15
 - space 6
 - space allocation 9
- security
 - checkpoint data set 4
 - user data set 3
- sending comments to IBM xi
- sequential data set 11
 - checkpoint/restart 13
- serially reusable resources 20
- SETDEV macro 19
- SETL macro 21
- SETLOCK system request 61, 77
- SETPRT macro 19
- shared resources 23
- shortcut keys 91
- SHOWCB macro 49
- SMS environment 40
- SMS-managed
 - data set 9, 30, 40, 41
 - sequential data sets 43, 47
- SPACE parameter 9
- space release
 - deferred checkpoint/restart 36
- SPOOL pack 27
- SSCR (subsystem checkpoint record) 66, 80, 84, 86
 - DCB (restart) 73
- step restart, automatic
 - accounting 57
 - JES3 58
 - resource variations 30
 - suppressed (JES3) 49
- step restart, deferred
 - checkpoint entries 12, 14
 - description 25
 - DFSORT 59
 - GDS (generation data set) 44
 - JCL requirements 33
 - job and job step accounting 57
 - requesting 2

- step restart, deferred (*continued*)
 - VIO data set 22
- STIMER macro 19, 63
- storage estimating 11
- STOW macro 11, 12, 46, 65, 72
- subparameter
 - NRC 51
- subpools 11
- subtasks 18
- Summary of changes xiii
- suppressing action 2
- SWA (scheduler work area)
 - control block 65, 68
 - manager 65, 69
 - pageable storage 3
 - read a JFCB 84
- SYNCDEV macro 65
- SYS1.IMAGELIB data set 19
- SYS1.PARMLIB member 2
- SYSABEND 43
- SYSABEND data set 47
- SYSCHK DD statement 2, 34, 36, 40, 41
- SYSCKEVO DD statement 2, 7
- SYSIN 43
- SYSIN data set
 - automatic checkpoint/restart 48
- SYSOUT 43
- SYSOUT data set
 - checkpoint/restart 48
- SYSOUT parameter 47, 48
- SYSRES (system residence volume) 30
- SYSRES volume 27
- system failure 26
 - automatic restart 38
 - during checkpoint entry 13
 - job journaling 58
 - saving checkpoint data (MOD) 7
- system generated
 - checkid (checkpoint entry identification) 1
- system generated name (data set) 31
- system operation 30, 34
- SYSUDUMP data set 47

T

- table terminal 48
- tape
 - 7-track 8
 - cartridge 53
 - cartridge protection 53
 - data set (ISO/ANSI label) 52
 - DCB (data control block)
 - parameters 8
 - files (VSE) 52
 - label conflict 54
 - labels (data set) 9
 - reel (data set) 54
 - volume 15, 22, 73
 - volume (DEQ) 21
 - volume, EOV (end-of-volume) 9
- Tape Library Dataserver 52, 54, 73
- normal mount 26
- operator message 31
- TCAM (telecommunications access method) 43, 48
- TCB (task control block) 18

- TCBUSER field 57
- TCDB (tape configuration data base)
 - Tape Library Dataserver 74
- termination
 - automatic restart 9
 - completed at system restart 58
 - records reprocessed 56
- time limit
 - job step 58
- TIME parameter 58
- trademarks 99
- TTR (track record address) 29

U

- UCS (universal character set) 6
 - SETPRT macro 19
- unallocation
 - dynamic
 - checkpoint 11
 - conditional disposit 36
 - data sets 33
 - JCL specified data sets 36
- unavailable virtual storage 37
- UNIT keyword 41
- UNIT parameter 8, 45, 52, 55, 58, 73
- unit record
 - data sets 48
 - devices 19, 54
- UNIX system services 74
- UNLOAD command 27
- unmovable attribute 47
- UNSTACK parameter 64
- UPDAT option 45
- update in place 55
- user
 - data set security 4
- user data sets 5
- user errors 56
- user interface
 - ISPF 91
 - TSO/E 91

V

- variable spanned record format
 - data set 47
- VARY command 27
- vector facility 86
- VIO (virtual I/O) 22, 77
 - data set allocation 22
 - data sets 3, 58
 - allocation, cancel 27
- virtual fetch 59
- virtual storage
 - buffers (data set records) 54
 - CHKPT macro 1, 15
 - contents written 1, 5
 - EOV (end-of-volume) 89
 - job journaling 3
 - operator considerations 32
 - requirements 26, 27, 32
 - requirements (TCAM) 48
 - writing, DASD 8
- VOL parameter 22, 34, 37

- volume
 - allocation 29
 - automatic restart 30
 - changing at restart 27
 - configuration 37
 - deferred step 34
 - direct access 5, 15, 59
 - mount (VSAM) 31
 - serial (restart) 33
- VSAM (virtual storage access method) 5, 43
 - completion codes 49, 50
 - data set 50
 - data set repositioning 49
 - ESDS (entry-sequenced data set) 51
 - hiperbatch 50
 - I/O errors 56
 - KSDS (key-sequenced data set) 51
 - RPL feedback word 75
 - RRDS (relative record data set) 51
 - serially reusable resources 20
 - volume mount 31
- VSE (virtual storage extended)
 - checkpoints 52
- VSM (virtual storage manager) 68
- VTAM
 - CHKPT (ACB) 78

W

- WAIT macro 21, 44, 64
- WRITE macro 8, 21, 44
- write validity check 8
- WTOR macro 19, 64



Product Number: 5650-ZOS

Printed in USA

SC23-6862-00

