

IBM z/VSE



Diagnosis Tools

Version 3 Release 1

IBM z/VSE



Diagnosis Tools

Version 3 Release 1

Note!

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

First Edition (March 2005)

This edition applies to Version 3 Release 1 of IBM z/Virtual Storage Extended (z/VSE), Program Number 5609-ZVS, and to all subsequent releases and modifications until otherwise indicated in new editions.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the addresses given below.

A form for readers' comments is provided at the back of this publication. If the form has been removed, address your comments to:

IBM Deutschland Entwicklung GmbH
Department 3248
Schoenaicher Strasse 220
D-71032 Boeblingen
Federal Republic of Germany

You may also send your comments by FAX or via the Internet:

Internet: s390id@de.ibm.com
FAX (Germany): 07031-16-3456
FAX (other countries): (+49)+7031-16-3456

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

© Copyright International Business Machines Corporation 1984, 2005. All rights reserved.

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

Contents

Figures	vii
--------------------------	------------

Notices	ix
Programming Interface Information	ix
Trademarks and Service Marks	x

About This Book	xi
Restriction	xi
Who Should Use This Book	xi
How to Use This Book	xii
Where to Find More Information	xii

Summary of Changes	xiii
-------------------------------------	-------------

Understanding Syntax Diagrams	xv
--	-----------

Part 1. Dumps of Virtual Storage. 1

Chapter 1. General Description 3

Dump Contents Overview	3
The ABEND Dump Function	3
Overview of ABEND Dump Function	4
Activation of the ABEND Dump Function	5
Contents of the ABEND Dump Output	5
The DUMP Command Dump	7
The Stand-Alone Dump (SADUMP) Program	7
Support of Integrated Console by SADUMP Program	7
Output of the Stand-Alone Dump Program	8
The SDAID Dump	9
Dump Requested by Macros	9
Info/Analysis	10

Chapter 2. Maintaining the Dump Library and File Environment 11

The Library and Files Required to Process Dumps	11
The SYSDUMP Sublibraries	11
Purpose of the SYSDUMP Library	11
Establishing the Dump Sublibraries	13

Chapter 3. Handling Dumps. 15

Options to Activate Dump Writing	15
Options to Deactivate Dump Writing	15
Identifying the Stored Dumps	15
Sending Dumps to IBM Support Electronically	16
Mailing Dumps That Are Stored on Tape to IBM Support	17
Handling a Dump Library Full Condition	18

Chapter 4. Requesting a Dump 19

Overview of Dump Requests	19
Options to Control the ABEND Dump	19
Options to Control the Dump Contents	20

Options to Control the Output Destination	20
Requesting a Dump by the CANCEL Command	21
Requesting a Dump by the DUMP Command	21
Taking a Stand-Alone Dump	21
Requesting a Dump on Event (SDAID Dump)	23
Requesting a Dump from a Program	23
Printing the Stored Dump	24
Archiving Expired or Unrequired Dumps	24

Chapter 5. The DOSVSDMP Utility 25

The DOSVSDMP Utility Functions	25
Functions of the DOSVSDMP Utility	25
Creating the Stand-Alone Dump Program	25
Scanning the Dump Files on Disk or Tape	28
Dumps Printed with DOSVSDMP	29
Printing an SDAID or DUMP Command Produced Tape	30

Part 2. Interactive Trace Program 33

Chapter 6. Interactive Trace Program 35

Introduction	35
Branch Trace	35
Instruction Trace	35
Storage Alteration Trace	35
ABEND Trace	35
Trace Activation	36
Interactive Trace Commands	36
TRACE Command	37
QUERY Command	37
DISPLAY Command	38
ALTER Command	38
GO Command	39
Tracing in a User Partition with Subtasks Attached	39
Scope of Tracing	40
Restrictions for Programs Using the PER Function	40
The Interactive Trace Program versus SDAID	40
Examples of the Interactive Trace Program	41
Trace Initialization Example	41
TRACE, TRACE END and QUERY Command Example	41
Batch Trace Example	42
DISPLAY and ALTER Command Example	43

Part 3. SDAID Trace 45

Chapter 7. SDAID Overview 47

The SDAID Session	47
Interaction SDAID versus Interactive Trace Program	47
How to Initialize an SDAID Trace	47
Initialization in Direct Input Mode	48
Initialization via Job Control Procedures	48
Initialization via Prompts in the Attention Routine	48

AR Commands to Start, Stop and End an Initialized Trace	49
Trace Type Summary	49
Trace Output	50
Performance Considerations	50
SDAID Space Requirements	51
Number of Traces per Session	52

Chapter 8. SDAID General Description 53

Defining the Output Device	53
Printer Defined as Output Destination	53
Tape Defined as Output Destination	53
Buffer Defined as Output Destination.	53
Steps to Define a Wraparound Buffer.	54
Exceptional Conditions on the Output Device	55
Summary of TRACE Types	56
BRANCH Trace	56
BRANCH Trace Output Example	57
BUFFER Trace	57
CANCEL Trace	57
CANCEL Trace Output Example	57
EXTERNAL Trace	58
SDAID Default Value	58
EXTERNAL Interrupt Trace Output Example	58
GETVIS / FREEVIS Trace.	58
SDAID Default Value	59
GETVIS / FREEVIS Trace Output Example	59
INSTRUCTION Trace	59
INSTRUCTION Trace Output Example	60
IO Trace (I/O Interrupt)	60
SDAID Default Value	60
I/O Interrupt Trace Output Example	61
LOCK / UNLOCK Trace	61
SDAID Default Value	61
LOCK / UNLOCK Trace Output Example	61
MONITORCALL Trace	62
MONITORCALL Trace Output Example.	63
PGMCheck Trace (Program Check)	63
PGMCHECK Trace Output Example	64
PGMLOAD (Fetch/Load) Trace.	64
SDAID Default Values.	65
PGMLOAD Trace Output Example	65
SSCH Instruction Trace	65
SDAID Default Value	65
STORAGE Alteration Trace	66
SDAID Default Value	66
SVC Trace (Supervisor Call)	67
VTAMBU Trace (VTAM Buffer).	68
VTAMIO Trace	68
SDAID Default Value	69
XPCC Trace	69
SDAID Default Value	69
XPCC Trace Output Example	69
Notational Conventions	69
Defining the Area to be Traced: AREA Definition.	70
Defining the Job to be Traced: JOBNAME Definition	70
Defining the Storage to be Traced: OFFset, ADDRess, PHase, LTA	71
Storage Definition for AREA and JOBNAME	71
Defining Additional Trace Output: OUTPut Definition	73

Writing the Trace Buffer	74
Recording the CCB or IORB	74
Recording the CCW	74
Recording the Partition Communication Region	75
Recording the Control Registers	75
Dumping Virtual Storage	75
Recording Floating-Point Registers	77
Recording General-Purpose and Access Registers	78
Recording PUB, LUB, ERBLOC, CHANQ	78
Dumping Processor Storage from X'00'to X'2FF'	79
Recording the Locktable Area	79
Recording the Logical Transient Area	79
Recording the Physical Transient Area	79
Recording Partition-Related Control Blocks.	79
Recording the Supervisor Area	80
Recording the System Communication Region.	80
Recording the Time-of-Day Clock	80
Recording Task-Related Control Blocks	80
Recording the XPCC Communication Control Block	80
Recording the XPCC Data Buffer	80
Defining the Trace Options: OPTiOn Definition	81
Defining the Traced I/O Devices	82

**Chapter 9. Initialize an SDAID Trace in
Direct Input Mode 83**

Initializing an SDAID Trace in Direct Input Mode	83
Selecting the SDAID Input Mode	84
Starting the SDAID Trace Initialization	85
Ending the SDAID Trace Initialization	86
Defining the Output Device in Direct Input Mode	86
Defining the Output Device	86
The TRACE Statement.	88
Summary of Trace Types	88
BRanch Trace.	89
Initialization Example	89
BUffer Trace	90
Initialization Example	90
CANcel Trace	90
Statement Example	91
Initialization Example	91
EXtErnal Trace	91
Statement Example	92
Initialization Example	92
GETVIS Trace.	93
Statement Examples	94
Initialization Example	94
INStRuction Trace	95
Statement Examples	95
Initialization Example	96
I/O Interrupt Trace.	96
Statement Examples	96
Initialization Example	97
LOCK Trace	97
Statement Examples	99
Initialization Example	99
MONitor Call Trace.	99
Statement Examples	100
PGMCheck Trace	100
Statement Examples	101
Initialization Example	101

Program Load Trace (Fetch/Load Trace)	101
Statement Examples	102
Initialization Example	102
SSCH Instruction Trace	103
Statement Examples	103
Storage Alteration Trace	103
Statement Example	105
Initialization Example	105
Supervisor Call Trace	105
Statement Examples	106
Initialization Examples	106
VTAM Buffer Trace	107
Statement Example	107
VTAMIO Trace	107
Statement Example	107
Initialization Example	107
XPCC Trace	108
Statement Examples	110
Initialization Example	110
Additional Definitions	110
ARea or JOBNAME Definition	111
ADDRESS Definition	111
OFFset Definition	112
PHase Definition	112
I/O Device Definition	113
OUTPut Definition	113
OPTion Definition	116

Chapter 10. Initialize an SDAID Trace via a Procedure 117

Introduction	117
Notational Conventions	117
Default Value Considerations	118
Writing Cataloged Procedures	119
The Statements of a Cataloged Procedure	119
Procedures to Initialize SDAID Traces	121
Summary of Trace Procedures	121
Branch Trace Initialization	121
Instruction Trace	122
SSCH and I/O Interrupt Trace	123
Fetch/Load Trace	124
Program Check Trace	125
Storage Alteration Trace	126
SVC Trace	127
Additional Keyword Operands in Trace Procedure Statements	128
Define the Output Device in a Procedure Statement	129
TERM=Keyword Operand	129

Chapter 11. Initialize a Trace in Prompt Input Mode 131

Overview	131
How to Initialize an SDAID Trace in Prompt Mode	131
The Various SDAID Commands	132
Sample SDAID Trace Initialization	132
Command Input Paths	135
OUTDEV Command Input Path	135
TRACE Command Input Path	136

Output Device Definition in Prompt Mode:	
OUTDEV Command	142
Possible Buffer Sizes	143
Specifying the Trace: TRACE Command	143
Defining the Trace Type	143
Summary of Trace Types	144
BRanch Trace	145
BUffer Trace	145
CAnceL Trace	145
EXtErnal (External Interrupt) Trace	145
GETVis (Getvis / Freevis Request) Trace	146
INStRuction (Instruction Execution) Trace	147
IO (I/O Interrupt) Trace	148
LOCK (Lock / Unlock of Resources) Trace	148
MONitorcall Trace	150
PGMCheck (Program Check) Trace	150
PGMLoad (Program Load) Trace	151
Start Subchannel Instruction Trace	152
STorage Alteration Trace	153
SVC (Supervisor Call) Trace	153
VTAMBU (VTAM Buffer) Trace	154
VTAMIO (VTAM I/O) Trace	154
XPCC (Partition Communication) Trace	154
AREA Definition	156
JOBNAME Definition	157
Prompts after AREA and JOBNAME Definitions	157
I/O Definition	158
Additional Output Definition	160
Option Definition	160

Chapter 12. Start/Stop and End the Trace 163

The Required Commands	163
STARTSD/STOPSD Commands: Starting and Stopping	163
ENDSD Command: Ending Execution	163
Attention Routine Command Example	164
How to Control the Trace under Exceptional Conditions	164
Tracing an Unintended Loop	164
Terminating SDAID Program Without the Attention Routine	165
Starting/Terminating Tracing in a System Wait Condition	165

Part 4. Info/Analysis 167

Chapter 13. Info/Analysis: Introduction 169

Operating Environment	169
The Dump Management File	169
The External Routines File	170
Label Information for Info/Analysis	172
Functional Overview	172

Chapter 14. Dump Symptoms 173

Types of Dump Symptoms	173
Environment	174
Required Symptoms	174
Optional Symptoms	174

Chapter 15. Invoking Info/Analysis	175
Submitting a Job to Invoke Info/Analysis	175
Standard Info/Analysis Job Stream	176
Control Statement Syntax	176
Entering Control Statements	176
Common Control Statements	177
SELECT - Specify a Function or End	
Info/Analysis	177
RETURN - End Current Function.	178
DUMP NAME - Specify or Add Current Dump	178
Recommendations (Restrictions) for the	
Generation of Dump Names	179
Dump Management	179
UTILITY - Initialize Dump Management File	180
DELETE - Delete Current Dump	180
PRINT - Print List of Managed Dumps	181
Printing Dump Information	183
Dump Symptoms	184
PRINT - Print Dump Symptoms	184
Dump Viewing	188
PRINT - Print Dump Data	188
CALL - Initiate Analysis Routine	190
The Stand-Alone Dump Analysis Routine	
IJBXCSMG	191
Activating the routine	191
The Stand-Alone Dump Analysis Routine	
IJBXDEBUG	192
Activating the Routine	192
Output of the Routine	192
The Stand-Alone Dump Analysis Routine IJBXSDA	198
Activating the Routine	198
Dump Offload	198
VOLID - Specify Output Volume	199
BYPASS - Skip Offload	199
ERASE - Delete or Retain Library Copy of	
Dump	200
Offloading a Dump to Tape	200
SELECT DUMP OFFLOAD versus SELECT	
DUMP MANAGEMENT DELETE	201
Dump Onload	202
VOLID - Specify Input Volume	202
FILE - Specify Dump File on Multiple-Dump	
Device.	203
Loading a Dump into a Dump Sublibrary	203
Printing a Dump Stored on Tape or Disk	205
Processing and Printing a Dump with	
Info/Analysis	205
Printing a Stand-Alone Dump with	
Info/Analysis	206

DUMP Command Dump Printed with	
Info/Analysis	210
Ending the Info/Analysis Job	211
Control Statement Sequence Examples	211
Control Statement Summary	213

Part 5. Appendixes 215

Appendix A. Symptom Records

Overview	217
Symptom Records Structure	217
Symptom Record Creation	218
Section 6	218
Locators	219
Linkage Descriptors	220
Formatting Descriptors	221

Appendix B. Other Diagnosis Tools 223

ACTION: Print Linkage Editor Map	223
Linkage Editor Map Warning Messages	223
DITTO: Dump a Disk, Diskette, or Tape	225
DSPLY/ALTER: Display or Alter Storage	226
LIBLIST: Display Library Chains	227
LIST: Print Language Translator Source Code	229
LISTIO: List I/O Device Assignments	229
LISTLOG: Display Console Communication	229
LOG: Print Job Control Statements	230
LSERV: Display Label Information Area	230
LVTOC: Display Volume Table of Contents	232
STOP/PAUSE: Suspend Program Execution	233

Appendix C. Hardware Service Aids 235

Controlling the Recovery Management Support	235
The ROD Command	235
Retrieval and Analysis of RMS Information	235
The EREP Program	235
Hardware Aids via the Operator Console	235
Hardware Alter/Display	236
Instruction Stepping Feature	236
Stop-on-Address-Compare Feature	236

Glossary 237

Index 251

Figures

1. Overview: Dump Contents	3	48. Example: Help and Cancel Initialization	134
2. Overview: The ABEND Dump Function	4	49. Sample Command Input Path	135
3. VSE Control Blocks in System Dump	6	50. OUTDEV Command: Syntax Diagram	135
4. The SYSDUMP Library Concept.	12	51. TRACE Command: Syntax Diagram	136
5. Example: Labels for the SYSDUMP Library	13	52. Prompting for a PGMLOAD Request	152
6. Example: Defining SYSDUMP with the LIBR Librarian Program	13	53. Attention Routine Commands to Start, Stop and End the Trace	164
7. Example: LIBDEF Statement for a Dump Sublibrary	14	54. Sample Job: Dump Management File Initialization	170
8. Dump Requesting Functions	19	55. Sample Job: Loading the External Routines File via DITTO	171
9. Sample: Stand-Alone Dump Program Generation	26	56. Sample Job: Loading the External Routines File via OBJMAINT	171
10. Sample: Directory of Dump Disk/Tape	29	57. Example: File Labels for Dump Processing	172
11. Sample: Dump Tape Printed with DOSVSDMP	30	58. Dump Symptoms Part	173
12. Sample Job: Printing SDAID Tape with DOSVSDMP	31	59. Sample Job: Invoke Info/Analysis.	176
13. Trace Example: Trace Initialization	41	60. Sample Job: Delete Dumps	181
14. Trace Example: TRACE, TRACE END and QUERY Command	41	61. Sample Job: List Managed Dumps	182
15. Trace Example: Batch Trace	42	62. Example: Dump Management PRINT DATA Output.	182
16. Trace Example: DISPLAY and ALTER Command	43	63. Overview: Dump Contents	184
17. Overview, Tracing Events into a Buffer	54	64. Sample Job: Print Dump Symptoms	185
18. BRANCH Trace Event Record	57	65. Example: Output of Print Dump Symptoms	185
19. CANCEL Trace Event Record	57	66. Sample Job: Print Selected Dump Areas	190
20. EXTERNAL Interrupt Trace Event Record	58	67. Sample Job: Print a Dump in Formatted Form	190
21. GETVIS / FREEVIS Trace Record	59	68. Sample Job: Call the Analysis Routines IJBXCSMG, IJBXDEBUG and IJBXSDA	191
22. Additional Fields Displayed By GETVIS / FREEVIS Trace	59	69. Example: WAITFFF Analysis Report	195
23. INSTRUCTION Execution Event Record	60	70. Example: Soft Wait Analysis Report	196
24. I/O-Interrupt Trace Event Record	61	71. Example: System Loop Analysis Report	197
25. LOCK Trace Record	62	72. Offloading a Dump to Tape.	201
26. Contents of LOCKTABLE	62	73. Summary: SELECT DUMP OFFLOAD and DELETE Operation	202
27. MONITORCALL Trace Event Record	63	74. Sample Job: Onload a Dump from Tape into a Dump Sublibrary	204
28. Program Check Trace Event Record	64	75. Overview: Print a Dump.	205
29. PGMLOAD Trace Event Records	65	76. Sample Job: Print a Stand-Alone Dump (Main Dump File)	206
30. SSCH (Start Subchannel) Trace Event Record	66	77. Sample Job: Print a Stand-Alone Dump (Additional Dump File)	207
31. Storage-Alter Trace Event Record	67	78. Sample: Symptom Part of the Stand-Alone Dump Output (Main Dump File)	208
32. SVC Trace Event Record	67	79. Summary of the DUMP VIEWING, PRINT FORMAT Operation Output	209
33. VTAMIO/VTAMBU Trace Record	68	80. Sample Job: Print the Output of a DUMP Command	210
34. XPCC Trace Record	69	81. Sample: Symptom Part of the DUMP Command Dump	211
35. Additional Fields Displayed By XPCC Trace	69	82. Summary: Areas to be Printed with DUMP VIEWING, PRINT FORMAT	211
36. Output of OUTPut=(CCWD=256)	75	83. Control Statement Sequence Example	212
37. Overview: Defining the Area to be Dumped	77	84. Control Statement Sequence Example	213
38. Printout of Floating-Point Registers	78	85. Symptom Records	217
39. Printout of General-Purpose Registers. . . .	78	86. Sample: Linkage Editor Output (ACTION MAP)	224
40. Printout of General-Purpose and Access Registers	78	87. Sample of the DSPLY and ALTER Commands	227
41. Printout of Low Address Storage	79		
42. Program-Check Event with Time of Day	80		
43. Trace Initialization Examples (Direct Input Mode)	85		
44. Example: Initializing an SDAID Trace. . . .	88		
45. SDINST Sample Procedure	118		
46. Example: Cataloged Procedure.	120		
47. Example: Prompt Mode Trace Initialization	133		

88. Example: Library Chain Listing	228	90. Control Statements to Invoke LVTOC	233
89. Sample: LSERV Output	232		

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. 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 of the intellectual property rights of IBM may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785, U.S.A.

Any pointers in this publication to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement. IBM accepts no responsibility for the content or use of non-IBM Web sites specifically mentioned in this publication or accessed through an IBM Web site that is mentioned in this publication.

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

IBM Deutschland GmbH
Department 0790
Pascalstr. 100
70569 Stuttgart
Germany

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

Programming Interface Information

This manual is intended to help the customer analyze and solve problems that may occur when IBM z/VSE is installed.

This manual also documents General-use Programming Interface and Associated Guidance Information provided by z/VSE.

General-use programming interfaces allow the customer to write programs that obtain the services of z/VSE.

General-use Programming Interface and Associated Guidance Information is identified where it occurs by the following marking:

General-use programming interface

General-use Programming Interface and Associated Guidance Information....

End of General-use programming interface

Diagnosis, Modification or Tuning Information is provided to help the customer to do diagnosis or tailoring of z/VSE.

Warning: Do not use this Diagnosis, Modification or Tuning Information as a programming interface.

Diagnosis, Modification or Tuning Information is identified where it occurs by the following marking:

Diagnosis, Modification or Tuning Information

Diagnosis, Modification or Tuning Information...

End of Diagnosis, Modification or Tuning Information

Trademarks and Service Marks

The following terms used in this publication, are trademarks or service marks of the IBM Corporation in the United States and/or other countries.

ACF/VTAM	RS/6000
CICS/VSE	SAA
DB2	System/390
ESA/370	S/390
ESA/390	TotalStorage
ECKD	VM/ESA
IBM	VSE/ESA
OS/390	VTAM
QMF	zSeries
Multiprise	z/VM
OS/2	z/VSE
Redbooks	

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

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

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

About This Book

z/VSE is the successor to IBM's VSE/ESA product. Many products and functions supported on z/VSE may continue to use VSE/ESA in their names.

z/VSE can execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities.

z/VSE is designed to exploit select features of IBM eServer zSeries hardware.

This manual is intended for customers who need to use the diagnosis tools of IBM z/VSE. These tools consist of the various dump types, the Interactive Trace Program, SDAID traces, and the Info/Analysis dump management facility of z/VSE. When to use these tools, and under what circumstances, is explained in *z/VSE Guide for Solving Problems*.

With interactive tools, you will find examples of the panels and the interactive dialogs that are used to invoke and run the programs. For tools running in batch mode, you will find examples of job control streams. If an explanation of the output is necessary, sample listings are included.

Readers of this publication should be familiar with the operational concept of the IBM VSE system.

IBM z/VSE includes diagnosis tools that help you in information gathering and problem diagnosis when a system or program malfunction occurs.

This publication describes the use of these tools.

Restriction

If any of these diagnosis tools writes the output to an IBM 3211 printer and this printer's indexing feature is being used, a number of characters may get lost on each line of the output. The system's dump and trace routines, for example, write output records of 120 bytes in length.

To avoid the loss of data, you should load another FCB (forms control buffer) image which disables the indexing feature before requesting the desired printout. For information on FCB loading, see *z/VSE System Control Statements*.

Who Should Use This Book

This manual addresses primarily the **system administrator**.

Note, however, that any of the following persons may be the first to encounter a problem:

- The system console operator.
- A display station user, including the system administrator.
- An application programmer.
- An application end user.

Most problems, however, will end up with the administrator. Whenever an application program seems to be at fault, the administrator may hand the problem over to the programmer responsible.

How to Use This Book

The conventions for showing the format of job control commands and statements used in the publication *z/VSE System Control Statements* apply also to this manual.

This publication is divided into the following parts:

- **Dumps of Virtual Storage**

Which describes the various dump functions in general and shows the file and library environment which is needed to store dumps. The methods to request and to print storage dumps which have been stored on tape or in a dump sublibrary are described. This part contains also the description of the DOSVSDMP utility.

- **Interactive Trace Program**

The Interactive Trace Program is the tracing tool for z/VSE application programs. This part describes how you can trace the execution of programs running in static or dynamic user partitions.

- **SDAID Trace**

Contains an overview of the SDAID trace program, describes all trace types and the various methods to initialize them. How you can start, stop, or terminate the initialized traces is also described in this part.

- **Info/Analysis**

Info/Analysis is the dump viewing and management facility of VSE. This part describes the use of Info/Analysis. It also describes the stand-alone dump analysis routines IJBXCSMG, IJBXDEBUG and IJBXSDA.

- **Appendixes**

Contain a description of the symptom record, various display and list aids such as the LVTOC or the LSERV program, and tells how to use some hardware diagnosis aids.

Where to Find More Information

You will need the following IBM publications when diagnosing a problem:

- *z/VSE Guide for Solving Problems*
- *z/VSE Messages and Codes*
- *z/VSE Operation*
- *z/VSE System Utilities*
- *VTAM Diagnosis*

Summary of Changes

- **z/VSE** is the successor to **VSE/ESA**. However, the names of many features and programs related to z/VSE remain unchanged (such as IBM Language Environment for VSE/ESA, IBM COBOL for VSE/ESA, or TCP/IP for VSE/ESA).
- Details have been included of how you can send dumps electronically to IBM Support. For details, see “Sending Dumps to IBM Support Electronically” on page 16.
- Details have been included of how you can mail dumps, that are stored on tape, to IBM Support. For details, see “Mailing Dumps That Are Stored on Tape to IBM Support” on page 17.
- An introduction has been included of how you can archive dumps. For details, see “Archiving Expired or Unrequired Dumps” on page 24.
- The SDAID GETVIS trace have been integrated into this manual. For details, see pages 58, 93, and 146.
- The SDAID LOCK trace have been integrated into this manual. For details, see pages 61, 97, and 148.
- The SDAID XPCC trace have been integrated into this manual. For details, see pages 69, 108, and 154.
- The ADDRESS definition, used when initializing an SDAID trace, has been modified. For details, see page 111.
- The LSERV system utility program has two new parameters, ALL and FREE. There are also changes in the use of the PARSTD parameter. For details, see “LSERV: Display Label Information Area” on page 230.
- The information concerning Unattended Node has been removed.

Understanding Syntax Diagrams

This section describes how to read the syntax diagrams in this manual.

To read a syntax diagram follow the path of the line. Read from left to right and top to bottom.

- The **▶—** symbol indicates the beginning of a syntax diagram.
- The **—▶** symbol, at the end of a line, indicates that the syntax diagram continues on the next line.
- The **▶—** symbol, at the beginning of a line, indicates that a syntax diagram continues from the previous line.
- The **—▶◀** symbol indicates the end of a syntax diagram.

Syntax items (for example, a keyword or variable) may be:

- Directly on the line (required)
- Above the line (default)
- Below the line (optional)

Uppercase Letters

Uppercase letters denote the shortest possible abbreviation. If an item appears entirely in uppercase letters, it can not be abbreviated.

You can type the item in uppercase letters, lowercase letters, or any combination. For example:

▶—KEYWOrd—▶◀

In this example, you can enter KEYWO, KEYWOR, or KEYWORD in any combination of uppercase and lowercase letters.

Symbols

You **must** code these symbols exactly as they appear in the syntax diagram


- * Asterisk
- :
- ,
- = Equal Sign
- Hyphen
- // Double slash
- () Parenthesis
- .
- + Add

For example:

* \$\$ LST

Variables

Highlighted lowercase letters denote variable information that you must substitute with specific information. For example:



Here you must code USER= as shown and supply an ID for user_id. You may, of course, enter USER in lowercase, but you must not change it otherwise.

Repetition

An arrow returning to the left means that the item can be repeated.



A character within the arrow means you must separate repeated items with that character.



A footnote (1) by the arrow references a limit that tells how many times the item can be repeated.



Notes:

- 1 Specify *repeat* up to 5 times.

Defaults

Defaults are above the line. The system uses the default unless you override it. You can override the default by coding an option from the stack below the line. For example:



In this example, A is the default. You can override A by choosing B or C.

Required Choices

When two or more items are in a stack and one of them is on the line, you **must** specify one item. For example:



Here you must enter either A or B or C.

Optional Choice

When an item is below the line, the item is optional. Only one item **may** be chosen. For example:



Here you may enter either A or B or C, or you may omit the field.

Required Blank Space


A required blank space is indicated as such in the notation. For example:

* \$\$ E0J

This indicates that at least one blank is required before and after the characters \$\$.

Part 1. Dumps of Virtual Storage

You may face system conditions in which you want to know what the contents of your system's storage is. For this, the storage data can be read out, saved in a library or on a tape, or can be printed. This processed storage data is called a **dump**.



This part of the manual describes how to retrieve a dump and how to use the saved dump for problem determination.

Which of the shown methods you use to retrieve information for problem determination depends on the error situation. For example, in case of a system wait or a system loop the stand-alone dump program would be the appropriate tool to save or print the storage contents.

Dumps of Virtual Storage

Chapter 1. General Description

This chapter describes the various types of dump, the functions you use to create dumps in general, and how to define dump sublibraries.

The types of dump described in this chapter are:

- The **ABEND dump**, initiated by
 - the system ABEND handling routines,
 - the programmer, issuing the macro DUMP,
 - the operator entering the CANCEL command;
- The **DUMP command dump**, initiated by the operator;
- The **Stand-alone dump**, initiated by the operator;
- The **SDAID dump**, initiated by the operator.

Dump Contents Overview

The output of the DUMP command, ABEND dump and stand-alone dump program contains two major parts.

- The symptom records.
- The data records.

The amount of information which is stored in these dump records depends on the function which requests the dump. Note that pages containing only zeros are not dumped explicitly.

Figure 1 gives an overview of a dump, which can reside either in a dump sublibrary or in a dump file on tape or disk.

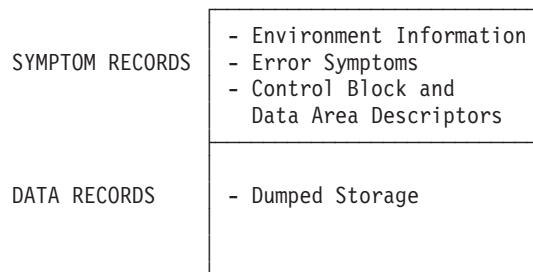


Figure 1. Overview: Dump Contents

The symptom records are built by the component which produces the dump. They contain information to format the dump data later on. The symptom records are described in Appendix A, "Symptom Records Overview," on page 217.

The ABEND Dump Function

The ABEND dump function is internally called when the system detects an ABEND condition or when a CANCEL command has been given.

Overview of ABEND Dump Function

What is an ABEND

ABEND stands for ABnormal END of task. This means that a program (task) is terminated prior to its completion because of an error that could not be resolved by system recovery facilities.

What is an ABEND Dump

The system's ABEND dump function is called by VSE/Advanced Functions:

- When an ABEND (abnormal termination) occurs;
- When a CANCEL command is issued.

When the function is called, it provides a dump of the storage areas in which the program was running.

Figure 2 shows that:

- The ABEND dump function is activated when an ABEND condition occurs;
- The output from the function is controlled by job control options. These are specified in STD OPT or OPTION statements.
- The options determine:
 - The contents of the dump;
 - To which I/O device the dump is written.

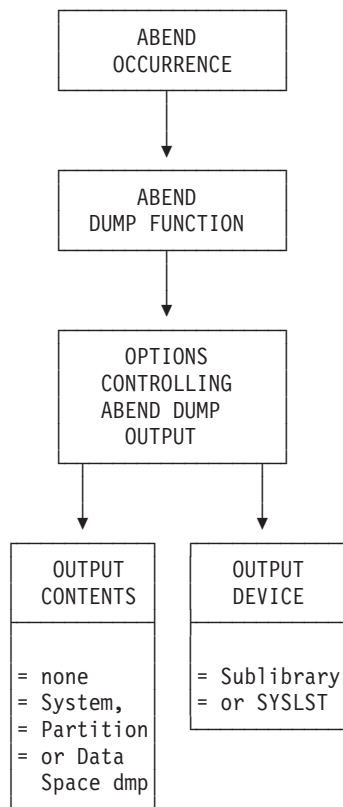


Figure 2. Overview: The ABEND Dump Function

The following ABEND dumps are shown in Figure 2 under 'OUTPUT CONTENTS':

- The **System Dump** dumps the whole supervisor area and the dump symptoms besides the partition area.
- The **Partition Dump** includes only selected VSE/Advanced Functions control blocks and the dump symptoms in addition to the partition area.
- The **Data Space Dump** includes a dump of one or more data spaces.

The output of the dump is either written into a dump sublibrary or on a printer device assigned to SYSLST.

Activation of the ABEND Dump Function

The ABEND dump function is activated when

- A program or task running in one of the system's partitions comes to an **ABnormal END**, and **no AB exit** routine is active. The macro DUMP activates the ABEND handling routines, too. See "Options to Control the ABEND Dump" on page 19.
- A **CANCEL command** is issued by the operator for one of the operating system's partitions. See "Requesting a Dump by the CANCEL Command" on page 21.

If the activation of the ABEND dump function leads to a dump writing operation (depending on the active job control options), the storage contents are dumped

- Before any end-of-job routine is executed.
- Before any of the attached subtasks is terminated.

Contents of the ABEND Dump Output

The output of the ABEND dump function (either in the sublibrary or on SYSLST) contains:

- A dump symptom part, which is always included.
- A system dump or a partition dump, depending on the options active at the time the dump was taken.
- A data space dump, if the corresponding option was specified.

Symptom Part of the ABEND Dump

This part of the output contains

- Control data from the symptom records, like information about the environment or the failure. For a description of the symptom records, see Appendix A, "Symptom Records Overview," on page 217.

System Dump

The system dump, which is produced if OPTION DUMP or STDOPT DUMP=YES is active, contains the following information besides the symptom part:

- The ending task PSW, general purpose registers, access registers and floating point registers.
- The entire supervisor area.
- The areas containing z/VSE control blocks listed in Figure 3 on page 6.
- The allocated portion(s) of the system GETVIS area.
- If the error occurred in the SVA, that part of the SVA which holds the phase responsible for the ABEND.
- The partition for which the ABEND dump function is active including areas acquired dynamically within the partition by GETVIS macros in your program.
- The dynamic space GETVIS area for dynamic partitions.

Partition Dump

A partition dump is produced when option OPTION PARTDUMP or STDOPT DUMP=PART is active. The dump output includes the following system areas besides the symptom part:

- The ending task PSW, general purpose registers, access registers, and floating point registers;
- The LOWCORE (low address storage);
- The areas containing VSE control blocks listed in Figure 3;
- The partition for which the ABEND dump function is active including areas acquired dynamically within the partition by GETVIS macros in the program;
- If the error occurred in the SVA, that part of the SVA which holds the phase responsible for the ABEND;
- The logical transient area (LTA), if the error causing the dump to be taken occurred in a task owning the LTA.

Data Space Dump

If OPTION DSPDUMP or STDOPT DSPDUMP=YES is active, a data space dump is to be taken in case of an abnormal program end. If the ABEND routine finds out that the failing program has access to a data space, it takes a dump of that space and enters it as a separate library member in the same dump library. The failing program must be in access register mode, and at least one of the access registers must contain the ALET (access list entry token) of that data space. The number of different ALET pointers in the access registers determines how many data spaces will be dumped.

The ABEND routine dumps an area of at least 4K of storage on either side of the address(es) pointed to by the matching general register(s). However, if the size of the data space does not exceed 128K of storage, the whole data space is dumped.

SUP	Supervisor
BG	Background partition
GETVIS24	24-bit GETVIS area
GETVIS31	31-bit GETVIS area
COMREG	Partition's communication region
SYSCOM	System communication region
PUBTABLE	Physical unit block table
PUBOWN	PUB ownership table
PUB2TAB	Physical unit block extension table
LUBTAB	Partition's logical unit block table
LUBEXT	Partition's LUB table extension
DIBTAB	Partition's disk information block
PIBTAB	Partition information block
PIB2TAB	Partition information block extension
PCB	Partition control block
AF-TIB	Task information block
AF-TCB	Task control block
LOADLIST	Partition's phase load trace table
LPT	Library pointer table
LDT	Library definition table
SDT	Sublibrary definition table
EDT	Extent definition table
DDT	Device definition table
LIB_ANC	Library anchor table
L-TASK-R	Librarian task LOT-row
LOTPPOOL	Library offset table pool

Figure 3. VSE Control Blocks in System Dump

The DUMP Command Dump

You can request a dump of parts of the virtual storage with the attention routine command DUMP.

For a detailed description of the DUMP command, refer to the manual *z/VSE System Control Statements*.

The Stand-Alone Dump (SADUMP) Program

If your system entered a hard or soft wait state or is in a continuous loop, no normal system operation is possible. In this case you can invoke the stand-alone dump program to get information about the problem. The stand-alone dump program records the supervisor and the SVA in one file and the page manager address spaces and the selected partitions and data spaces in separate files on one or more stand-alone dump tapes or on a disk device. (On disk, these files reside in one physical extent containing several 'logical' files.)

The output device on which the stand-alone dump program is created must be a tape or disk unit. The stand-alone dump program writes its dump output on one or more tapes or on a disk device. It is not possible to write the dump output directly on a line printer.

The stand-alone dump program dumps selected parts of virtual storage of your VSE system on tape or disk. The // OPTION SADUMP job control option allows to include important pieces of virtual storage into the stand-alone dump. It is usually not necessary or practical to dump the complete system; DUMP command dumps and partition dumps should be used when possible.

Before you can use the stand-alone dump program, it has to be created with the DOSVSDMP utility. See "Creating the Stand-Alone Dump Program" on page 25.

The creation of the stand-alone dump program should be done shortly after system installation via IUI panels in order to have the program available in case of a system error.

After the dump is taken, the operator has to perform a manual IPL from SYSRES. If the dump is on SYSRES, no manual IPL is required.

For a description of how to request a stand-alone dump, see "Taking a Stand-Alone Dump" on page 21.

Support of Integrated Console by SADUMP Program

The stand-alone dump program (SADUMP) supports the integrated console as system console in addition to 3215 and 3270 type devices. The selection of the system console depends on device availability and the IPL load parameter.

The selection criteria are:

1. If a console is specified as load parameter, the system will route messages to that preferred device.
 - In case the integrated console is specified in the load parameter, SADUMP will route messages to the integrated console.
 - In case a local console is requested in the load parameter, or the integrated console is not available, SADUMP will route messages to the local console.

2. If the stand-alone dump program is activated without specification of a communication device type, SADUMP will route messages to the device which is found in the SYSCOM.
 - If this device is not operational, SADUMP will route messages to the integrated console.
 - If an integrated console is not available, SADUMP will abnormally terminate.

IPL Load Parameter

The IPL load parameter must be used to specify the preferred communication device.

To determine the communication path, SADUMP first analyzes the load parameter. If the hardware does not support the load parameter, selection of the communication device is determined during the creation of the Dump program by the DOSVSDMP utility.

If the load parameter is specified, its first byte indicates the **console type**.

For details of the IPL load parameter, refer to the manual *z/VSE System Control Statements*.

Output of the Stand-Alone Dump Program

The stand-alone dump program stores the dump information on the tape or where it has been loaded or on a disk extent. It produces a main dump file of the system areas and additional dump files for the page manager address spaces and for each partition and data space to be dumped.

Main Dump File

The 'main' dump file is always file 3 on tape (files 1 and 2 are used by the system) or file 1 on disk. The stand-alone dump program writes the following information into the main dump file:

- The symptom record, which holds information on the hardware and software environment, error symptoms, and control block locators.
- The dump data, which consists of retrieved pages from processor storage, or from the page data set. It includes the shared area (supervisor, system GETVIS area, and SVA) and control block locators for supervisor control blocks.
- If certain system information needed for accessing the page data set is not available, you get a dump of the data in processor storage only.
- The last 200 messages from the hardcopy file.

The main dump file can be onloaded into a VSE dump sublibrary from which it can be processed by Info/Analysis. The Info/Analysis exit routines IJBXCSMG, IJBXDEBUG and IJBXSDA can be invoked to analyze the main dump file.

Page Manager Address Space (PMRAS) Dump Files

The first file (PMRAS-R) contains real storage areas which are used by the Page Manager but are not mapped in any of the virtual spaces. The following files (PMRAS-*nn*) contain the Page Manager Address Spaces (segment tables, page tables etc.), where *nn* is the space id.

Partition and Data Space Dump Files

The partitions and data spaces are dumped as separate files.

With the job control // OPTION SADUMP=*n* statement you can indicate the order or priority in which the partitions should be dumped. The format // OPTION SADUMP=(*[n],m*) indicates (for the duration of the particular job) the order or priority (*n*) in which the partition and/or the owned *data spaces* (*m*) should be dumped. The values for *n* and *m* can be 0 to 9. 0 (which is also the default) indicates that the partition or data space should **not** be dumped when the stand-alone dump is taken, 9 indicates the highest priority. The partitions and the data spaces with the highest priority (starting with 9) are always dumped first, then those with the next lower priority, until all partitions and data spaces with SADUMP not equal to zero have been dumped.

With the job control STDOPT SADUMP command you can specify a priority for **all** partitions and/or data spaces in the system.

Examples:

```
F1 ... SADUMP=(5,5)
F2 ... SADUMP=4
```

Dumps: F1 partition, F1-owned data space(s), F2 partition

```
F1 ... SADUMP=(5,3)
F2 ... SADUMP=4
F3 ... SADUMP=(,9)
```

Dumps: F3-owned data space(s), F1 partition, F2 partition, F1-owned data space(s).

Note that for stand-alone dumps to disk the stand-alone dump program stops dumping when the dump data set becomes full. Therefore, it is possible that one or more of the partitions or data spaces with SADUMP not equal to 0 will not be dumped or that the last dump file may be incomplete. This does not apply for stand-alone dumps to tape, since the output can be written to several tapes.

All dump files can be processed via DOSVSDMP which allows to print the contents of an appended dump file on SYSLST.

All dump files can also be onloaded into a VSE dump sublibrary from which they can be processed by the Info/Analysis program. You can display the symptom string or print selected parts of the storage dump.

The Info/Analysis exit routines IJBXCSMG, IJBXDEBUG and IJBXSDA cannot be invoked to analyze the appended dump files.

A description of how to print the stand-alone dump program output can be found under "Printing a Dump Stored on Tape or Disk" on page 205.

The SDAID Dump

The SDAID program can also be used to dump virtual storage. You may use this program for example if you need a dump of a certain part of storage at a defined event.

For a short description of this SDAID function, see "Requesting a Dump on Event (SDAID Dump)" on page 23.

Dump Requested by Macros

A dump of virtual storage can also be requested through dump macros.

For a short description of this method of requesting a dump, see “Requesting a Dump from a Program” on page 23.

Info/Analysis

Info/Analysis is a component of VSE. It is a tool to:

- Manage the dump files
- Print or display dump information.

With Info/Analysis, you can simplify the task of using dump data to solve software problems. Info/Analysis assists you in this task through the following functions:

- Dump management - to list the dumps being managed by Info/Analysis, to add or delete dumps from that list, and to delete dumps from the system.
- Dump symptoms - to display problem failure information collected by the dumping component and by subsequent analysis routines.
- Dump viewing - to display dump data in hexadecimal and character format, to format control blocks and other dump data that may be relevant to the problem, to invoke dump analysis routines, and to display the results of those routines.
- Dump offload - to copy a dump to tape for later retrieval.
- Dump onload - to copy a dump to a dump sublibrary (a stand-alone dump for example).

You enter input either from SYSIN or from SYSLOG. Output always goes to SYSLST. For an example of a job to invoke Info/Analysis, see Figure 59 on page 176.

For more information on Info/Analysis refer to: Part 4, “Info/Analysis,” on page 167.

Chapter 2. Maintaining the Dump Library and File Environment

Various files are used to process and evaluate dumps stored either on a tape or disk volume or in a dump sublibrary.

The Library and Files Required to Process Dumps

The libraries and files required to process and use dump information are:

1. The dump sublibraries (in the library SYSDUMP)
2. The dump management file (for Info/Analysis)
3. The external routines file (for Info/Analysis).

The SYSDUMP Sublibraries

The system uses the defined dump sublibraries to store dumps for later processing. These sublibraries are also used to onload dumps which have been stored on tape either by the system's dump functions or by a previous Info/Analysis offload operation.

Dumps can be processed using the print, analyze, and management functions of the Info/Analysis program once the dumps have been onloaded into a dump sublibrary. How the SYSDUMP library is defined and used is described in the following section.

Purpose of the SYSDUMP Library

The library named SYSDUMP is used to store the various dump types for further processing. It contains one or more dump sublibraries. Each dump sublibrary should be assigned to one partition and may contain one or more dumps. A separate dump sublibrary is used for all dynamic partitions. Figure 4 gives an overview of the SYSDUMP library concept.

Maintaining Dump Library and Files

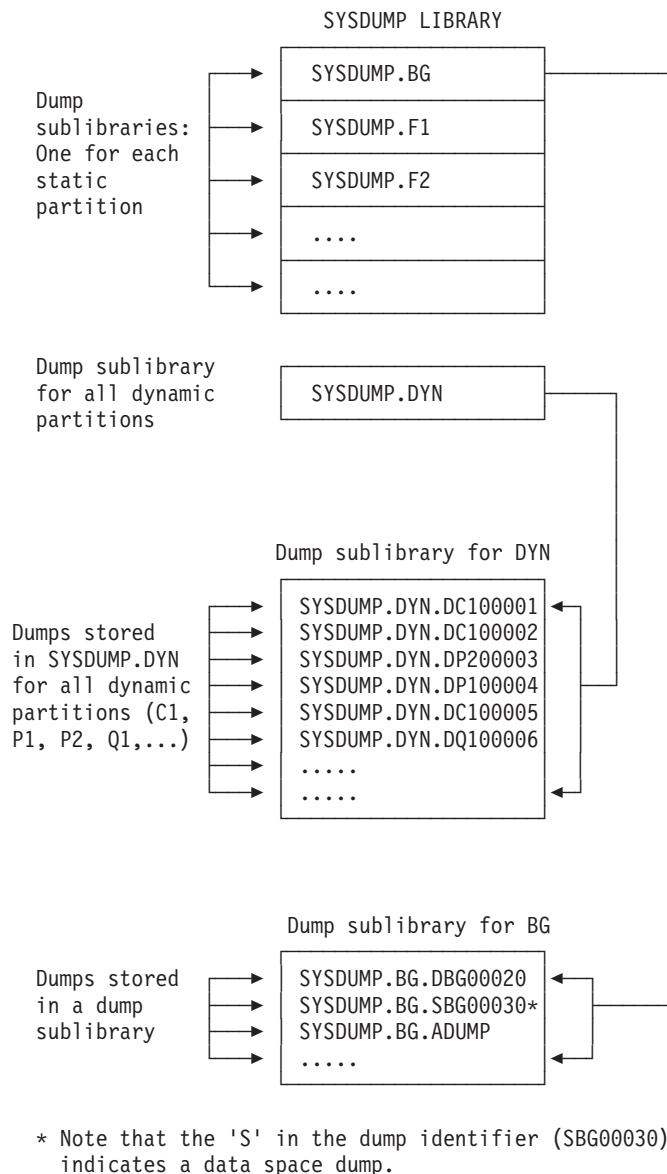


Figure 4. The SYSDUMP Library Concept

These dump sublibraries are used by the system and by you.

VSE/Advanced Functions stores dumps for later processing from

- ABEND events
- CANCEL commands.

You can use the dump sublibraries to onload dumps which have been stored on tape or disk in order to process them with Info/Analysis functions. You may use the dump sublibraries to store the following:

- DUMP command dumps
- Stand-alone dumps (from tape or disk)
- Dumps which have been offloaded to tape.

Establishing the Dump Sublibraries

Before dumps can be stored, the dump sublibraries have to be created. The following describes, what job control label information is required for the SYSDUMP library and how the dump sublibraries can be defined via the librarian program LIBR.

The following *requirements* have to be met if you want to use the dump library and its sublibraries:

1. The DLBL and EXTENT labels for the library SYSDUMP have to be specified.
2. The library SYSDUMP and its sublibraries have to be defined with the LIBR program.
3. LIBDEF statements have to be given.
4. The SYSDUMP option has to be set in order to get ABEND dumps written into the dump sublibraries.

Label Information for SYSDUMP

Figure 5 shows an example of the label and extent information you have to submit if you want to define the dump library SYSDUMP.

The standard label area should be used to store this information.

```

... ..
// DLBL SYSDUMP,'VSE.DUMP.LIBRARY',1999/365,,DSF
// EXTENT SYS010,,1,0,3150,600
... ..

```

Figure 5. Example: Labels for the SYSDUMP Library

Note: IBM recommends securing the dump library. Securing the dump library prevents overwriting of part of the file(s) as a result of a faulty response to an OVERLAPPING EXTENT message. For information about using the access control function, see the chapter “Protecting Data” in the manual *z/VSE Guide to System Functions*.

Defining the Dump Library

You define the dump library (normally named SYSDUMP) with the LIBR program. Figure 6 shows an example of such a definition.

```

// JOB DEFINE
// EXEC LIBR
DEFINE L=SYSDUMP
DEFINE S=SYSDUMP.BG -
        SYSDUMP.F1 -
        SYSDUMP.F2 -
        SYSDUMP.Fn -
        SYSDUMP.DYN REUSE=IMM
... ..
/*
/&

```

Figure 6. Example: Defining SYSDUMP with the LIBR Librarian Program

Maintaining Dump Library and Files

LIBDEF Statement for Dump Sublibraries

To get the dumps stored into the sublibrary assigned to the partition, the ASI Job Control procedure for each partition should contain a LIBDEF statement as shown in Figure 7. In the example given in Figure 7 a dump sublibrary is connected to the BG partition.

```
// LIBDEF DUMP,CATALOG=SYSDUMP.BG,PERM
```

Figure 7. Example: LIBDEF Statement for a Dump Sublibrary

Chapter 3. Handling Dumps

Options to Activate Dump Writing

The system writes the output of an automatically invoked ABEND dump into the dump sublibrary for the partition if you submit either of the following statements:

```
// STDOPT SYSDUMP=YES  
// OPTION SYSDUMP
```

With the // STDOPT SYSDUMP statement you request the system to write dumps of the next and all subsequent jobs or job-steps into the dump sublibrary for the particular partition until the SYSDUMP option is deactivated. The STDOPT statement must be given in the BG partition and is active for all partitions.

You can display the current settings for the *permanent options* using the QUERY STDOPT command. The // OPTION statement is active only for the duration of the particular job (this is the *temporary option*). After EOJ, the permanent option given in a previous STDOPT statement will be active again.

Options to Deactivate Dump Writing

The SYSDUMP option is deactivated by:

```
// STDOPT SYSDUMP=NO  
// OPTION NOSYSDUMP  
UNBATCH (to deactivate the partition)  
LIBDROP DUMP,PERM
```

Identifying the Stored Dumps

Once the dump library and dump sublibraries have been defined, dumps from various sources can be stored there. The dumps stored by the ABEND dump routines have an identifier of the following format:

```
SYSDUMP.partition_id.nnnnnnnn
```

SYSDUMP

Dump library name.

partition_id

Sublibrary name, normally the partition identifier, like BG or F3 or, for dynamic partitions, DYN.

nnnnnnnn

Dump identifier of the form:

Dppnnnnnn

for address space dumps, or

Sppnnnnnn

for data space dumps.

pp = partition identifier of the static or dynamic partition.

n = integers between 0 and 9 which are maintained by the system automatically with every new store dump operation.

For example:

```
SYSDUMP.F4.DF400002
```

is the name of a dump residing in the dump sublibrary for the F4 partition of the library SYSDUMP, with the identifier DF400002.

Note: When you onload a dump into the dump library via Info/Analysis, you select a dump name by your choice. The rules for creating a dump name are explained in “Recommendations (Restrictions) for the Generation of Dump Names” on page 179

Sending Dumps to IBM Support Electronically

Dumps are normally stored by z/VSE in the VSE dump library. However, you may need to transmit dumps *electronically* to other locations, such as to IBM Support (previously, you were required to send dumps to IBM on a physical tape).

Step 1. Locate the Dump Location on z/VSE. If the dump is a standalone dump, it will be stored on tape. If the dump is a system dump (for example, an Attention Routine dump), it will be stored on a disk or tape. In both cases, you must upload the dumps to the z/VSE dump library.

Step 2. Upload the Dumps to the z/VSE Dump Library. You use the *Storage Dump Management* dialog (Fastpath 43) to do so. For details, refer to the manual *z/VSE Guide for Solving Problems*. For a standalone dump, there will be multiple dumps consisting of one for each selected partition, and one for the supervisor and SVA.

Step 3. Format and Print Dumps (Optional). On the z/VSE dump library, you might be required to use the *Interactive Interface* tools to format these dumps (for example CICS dumps). However, in most cases IBM Support will require *unformatted* dumps.

Step 4. Download the Dumps From z/VSE to Your PC. To download dumps from z/VSE to your Personal Computer (PC), you use the *File Transfer* utility of the Interactive User Interface (IUI). For the example of “Identifying the Stored Dumps” on page 15, you would enter this command at the PC:

```
receive DFH400002.dump a:DFH400002 dump (file=lib l=sysdump s=f4 binary
```

In this example, “a:” is the emulation session where you are signed on to CICS.

As an alternative to using the File Transfer utility of the IUI, you can use the FTP of TCP/IP to transfer dump to transfer the dump from the z/VSE dump library to your PC. For the example used in this procedure, a transfer using FTP would appear like this:

a. Define the z/VSE Dump Library to TCP/IP:

```
DEFINE FILE,TYPE=LIBRARY,DLBL=SYSDUMP,PUBLIC='SYSDUMP',ALLWSITE=NO
```

b. Start the File Transfer (in Binary):

```
C:\>ftp 9.164.155.2 <----- Your IP address  
Connected to 9.164.155.2.
```

```
:
```

```
User (9.164.155.2:(none)): sysa <----- your user id  
331 User name okay, need password.
```

```
Password: <----- your password
```

```
200 Command okay.
```

```
ftp> cd sysdump
```

```
250 Requested file action okay, completed.
```

```

ftp> cd f4          <----- Sublib where your dump resides
250 Requested file action okay, completed.
ftp> bin           <----- switch to binary mode
200 Command okay.
ftp> get DF400002.dump <----- Name of the dump
200 Command okay.
150-About to open data connection
File: SYSDUMP.F4.DF400002.DUMP
Type: Binary Recfm: S Lrecl: 4096
CC=ON UNIX=OFF RECLF=OFF TRCC=OFF CRLF=ON
150 File status okay; about to open data connection
226-Bytes sent: 23,273,920
Records sent:          711
Transfer Seconds:      6.91 ( 3,788K/Sec)
File I/O Seconds:     5.30 ( 4,545K/Sec)
226 Closing data connection.
ftp: 23273920 bytes received in 7,51Seconds 3098,64Kbytes/sec.

```

Step 5. Send the Dump to IBM Support. Depending on the size of the dump stored on your PC, you might need to compress the dump using the PKZIP utility. To send the dump to IBM Support, you should send the dump to an IBM FTP server using the FTP of TCP/IP. For the example used in this procedure, you would send the dump to the public server at IBM Boulder using these commands:

```

C:\>ftp testcase.boulder.ibm.com <----- IBM Boulder IP address
Connected to testcase-blue.boulder.ibm.com.

:
User (testcase-blue.boulder.ibm.com:(none)): anonymous
331 Guest login okay, need password.
Password:          <----- your Internet e-mail address
200 Command okay.
ftp> cd /vse/toibm <----- cd vse [enter] cd toibm [enter]
250 Requested file action okay, completed.
ftp> bin           <----- switch to binary mode
200 Command okay.
ftp> put DF400002.dump <----- Name of the dump
200 Command okay.
150-About to open data connection
File: SYSDUMP.F4.DF400002.DUMP
Type: Binary Recfm: S Lrecl: 4096
CC=ON UNIX=OFF RECLF=OFF TRCC=OFF CRLF=ON
150 File status okay; about to open data connection
226-Bytes sent: 23,273,920
Records sent:          711
Transfer Seconds:      6.91 ( 3,788K/Sec)
File I/O Seconds:     5.30 ( 4,545K/Sec)
226 Closing data connection.
ftp: 23273920 bytes received in 7,51Seconds 3098,64Kbytes/sec.

```

Mailing Dumps That Are Stored on Tape to IBM Support

Dumps are normally stored by z/VSE in the VSE dump library. However, you may need to mail dumps that are stored *on tape* to other locations, such as to IBM Support.

Step 1. Locate the Dump Location on z/VSE.

- If the dump is stored on disk, go to Step 2 below.
- If the dump is stored on tape and you wish to format the dump, go to Step 2 below.
- If the dump is stored on tape and you do **not** wish to format the dump, mail the tape to the address provided by IBM Support.

- Step 2. **Upload the Dump From Disk to the z/VSE Dump Library.** You use the *Storage Dump Management* dialog (Fastpath 43) to do so. For details, refer to the manual *z/VSE Guide for Solving Problems*.
- Step 3. **Format and Print the Dump (Optional).** On the z/VSE dump library, you might be required to use the *Interactive Interface* dialog to format these dumps (for example CICS dumps). However, in most cases IBM Support will require *unformatted* dumps.
- Step 4. **Download the Dump From the z/VSE Dump Library to Tape.** To download a dump from the z/VSE dump library to tape, you use the *Interactive Interface* dialog.
- Step 5. **Mail the Tape to IBM Support.** You should mail the tape to the address provided by IBM Support.

Handling a Dump Library Full Condition

Information is written to the dump sublibraries when:

1. A dump is taken automatically by the system;
2. A dump is stored in a sublibrary via the onload process of Info/Analysis dump management;
3. Dumps are examined with an Info/Analysis analysis routine function.

In all three cases, space is needed in the dump sublibrary for page maps and analysis information, in addition to the dumps themselves. How the system reacts to a library-full condition depends on which routine caused the condition, and what kind of information was being written at the time.

If the library becomes full while:

- The system is writing an ABEND dump into it, the whole dump is printed on SYSLST and a dump-library-full information message is issued on SYSLOG;
 - If you want to *ignore* the dump (rather than printing it on SYSLST) you must specify the *SYSDUMPC* option in addition to specifying *SYSDUMP*.
 - Refer to the IBM manual *z/VSE System Control Statements* for details of how to specify the *SYSDUMPC* option.
- The Info/Analysis dump management function is writing a dump, the dump is flagged “to be onloaded”. In spite of this flag, the dump may have been stored in the sublibrary. This can happen when the library-full condition arises while additional information is being stored after the dump itself has been written.
- An Info/Analysis dump viewing function is being used, the function fails.

In all three cases, the amount of free space in the sublibrary is kept as it was before the dump write operation was started.

You can clear sublibrary space to make room for new dumps by deleting dumps that are no longer required. How to delete a dump is described under “DELETE - Delete Current Dump” on page 180.

Note: Do not delete a dump with a delete function other than the Info/Analysis delete function.

Chapter 4. Requesting a Dump

VSE/Advanced Functions offers various functions with which storage areas can be dumped. These functions differ in their output contents, output device, and way of activation. You may use these functions to isolate system program or application program errors.

Overview of Dump Requests

The table in Figure 8 summarizes the dump functions offered by VSE/Advanced Functions. The table may help you to find the dump request function which is the most effective one for your particular error situation.

Initiated by/via	Output Contents	Output Device	Requesting Function
System (ABEND)	System, Part., Data Space Dmp	Dump Sublib. or SYSLST	OPTIONS to Request the Dump
Operator (Console)	System, Part., Data Space Dmp	Dump Sublib. or SYSLST	CANCEL Command
Operator (Console)	Selected Storage Areas	Tape or Printer	DUMP Command
Operator (Console)	System Storage	Tape or Disk	STAND-ALONE DUMP Program
Programm./Oper. (Defined Event)	Selected Storage Areas	Tape, Printer or Buffer	SDAID Dump Trace
Programmer (Macro)	Macro Dependent	Macro Dependent	MACROS (PDUMP, DUMP, JDUMP, SDUMP, SDUMPX)

Figure 8. Dump Requesting Functions

Note that the base structure of the ABEND dump, the DUMP command dump, and the stand-alone dump is shown under "Dump Contents Overview" on page 3.

Each of the dump requesting functions listed in Figure 8 is described in the following sections.

Options to Control the ABEND Dump

The ABEND dump function is internally called when the VSE/Advanced Functions system detects an ABEND condition or when a CANCEL command has been given (see the following section).

Using the job control options shown below you can define whether you want to suppress a dump, or which kind of dump you want to take, and whether you want the dump to be stored in a dump sublibrary or printed on a particular output device.

Use the STDOPT command or statement to specify options for **all** jobs in the system. This must be entered in the BG partition, but it affects all partitions. Use the // OPTION statement to override these options for one job. For a detailed description of the OPTION statement, please refer to *z/VSE System Control Statements*

Options to Control the Dump Contents

The options controlling the **dump contents** can be set with the STDOPT or the OPTION statement.

STDOPT DUMP=YES

Requests a system dump.

STDOPT DUMP=PART

Requests a partition dump.

STDOPT DUMP=NO

Suppresses the ABEND dump.

STDOPT DSPDUMP=YES

Requests a data space dump.

OPTION DUMP

Requests a system dump.

OPTION PARTDUMP

Requests a partition dump.

OPTION NODUMP

Suppresses ABEND dump.

OPTION DSPDUMP

Requests a data space dump.

Note: You will not get any dump output if you use the STXIT PC or STXIT AB macro, even if you include the DUMP or PARTDUMP option.

Options to Control the Output Destination

The options controlling the **output destination** can be set with the STDOPT or the OPTION statement.

STDOPT SYSDUMP=NO

Dump to SYSLST

STDOPT SYSDUMP=YES

Dump to Library

OPTION NOSYSDUMP

Dump to SYSLST

OPTION SYSDUMP

Dump to Library

The output of the ABEND dump is either written into the dump sublibrary for the partition or it is printed on SYSLST.

The ABEND dump function writes the output to a *dump sublibrary* if the:

- **Dump Library**(named SYSDUMP) and appropriate **sublibrary** has been created.
- **LIBDEF statement**for the dump sublibrary has been submitted (usually during the ASI procedure for the partition).

- **Job Control option**
STDOPT SYSDUMP=YES or OPTION SYSDUMP has been specified.
- Associated dump library is **not full**.

If one of the above is not true the dump is printed on *SYSLST*.

Note: The dump is lost if it cannot go to SYSDUMP, and SYSLST has not been assigned. Also, the output of the ABEND dump routine is suppressed if SYSLST is assigned to a CKD-type disk device.

The contents of a system or a partition dump are described under “Contents of the ABEND Dump Output” on page 5.

How the SYSDUMP library can be defined is described under “Establishing the Dump Sublibraries” on page 13.

For a description of how to print the ABEND dump from a dump sublibrary see “Printing Dump Information” on page 183.

Requesting a Dump by the CANCEL Command

The CANCEL command, when used as a job control command cancels the execution of the current job in the partition in which the command is given. No dump is produced by the CANCEL job control command.

A detailed description of the options for the CANCEL command is given in the manual *z/VSE System Control Statements*. How to print a CANCEL command dump (ABEND dump) from a dump sublibrary is described under “Printing Dump Information” on page 183.

Requesting a Dump by the DUMP Command

The DUMP command causes selected areas of virtual address space or data space storage to be dumped.

A detailed description of the options for the DUMP command is given in the manual *z/VSE System Control Statements*. How to print the DUMP command output from tape is described under “Printing a Dump Stored on Tape or Disk” on page 205.

Taking a Stand-Alone Dump

The following steps describe how to invoke the dump process using the stand-alone dump program. Note, however, that the procedure outlined below is only a generalized description of the dump process. For detailed information on the actual steps to be performed please consult the appropriate manual of your processor.

Caution

Do not reset (clear) the processor storage before taking the dump.

1. Do a STORE STATUS.

Note: If your z/VSE system runs under z/VM, you must first issue the CP SET RUN OFF command and then the CP STORE STATUS command.

With the STORE STATUS step you save machine information that would otherwise be lost. This information is essential for error diagnosis.

- Record the contents of low-address storage bytes X'00' to X'17'. Use the hardware DISPLY/ALTER function outlined under "Hardware Alter/Display" on page 236. To interpret the data stored in these bytes refer to "VSE/Advanced Functions Codes and SVC Errors" in the *z/VSE Messages and Codes* manual.
- Mount a stand-alone dump tape (if the output is to be written on tape).
- IPL the stand-alone dump tape or disk.

Note: SADUMP supports the IPL load parameter. It may be used to specify the preferred communication device.

To determine the communication path, SADUMP first analyzes the load parameter. If the hardware does not support the load parameter, selection of the communication device is determined during the creation of the dump program by the DOSVSDMP utility.

Caution

Do not reset (clear) the processor storage at this point.

The system now takes a stand-alone dump. The following message will be issued:

```
4G34I z/VSE STAND-ALONE DUMP IN PROGRESS ON TAPE cuu | DISK cuu
```

The following message indicates the end of the dump operation:

```
4G10I STAND-ALONE DUMP COMPLETE
```

If a problem occurred during processing, the following message is issued:

```
4G35I PROBLEM ENCOUNTERED DURING SA DUMP PROCESSING. REASON CODE nnnn
```

If the dump is on tape or on a work disk, the system enters a hard wait at dump completion. If the dump is on SYSRES, VSE is re-IPLed.

You need not regenerate the stand-alone dump program after it has been used. The dump program remains useable for all subsequent stand-alone dump requests.

Note: A stand-alone dump tape created on an IBM 9346 tape drive can only be used once. After the tape has been IPLed, the DOSVSDMP create function must be used to recreate the stand-alone dump tape before it can be IPLed again.

Incorrect information in the system may result in only a dump of processor storage being taken. The stand-alone dump program collects only those pages which are in processor storage at that moment, without address translation. Possible causes are, among others:

- Low core overlaid
- SYSCOM overlaid
- Page or segment tables not available or invalid.

Incorrect information in the system may also prevent the program from issuing messages on SYSLOG.

The output of the stand-alone dump program can be written on more than one dump tape. At end-of-volume the stand-alone dump tape will be rewound and unloaded. An information message will be issued to the console:

```
4G36I END OF VOLUME ON DUMP TAPE cuu. MOUNT NEW TAPE OR RE-IPL VSE
```

The stand-alone dump program will not wait for a reply. As soon as the new tape becomes ready, the dump will continue. If the operator decides to terminate stand-alone dump processing, he just re-IPLs VSE.

This multiple-tape support also allows stand-alone dump processing to continue if a tape error occurs in the middle of a tape. The operator will receive the following message:

```
4G37I ERROR ON DUMP TAPE cuu. MOUNT NEW TAPE OR RE-IPL VSE
```

The stand-alone dump program and its output are described under “The Stand-Alone Dump (SADUMP) Program” on page 7.

The creation of the stand-alone dump program is described under “Creating the Stand-Alone Dump Program” on page 25.

A description of how to print the stand-alone dump can be found under “Printing a Stand-Alone Dump with Info/Analysis” on page 206.

Requesting a Dump on Event (SDAID Dump)

You may define that a dump has to be produced whenever a certain trace event occurs. The OUTPUT definition of the SDAID program is used for this purpose.

The following SDAID specifications for dump areas are possible:

- Partition
- Phase
- Area specified by storage addresses
- Area addressed by a register
- Area addressed by a pointer
- Control blocks or tables addressed by name.

The SDAID program is fully described in Part 3, “SDAID Trace.”

Requesting a Dump from a Program

DUMP Macro, JDUMP Macro, PDUMP Macro, SDUMP Macro, SDUMPX Macro

VSE supports the requesting of address space or data space dumps through dump macros. These macros may be issued in any program written in assembler language.

If your program issues the macros DUMP or JDUMP, VSE/Advanced Functions terminates task processing and dumps the contents of the entire supervisor plus the used part of the system GETVIS area, or, if the options DUMP=NO (NODUMP) or DUMP=PART (PARTDUMP) are active, some supervisor control blocks plus the registers and the contents of the partition that issued the macro.

The PDUMP macro provides a dump of the general registers and of the storage area you defined with the macro operands on SYSLST regardless of the active options. Note however, if SYSLST is assigned to a CKD-type disk device, no output will be produced.

Detailed information on the output device and the output contents of the dump macros and the STXIT macro, are given in *z/VSE System Macros Reference*.

Printing the Stored Dump

To print dumps stored on tape/disk or in a partition's dump sublibrary use Info/Analysis.

For information on Info/Analysis refer to Part 4, "Info/Analysis," on page 167.

Archiving Expired or Unrequired Dumps

You may place dumps in the dump archives that is provided by the *Storage Dump Management* dialog.

You can also use REXX procedure DMPMGR to regularly delete expired dumps or dumps that are no longer required.

For details of both these facilities, refer to the manual *z/VSE Guide for Solving Problems*.

Chapter 5. The DOSVSDMP Utility

This chapter describes the functions of the DOSVSDMP utility that is used in problem determination.

The DOSVSDMP Utility Functions

The DOSVSDMP utility is used to create the stand-alone dump program with which virtual storage can be dumped. The utility can also be used to print the output of the DUMP command, the stand-alone dump program (from tape or disk), the SDAID program, and IPL diagnostic information.

Run the DOSVSDMP utility in a partition with at least 192K of virtual storage.

Functions of the DOSVSDMP Utility

The DOSVSDMP utility includes the following functions:

- “Creating the Stand-Alone Dump Program” (see below)
- “Dumps Printed with DOSVSDMP” on page 29, which describes how a DUMP command dump or a stand-alone dump can be printed.
- “Printing an SDAID or DUMP Command Produced Tape” on page 30.

Creating the Stand-Alone Dump Program

The stand-alone dump program is mainly used in case of a hard or soft wait or if a system loop occurred. You can generate the stand-alone dump program to reside on magnetic tape or disk (a virtual disk, or a card or diskette unit is not valid as program residence.)

It is recommended to create the stand-alone dump program on tape or on a work disk. If you create the stand-alone dump program on your SYSRES disk, then any IPL request first causes a stand-alone dump to be taken. When the dump program has completed execution, it transfers control to the IPL program. If a dump is not needed, you can avoid the time consuming stand-alone dump processing by selecting the option CLEAR on the program load panel. The option CLEAR defines a fast path through the stand-alone dump program which will immediately transfer control to the IPL program of z/VSE.

If you create the stand-alone dump program on disk, two data sets (IJSYSDI and IJSYSDU) are required, as described under “Dump Program File and Dump Data Set” on page 27.

For processing the dump see “Printing a Dump Stored on Tape or Disk” on page 205.

To generate a stand-alone dump program, invoke DOSVSDMP by entering
`// EXEC DOSVSDMP`

The program, once it receives control, prompts you for further control information as shown in Figure 9 on page 26.

Prompt Message

```
4G01D SELECT ONE OF THE FOLLOWING FUNCTIONS:
1 CREATE STAND ALONE DUMP PROGRAM
2 SCAN DUMP TAPE/DISK
3 PRINT DUMP TAPE/DISK
4 PRINT SDAID TAPE
5 PRINT IPL DIAGNOSTICS
R END DOSVSDMP PROCESSING
```

Enter **1** to create a stand alone dump program on tape or disk. The DOSVSDMP utility responds with

Prompt Message

```
4G04D SPECIFY ADDRESS OF DUMP DEVICE (CUU OR SYSNNN)
```

The device defined with SYSNNN or CUU can be a tape or disk.

Note: Neither the utility DOSVSDMP nor the generated stand-alone dump program supports streaming mode on tape devices.

If the specified device address is that of a disk unit, DOSVSDMP responds with

Prompt Message

```
4G02D CREATE THE STAND ALONE DUMP PROGRAM
1 ON A WORK DISK
2 ON A SYSRES DISK
R END DOSVSDMP PROCESSING
```

Enter **1** if you want to create the stand-alone program on a (non-SYSRES) work disk. In this case DOSVSDMP creates a VTOC entry for a dump program file IJSYSDI, for which you have to specify labels (see "Dump Program File and Dump Data Set" on page 27).

Figure 9. Sample: Stand-Alone Dump Program Generation (Part 1 of 2)

Enter **2** if you want to create the stand-alone program on a SYSRES disk. In this case, no labels are required for IJSYSDI. DOSVSDMP creates the dump program within the disk extent reserved for the system library. Note, however, that if you create the stand-alone dump program on the SYSRES disk, a new stand-alone dump is taken with every subsequent IPL (unless you specify CLEAR). In both cases you have to specify labels for a dump data set IJSYSDU (see "Dump Program File and Dump Data Set" on page 27). You can remove the stand-alone dump program from the system disk by entering option 3 (Remove Stand-Alone Dump Program from a SYSRES disk) from the **Dump Program Utilities** panel of the Interactive Interface.

Figure 9. Sample: Stand-Alone Dump Program Generation (Part 2 of 2)

The completion message

Completion Message

```
4G09I DUMP PROGRAM HAS BEEN CREATED
```

indicates the successful generation of the dump program.

If the dump file is on disk, the completion message is followed by a message indicating the dump file capacity:

Capacity Message

```
4G27I DUMP FILE CAPACITY IS nnnnnnn K BYTES
```

Note: If the stand-alone dump program was created on the DOSRES or SYSWK1 disk, you have to recreate it after indirect service application. This is because during service application, the stand-alone dump program is overwritten by IPL records.

The description of how the stand-alone dump program is executed can be found under “The Stand-Alone Dump (SADUMP) Program” on page 7.

Dump Program File and Dump Data Set

Two data sets are required to create a stand-alone dump program on a disk pack: the dump program file and the dump data set. These files have to be defined on the same disk pack.

Dump Program File (IJSYSDI): If the dump program is to be created on a SYSRES disk, the dump program becomes part of the system library and you need not specify labels for the dump program file. If the dump program is to be created on a non-SYSRES disk, you have to define the required disk space explicitly and create the following labels for IJSYSDI:

```
// DLBL IJSYSDI,'VSE.DUMP.PROGRM'  
// EXTENT ,,,1,7 (for CKD)  
// EXTENT ,,,2,128 (for FBA)
```

The dump program occupies the first eight tracks of a CKD disk or the first 130 blocks on an FBA disk. Track 0 of a CKD disk and blocks 0 and 1 of an FBA disk are used for IPL records.

Dump Data Set (IJSYSDDU): The dump data set may be defined anywhere on the disk pack. Labels for IJSYSDDU are required for stand-alone dump program creation, for printing or scanning the dump data set, and for the dump onload function:

```
// DLBL IJSYSDDU,'VSE.DUMP.FILE'  
// EXTENT ,,,rel-track,no-of-tracks (for CKD)  
// EXTENT ,,,block,no-of-blocks (for FBA)
```

You need to define enough space to dump the supervisor, the shared virtual area (SVA), and space for those partitions and/or data spaces that you want to dump. If there is not sufficient space, the areas will be dumped until the space is full.

To make sure that the dump data set is large enough, calculate the amount of storage you want to have dumped, add 5% to the result, and compare it to the size provided by message 4G27I. If the size is too small, increase it and rerun the job.

If the dump data set is too small to contain a complete stand-alone dump, the remainder of the dump is dropped. The dump data set will contain only one stand-alone dump at a time. Any subsequent dump will overwrite the previous dump. ABEND dumps or attention routine dumps cannot be written into the dump data set.

Scanning the Dump Files on Disk or Tape

The SCAN function of DOSVSDMP provides a file directory of the dump tape or disk. After having invoked DOSVSDMP by entering

```
// EXEC DOSVSDMP
```

the program prompts you for further information as shown in Figure 10 on page 29.

Prompt Message

```
4G01D SELECT ONE OF THE FOLLOWING FUNCTIONS:
1 CREATE STAND ALONE DUMP PROGRAM
2 SCAN DUMP TAPE/DISK
3 PRINT DUMP TAPE/DISK
4 PRINT SDAID TAPE
5 PRINT IPL DIAGNOSTICS
R END DOSVSDMP PROCESSING
```

Enter 2 to scan the dump tape or the dump data set on disk. DOSVSDMP prints the following information on SYSLST:

1. For SCAN DUMP DISK:

DUMP FILE	DUMP TYPE	NAME	DATE	DATA DUMPED
001	SADUMP			SUPERVISOR+SVA
002	SADUMP			PMRAS-R
003	SADUMP			PMRAS-00
004	SADUMP	CICSICCF	94/11/17	F2-PARTITION
005	SADUMP	VTAMSTRT	94/11/17	F3-PARTITION
006	SADUMP	POWSTART	94/11/17	F1-PARTITION
007	SADUMP	JOB0815	94/11/17	BG-PARTITION
008	SADUMP	DATA007	94/11/17	BG-DATA_SPACE

2. For SCAN DUMP TAPE:

DUMP FILE	DUMP TYPE	NAME	DATE	DATA DUMPED
001				DOES NOT CONTAIN DUMP DATA
002				DOES NOT CONTAIN DUMP DATA
003	SADUMP			SUPERVISOR+SVA
004	SADUMP		94/11/17	PMRAS-R
005	SADUMP		94/11/17	PMRAS-00
006	SADUMP	CICSICCF	94/11/17	F2-PARTITION
007	SADUMP	VTAMSTRT	94/11/17	F3-PARTITION
008	SADUMP	POWSTART	94/11/17	F1-PARTITION
009	SADUMP	JOB0815	94/11/17	BG-PARTITION
010	SADUMP	DATA007	94/11/17	BG-DATA_SPACE

Figure 10. Sample: Directory of Dump Disk/Tape

Dumps Printed with DOSVSDMP

How to print the dumps produced by the stand-alone dump program and the DUMP command in unformatted form is discussed in this section. Normally Info/Analysis is used to process and print dump tapes. In exceptional cases the use of the DOSVSDMP utility may be necessary, for example:

- If none of your dump sublibraries are big enough to hold the stand-alone dump;
- If the dump was taken with the DUMP BUFFER,cuu command.

The printed output of the DOSVSDMP utility contains for both DUMP command tape or stand-alone dump tape/disk, the following:

- Symptom record.
- Unformatted dump data.

Sample DOSVSDMP Print Setup

To print a dump from tape or disk using the DOSVSDMP utility, invoke DOSVSDMP by submitting the control statements shown in Figure 11.

The utility prompts you by messages for further control information, which you enter at SYSLOG.

```
// JOB DOSVSDMP
// EXEC DOSVSDMP
```

DOSVSDMP prompts you by messages at SYSLOG to define the operation you want to perform, with:

Prompt Message

```
4G01D SELECT ONE OF THE FOLLOWING FUNCTIONS:
1 CREATE STAND ALONE DUMP PROGRAM
2 SCAN DUMP TAPE/DISK
3 PRINT DUMP TAPE/DISK
4 PRINT SDAID TAPE
5 PRINT IPL DIAGNOSTICS
R END DOSVSDMP PROCESSING
```

Enter 3 to invoke DOSVSDMP Print Dump Tape/Disk processing.

The DOSVSDMP utility response is:

Prompt Message

```
4G04D SPECIFY ADDRESS OF DUMP DEVICE (CUU OR SYSNNN)
```

Enter 280, for example, if the dump tape is mounted on the tape drive 280.

If you have selected option 3, DOSVSDMP also prompts you for the number of the dump file that you want to print. (Option 2 - SCAN DUMP TAPE/DISK - gives you a directory of the dump files on the dump.

See also "Scanning the Dump Files on Disk or Tape" on page 28).

Prompt Message

```
4G30D SPECIFY FILE NUMBER
```

Enter 4, for example, if you want to print file 4.

Now the DOSVSDMP utility starts printing the dump on SYSLSLST.

After print completion, control is returned to Job Control.

Figure 11. Sample: Dump Tape Printed with DOSVSDMP

Printing an SDAID or DUMP Command Produced Tape

You may specify that the SDAID trace information is to be recorded on tape. DOSVSDMP can be used to retrieve this information from tape and to print it on SYSLSLST. This is done by responding to DOSVSDMP prompts as shown in the

Figure 12. Always use this option of DOSVSDMP to print dumps produced in response to the attention routine command

```
DUMP BUFFER, cuu
```

When the utility gets control, it prompts you for further definitions via SYSLOG, as shown in the example in Figure 12.

```
// JOB SDAID  
// EXEC DOSVSDMP
```

DOSVSDMP prompts you to define the operation you want to perform:

Prompt Message

```
4G01D SELECT ONE OF THE FOLLOWING FUNCTIONS:  
1 CREATE STAND ALONE DUMP PROGRAM  
2 SCAN DUMP TAPE/DISK  
3 PRINT DUMP TAPE/DISK  
4 PRINT SDAID TAPE  
5 PRINT IPL DIAGNOSTICS  
R END DOSVSDMP PROCESSING
```

Enter 4 to invoke DOSVSDMP Print SDAID Tape processing.

The DOSVSDMP utility responds with:

Prompt Message

```
4G05D SPECIFY ADDRESS OF SDAID TAPE (CUU OR SYSNNN)
```

Enter 280, for example, if the SDAID output tape is mounted on the device 280.

The DOSVSDMP utility now responds with:

Prompt Message

```
4G30D SPECIFY FILE NUMBER
```

Enter 2, for example, if the second file contains the SDAID output you want to print.

The file number is determined by the number of STOPSD commands given in the SDAID session. (Every STOPSD command writes a tapemark on the tape if there was any trace event.)

If, for example, you issue three times STARTSD/STOPSD within an SDAID session, you get three trace files on your trace output tape.

DOSVSDMP prints the tape on the device assigned to SYSLSLST. After print completion, control is returned to Job Control.

Figure 12. Sample Job: Printing SDAID Tape with DOSVSDMP

Part 2. Interactive Trace Program

Chapter 6. Interactive Trace Program

Introduction

The interactive trace program is the tracing tool for z/VSE application programs. It traces the execution of application programs running in static or dynamic partitions. The interactive trace program is activated via the // EXEC statement and controlled interactively from the z/VSE master console or from a user console. It operates at the level of machine instructions and virtual storage addresses, similar to the CP debugging facilities in z/VM.

The interactive trace program provides the following traces:

- Branch trace
- Instruction trace
- Storage alteration trace
- ABEND trace.

Branch Trace

The branch trace monitors branch instructions. The trace program displays all branch instructions which transfer control to an address which is located within a specified storage area. That means, branches are only recorded if the **target** address of the branch is located within the specified address range.

Instruction Trace

The instruction trace monitors the instructions executed within a specified storage area. An instruction is traced if the first byte of the instruction is contained in the specified storage area. The trace program displays also EXECUTE instructions if the first byte of the target of an EXECUTE is within the designated storage area.

Storage Alteration Trace

The storage alteration trace monitors storage alterations within a specified storage area. A storage alteration event occurs even if the value stored is the same as the original value. However, monitoring does not apply if data is altered by a channel program or by system control programs.

ABEND Trace

The ABEND "abnormal end" trace allows interactive debugging if a user program terminates abnormally. In case of an ABEND, the termination routines display the cancelation message on the screen and transfer control to the console operator. The operator can inspect storage data or register contents to determine the cause of the cancelation. It is, however, not possible to change the program status and return to normal operation via the GO command (see "GO Command" on page 39) with a branch address. The task termination is already in progress at that time. The GO command can only be used to resume the termination process. It is also possible to modify the dump option to DUMP, PARTDUMP, or NODUMP.

The ABEND trace does not become active if an AB exit routine (STXIT routine) with the options EARLY or NODUMP is defined. In these cases control is

transferred to the user exit routine before the trace is invoked. The code of the exit routine can, however, be traced via an instruction trace if the code segment of the exit routine is defined as tracing range.

If the branch trace, the instruction trace, or the storage alteration trace display an instruction on the screen, this instruction has already been executed. If the trace displays an interruptible instruction, like an MVCL, the PSW, the general registers, and storage data at the time of the interruption are displayed. If the MVCL is partly processed, the PSW still points to the MVCL instruction. If the execution of the MVCL is completed, the PSW points to the instruction after the MVCL instruction. If the trace program displays an SVC instruction, the related supervisor service has not been started yet.

Trace Activation

```
▶▶ // EXEC _____ progname, TRACE _____ ▶▶  
      |  
      └─ PGM= ─┘
```

You start the interactive trace program with the parameter TRACE in the EXEC statement. An example of the trace initialization is shown in Figure 13 on page 41. The invoked trace function is active for the duration of one VSE job step.

The parameter TRACE implicitly defines an instruction trace and an ABEND trace. These trace definitions allow the console operator to get interactive control over the program to be traced. The instruction trace passes control to the console operator at the beginning of a user program, the ABEND trace allows debugging when a program terminates abnormally.

The instruction trace defined implicitly via the TRACE parameter traces all instructions executed within the partition. Trace boundaries are the partition begin and end address. The user program stops after the first instruction has been executed. The trace program displays the interrupted instruction (preceded by a reply identification) and waits for an operator response. The operator answers with an interactive trace command. The operator may use the implicitly defined traces to step through all instructions of the program, or replace these implicitly defined traces by specific trace definitions. The implicitly defined traces remain in effect until they are explicitly deleted by the operator.

Interactive Trace Commands


This chapter describes the following interactive trace commands:

- TRACE
- QUERY
- DISPLAY
- ALTER
- GO

Most of these commands may be abbreviated. The possible abbreviation is shown through lowercase letters. An example of an abbreviated command is

```
tr i addr=4037ac.1c
```


TRACE Command



The TRACE command defines the type of trace to be activated (branch trace, instruction trace, storage alteration trace, or ABEND trace). It is possible to issue up to 100 trace definition statements within one interactive tracing session.

The parameter ADDRESS defines the tracing range. It has different meanings for the different types of traces.

Branch trace:

The parameter ADDRESS= defines the branch target area.

Instruction trace:

The parameter ADDRESS= defines the storage area within which instruction execution is monitored. The address range must be part of the partition area. A tracing area outside the user partition is rejected. If the specified tracing range crosses partition boundaries, only the range located within the partition is accepted.

Storage alteration trace:

The parameter ADDRESS= defines the area within which storage alteration is to be monitored.

ABEND trace:

The parameter ADDRESS= is not applicable for the ABEND trace.

If the parameter ADDRESS is omitted, the whole user partition is assumed as tracing range. The following examples explain the different forms of the address parameter.

ADDRESS=410C1F

specifies a one-byte storage interval.

ADDRESS=460C1F.C

specifies a storage interval of 12 bytes.

ADDRESS=40031C:400328

specifies a storage interval by its virtual start address and end address.




The TRACE END command deletes one or all traces specified for a partition. The parameter 'n' addresses a trace statement by its trace identification (obtained via the QUERY command). The keyword ALL (default) deletes all traces specified for a partition.

QUERY Command




The QUERY command displays a list of all traces active for a user partition. The displayed trace-identification may be used in a subsequent TRACE END command to delete one of the specified traces. An example of the QUERY command is shown in Figure 14 on page 41.

DISPLAY Command



The DISPLAY command displays either storage data, or the general purpose registers (GR, Gn), the floating point registers (FR), the access registers (AR), or the Program Status Word (PSW). The specification DISPLAY GR displays **all** general purpose registers, the specification DISPLAY Gn, displays a particular general purpose register. An example of the DISPLAY command is shown in Figure 16 on page 43.

ALTER Command



The ALTER command allows to alter storage data or the contents of a general purpose register.

- **Altering storage data:** The *address* parameter denotes the storage address where data is to be altered. The DATA parameter describes the new storage data by its hexadecimal representation. Any two hexadecimal digits describe the contents of one byte in storage. The specified data is not padded; that means, it is required to enter an even number of hexadecimal digits. It is possible to enter up to 16 hexadecimal digits in order to alter up to 8 bytes. It is not possible to alter storage locations outside the user partition, or to alter the mask portion of the stored PSW. Example:

The specification

```
ALTER 400312 DATA=03FEC7
```

alters the contents of addresses 400312, 400313, 400314 to the values 03, FE, C7 respectively.

- **Altering the contents of a general purpose register:** The parameter Gn denotes the general purpose register *n*. The DATA parameter describes the new contents

of the specified general purpose register. The entered data is padded on the left with binary zeros. It is possible to enter up to 8 hexadecimal digits.

Example:


The specification of

```
ALTER G5 DATA=12C
```

enters the value of 0000012C into general purpose register 5.

(Another example of the ALTER command is shown in Figure 16 on page 43.)

GO Command



The GO command reactivates the stopped user program. The program continues processing at the specified address. This address is not checked for validity. If the address parameter is omitted, the program continues processing with the next sequential instruction. The interactive trace command ignores the specified branch address if the last displayed instruction is an SVC instruction, or if a cancel condition has occurred and the termination routines have already issued the termination messages. In this case the GO command has the only effect to resume the termination process.

The parameter `OUTPut=SYSLST` switches the tracing mode from interactive tracing to batch tracing. The trace output lines are printed on a printer device or they are written into the VSE/POWER list queue. Switching to batch mode is not possible if the logical unit SYSLST is unassigned or assigned to a tape or disk device.

The parameter `OPTion=` modifies the temporary dump options. In case of an abnormal termination, the dump routines will print either a full dump (DUMP), a partition dump (PARTDUMP), or no dump at all (NODUMP). If the parameter `OPTION` is omitted, the dump options remain unchanged.

If the parameters `address`, `OUTPUT`, and `OPTION` are omitted, the command name GO can be omitted, too. A reply-ID without any parameter is processed as a 'GO' command.

Tracing in a User Partition with Subtasks Attached

The interactive trace program traces the main task and all attached subtasks. If two tasks execute in the specified tracing range concurrently, both tasks are traced. Different VSE users can activate the trace program independently in different partitions. It is possible that several partitions present trace messages at the same time. Within one partition, however, only one task can present a tracing event at a time. The first task arriving at a specified tracing event locks the tracing routine for exclusive usage. The interrupted task writes a console message and waits for an operator response. The other tasks continue processing until they arrive at a tracing point. As long as the operator issues interactive trace commands, no other task can issue a tracing message. When the operator resumes program operation (via a GO command), all tasks with a pending tracing message will be activated. The subtask with the highest priority will present its tracing message at the screen.

Scope of Tracing

The interactive trace program is designed to trace user programs. It cannot be used to trace supervisor routines, Job Control statements, or attention routine commands. These restrictions have been introduced to keep the impact of tracing on the operations in other partitions to a minimum. The system routines often lock non-reentrant system resources. Any other task in the system competing for the same resource might enter a wait state until the locked resource becomes free. Therefore it is not tolerable to interrupt a system routine for interactive tracing.

A program routine which runs in a user partition and owns the Logical Transient Area (LTA), cannot be traced interactively. An example of such a user routine is an OPEN exit routine which is called from the LTA via a CALL/RETURN interface. The trace program does not stop for exit routine instructions. Programs running in other partitions might wait for the LTA to become free. It is, however, possible to trace such a routine in batch mode. You may use the GO command with the parameter OUTPut=SYSLST to trace all or selected instructions of the exit routine on SYSLST.

In exceptional cases it is possible that an interactive trace of a user written routine may have an impact on the performance of programs running in other partitions. For example, such an interference with other partitions may occur if the program uses the LOCK macro or the track-hold option of a DTF macro to synchronize processing with programs in other partitions.

Restrictions for Programs Using the PER Function

The interactive trace program uses the Program Event Recording (PER) function of the ESA/370 or the ESA/390 Architecture. It is possible to run the interactive trace program in several partitions at the same time. However, it is not possible to run the interactive trace program concurrently with another program which uses the PER function.

The Interactive Trace Program versus SDAID

Some trace types of the SDAID program use the Program Event Recording function. These trace types are the branch trace, the instruction trace and the storage alteration trace. They **cannot** run concurrent with the interactive trace program. The SDAID initialization routine checks whether an interactive trace for any partition is already active. The SDAID STARTSD command is rejected if the interactive trace program is active for any partition. (For a description of the STARTSD command see Chapter 12, "Start/Stop and End the Trace," on page 163.)

An SDAID session which does not use the PER function can run concurrently with the interactive trace program. These trace types do not interfere with the interactive trace program:

- CANCEL
- EXTERNAL
- GETVIS
- IO
- LOCK
- MON
- PGML
- PGMC
- SSCH
- SVC

- VTAMBU
- VTAMIO
- XPCC

Examples of the Interactive Trace Program

Trace Initialization Example

```
// JOB ABC
// EXEC PROG1,TRACE
/ &

BG 0000 4I01I TRACE STARTED FOR PROGRAM PROG1
BG-0000 00600078 BALR 05C0 CC 0
```

Figure 13. Trace Example: Trace Initialization

Figure 13 shows a job stream which initializes an interactive trace session for program PROG1.

The last two lines show the system response. Message 4I01I indicates that trace initialization was successful. The traced program stops its execution after the first instruction has been executed. The trace program displays the first instruction on the screen and waits for an operator response. The operator may now use the implicitly defined instruction trace to step through all instructions of the program, or replace this instruction trace by specific trace definitions.

Figure 14 shows the available commands to modify the trace environment.

TRACE, TRACE END and QUERY Command Example

```
BG+0000 00400078 BALR 05C0 CC 0
0 query
BG 0000 001 TRACE INST ADDRESS=00600000:006AFFFF
BG-0000 002 TRACE ABEND
0 trace end 1
BG-0000 4I09D SPECIFIED TRACE ENDED
0 trace inst address=403BA0.70
BG-0000 003 TRACE INST ADDRESS=00603BA0:00603C0F
0 trace stor address=4002ad
BG-0000 004 TRACE STOR ADDRESS=006002AD:006002AD
0 trace inst address=4017cc:4017ff
BG-0000 005 TRACE INST ADDRESS=006017CC:006017FF
0 query
BG 0000 002 TRACE ABEND
BG 0000 003 TRACE INST ADDRESS=00403BA0:00403C0F
BG 0000 004 TRACE STOR ADDRESS=004002AD:004002AD
BG-0000 005 TRACE INST ADDRESS=004017CC:004017FF
```

Figure 14. Trace Example: TRACE, TRACE END and QUERY Command

Figure 14 shows how a tracing environment can be modified. The operator deletes the implicitly specified instruction trace (001) and defines two new instruction traces (003 and 005) and one storage alteration trace (004). Trace 002 continues unaltered.

Batch Trace Example

```
0 exec testtrac,trace
BG 0000 4I01I TRACE STARTED FOR PROGRAM TESTTRAC
BG-0000 00400078 BALR 0530 CC 0
0
BG-0000 0040007A B 47103024 -> 0040009E CC 0
0
BG-0000 0040009E NOPR 0700 CC 0
0
BG-0000 004000A0 BAL 4510303E -> 004000B8 CC 0
0
BG-0000 004000B8 LR 1801 CC 0
0
BG-0000 004000BA SVC 0A26 CC 0
0
BG-0063 004001A4 LA 41603313 = 0040038D CC 0
63 go output=syslst
BG 0000 4I20I TRACING TERMINATED
BG-0000 1I00D READY FOR COMMUNICATIONS.
```

Figure 15. Trace Example: Batch Trace (Part 1 of 2)

```
***** START OF BATCH TRACE *****
0063 004001A8 STCM BE67359F >> 00400619 CC 0
0063 004001AC LA 41600020 = 00000020 CC 0
0063 004001B0 STC 426035A5 >> 0040061F CC 0
0063 004001B4 L 58103606 00400680 CC 0
0063 004001B8 SVC 0A00 CC 3
0063 004001BA L 58103606 00400680 CC 0
0063 004001BE TM 91801002 004005DE CC 0
0063 004001C2 BO 4710314E 004001C8 CC 0
0063 004001C6 SVC 0A07 CC 0
0000 004000BC LTR 1211 CC 2
0000 004000BE BM 47403104 0040017E CC 2
0000 004000C2 NOPR 0700 CC 2
0000 004000C4 BAL 45103062 -> 004000DC CC 2
0000 004000DC LR 1801 CC 2
0000 004000DE SVC 0A26 CC 2
0064 0040020C LA 41603333 = 004003AD CC 2
0064 00400210 STCM BE6735AF >> 00400629 CC 2
0064 00400214 LA 41600020 = 00000020 CC 2
```

Figure 15. Trace Example: Batch Trace (Part 2 of 2)

Figure 15 shows a program with sub tasks attached. The main task and the sub tasks have different reply identifications (0000, 0063, 0064). After the instruction at location 4001A4 has been executed, the operator issues the command `go output=syslst` to switch from the interactive tracing mode into the batch tracing mode. A fragment of the batch output on SYSLST is shown in the second part of the figure.

DISPLAY and ALTER Command Example

```
5 display psw
F5-0005 PSW = 470D0000 00400EE6
5 display gr
F5 0005 GPR 0 = 00000000 00400648 00003110 00400EF4
F5 0005 GPR 4 = 00403518 004001AE 004010E0 004000E0
F5 0005 GPR 8 = 00407000 004020E0 00000590 004030E0
F5-0005 GPR 12 = 00000006 00407894 804017FC 90000004
5 alter gc data=00000007
F5-0005 GPR 12 = 00000007
5 display 403017.20
F5 0005 00403010 00E1D8E0 00E1474E 00006FBC 00000644 *..Q....+..?.....*
F5 0005 00403020 00009DA0 00E263C8 00E18598 00FC4540 *.....S.H..eq....*
F5-0005 00403030 00004930 000094E8 00E623F8 00E26414 *.....mY.W.8.S..*
5 display 403017
F5-0005 00403010 00E1D8E0 00E1474E 00006FBC 00000644 *..Q....+..?.....*
5 alter 403017 data=fefefe
F5-0005 00403010 00E1D8E0 00E147FE FEFE6FBC 00000644 *..Q.....?.....*
```

Figure 16. Trace Example: DISPLAY and ALTER Command

Part 3. SDAID Trace

In order to isolate a problem in a system or in an application program, you may need to know the exact sequence of execution steps which have been performed. To find out about the execution steps performed, you have to trace specific events. VSE offers a program which helps you in tracing specific events in your system. This program is called SDAID.

SDAID traces user and system programs running either below or above the 16MB line. The tracing range and the dump area are specified as 31-bit addresses.

You can initialize the SDAID traces by:

- Answers to prompts in the attention routine (AR);
- Direct input statements to the AR or a partition;
- Job control procedures.

This part of the book describes each of the methods to initialize a trace and how the initialized trace is started/stopped or ended. Trace output examples are shown for each of the trace types.

If the SDAID program is new for you, read Chapter 7, "SDAID Overview," on page 47 as an introduction.

If you are familiar with the SDAID program conventions, read the Summary in Table 2 on page 49. This Summary shows all trace types with references to their format descriptions for the various initialization methods.

SDAID Trace

Chapter 7. SDAID Overview

This chapter gives an overview of the SDAID program and the various ways to initialize a trace; it includes considerations on the performance and the environment.

The SDAID Session

Basically you will do two (or three) things:

1. Initialize a trace.
2. Start, stop, or terminate the initialized trace.
3. Print trace data via DOSVSDMP (if trace output is on tape).

Interaction SDAID versus Interactive Trace Program

The SDAID branch trace, the instruction trace and the storage alteration trace use the Program Event Recording (PER) function of the 370/ESA or the 390/ESA Architecture. The interactive trace program also uses this hardware function. SDAID sessions which contain one of the above mentioned traces cannot be started if an interactive trace for any user partition is already active.

An SDAID session which does not use the PER function can run parallel to the interactive trace program. These trace types do **not** interfere with the interactive trace program:

- CANCEL
- EXTERNAL
- GETVIS
- IO
- LOCK
- MON
- PGMC
- PGML
- SSCH
- SVC
- VTAMBU
- VTAMIO
- XPCC

How to Initialize an SDAID Trace

You initialize a trace with mainly four statement types which have to be submitted to the SDAID program:

1. The SDAID statement to start the initialization process.
2. The OUTDEV specification to define the output device for the trace data.
3. The TRACE statements to define all necessary information for the trace, like the trace type and the area to be traced.
4. A statement which signals the end of the initialization process (/ * or READY).

You submit the SDAID statements with one of the following methods:

- Direct input mode in the attention routine or partition.

- Job control procedures in a partition.
- Prompts in the attention routine (AR).

Initialization in Direct Input Mode

In direct input mode the SDAID information is entered in the form of commands to the attention routine or as SYSIN statements in a partition.

The SDAID program identifies the mode of initialization via the format of the TRACE and OUTDEV statement. In direct input mode these statements must contain at least one operand.

The following examples show two initialization jobs, one entered in a partition the other one entered via SYSIN.

Example of a trace initialization in direct input mode in the attention routine:

```
sdaid
AR 4C05I PROCESSING OF 'SDAID'    COMMAND SUCCESSFUL.
AR 1I40I READY
outdev tape=280
AR 4C05I PROCESSING OF 'OUTDEV'   COMMAND SUCCESSFUL.
AR 1I40I READY
trace ssch unit=009
AR 4C05I PROCESSING OF 'TRACE'    COMMAND SUCCESSFUL.
trace io unit=009 output=ccw
AR 4C05I PROCESSING OF 'TRACE'    COMMAND SUCCESSFUL.
ready
AR 4C05I PROCESSING OF 'READY'    COMMAND SUCCESSFUL.
AR 1I40I READY
```

Example of trace initialization by direct input mode statements read in from SYSIN:

```
// EXEC SDAID
OUTDEV TAPE=280
TRACE SSCH UNIT=009
TRACE IO UNIT=009 OUTPUT=CCW
/*
```

The direct input mode is described in Chapter 9, "Initialize an SDAID Trace in Direct Input Mode," on page 83.

Initialization via Job Control Procedures

The easiest way to initialize a trace is to use catalogued procedures.

An example of such a trace procedure statement is shown below.

```
// EXEC PROC=SDIO,UNIT=009,TAPE=280
```

Initialization by procedures is described in Chapter 10, "Initialize an SDAID Trace via a Procedure," on page 117.

Initialization via Prompts in the Attention Routine

You start the initialization process with the attention routine command 'SDAID'. The necessary trace definitions are given in response to promptings after you entered the TRACE or OUTDEV statement without an operand.

You enter the prompt mode whenever you define these two commands without an operand. Example of a trace initialization via prompts in the attention routine:

```

sdaid □
AR 4C05I PROCESSING OF 'SDAID' COMMAND SUCCESSFUL
outdev □
AR 4C08D SPECIFY OUTPUT DEVICE.+
tape □
...

```

□ indicates the ENTER key pressed

Note that you enter the prompt mode also if you specify direct input mode statements combined with prompt mode statements like a question mark (? requests the help function of SDAID). The example below shows, how you can combine the two input modes. You would be prompted after the question mark has been processed.

```

TRACE SSCH AREA=BG ?
- direct input→|← prompt mode

```

The prompt input mode is described in Chapter 11, "Initialize a Trace in Prompt Input Mode," on page 131.

AR Commands to Start, Stop and End an Initialized Trace

You can start, stop, or end an initialized trace via attention routine (AR) commands. The table below shows an overview of these commands. A more detailed description about stopping, starting, and ending a trace is given in Chapter 12, "Start/Stop and End the Trace," on page 163.

Table 1. AR Commands to Start/Stop and End an Initialized Trace

STARTSD	Starts SDAID execution; may follow READY or STOPSD. Note: The old form of the command (STRTSD) is still accepted.
STOPSD	Suspends SDAID execution; allowed only after STARTSD.
ENDSD	Ends SDAID session; releases all system resources used by SDAID at any time.

Trace Type Summary

Find the trace type which you want to initialize in the following trace command summary. You will find page references to the description of the trace type and to the format of the trace initialization statements for the various initialization methods. Locate the section according to the initialization method you choose.

The page references to the available descriptions are under the following headings:

- Des** Page reference to the *description* of the trace type
- Dir** Page reference to the format description for initialization in *direct input mode*.
- Prc** Page reference to the format description for the initialization via *procedures*.
- Prp** Page reference to the description for the initialization in *prompt mode*.

Table 2. Trace Type Summary

Trace Type	Provides a Trace of:	Des	Dir	Prc	Prp
BRANCH*	Successfully executed branch instructions	56	89	121	145

Table 2. Trace Type Summary (continued)

Trace Type	Provides a Trace of:	Des	Dir	Prc	Prp
BUFFER	The trace buffer when it is full	57	90	-	145
CANCEL	Program (main task) cancel or EOJ	57	90	-	145
EXTERNAL	External interrupts	58	91	-	145
GETVIS	GETVIS / FREEVIS requests	58	93	-	146
INSTRUCTION	Selected or all instruction(s) execution	59	95	122	147
IO	I/O interrupts	60	96	123	148
LOCK	LOCK / UNLOCK requests	61	97	-	148
MONITORCALL	MC instructions	62	99	-	150
PGMCHECK	Program checks	63	100	125	150
PGMLOAD	Phase load requests, or actual load	64	101	124	151
SSCH	Start Subchannel instructions	65	103	123	152
STORAGE	Storage alterations	66	103	126	153
SVC	Executed supervisor calls	67	105	127	153
VTAMBU	Usage of VTAM buffers	68	107	-	154
VTAMIO	VTAM I/O operations	68	107	-	154
XPCC	Cross partition communication	69	108	-	154

* See, however, "System Performance Degradation Caused by PER Traces" on page 51.

Trace Output

The trace output, an **event record**, is supplied for each occurrence of a traced event, according to your setup instructions.

You may request the event records to be written to a line printer, onto magnetic tape, or into a wraparound buffer. How to define the output device is described together with each type of initialization process.

Sample event records are shown for each trace type under "Summary of TRACE Types" on page 56.

Performance Considerations

System Performance Degradation

The tracing of events with SDAID may affect overall system performance. This may especially affect time dependent programs (such as programs doing input/output via telecommunication lines).

As long as SDAID processes a tracing event, all external and I/O interrupts are disabled and remain pending until trace data collection is complete. The supervisor may recognize an attention interrupt from the system console immediately or with a significant delay. Specification of an output tape (OUTDEV TAPE=cuu) reduces the possible time delays.

When you invoke SDAID in prompt mode, console input is blocked during the processing of each SDAID command.

System Performance Degradation Caused by PER Traces

The following SDAID traces use the program event recording (PER) feature:

- branch trace (BRANCH)
- instruction trace (INSTRUCTION)
- storage alter trace (STORAGE).

The PER feature allows the SDAID to limit the trace address range via the control registers 10 and 11. For this, the use of the address specification (ADDRESS= in direct input mode for example) helps to achieve better performance.

SDAID Space Requirements

Space Requirements during Initialization in the AR

The SDAID setup phases are loaded into the system GETVIS space. The SDAID setup phases require approximately 100K bytes (K equals 1,024) of virtual storage. When initialization is complete (the READY command is processed successfully), that GETVIS space is released.

Space Requirements during Initialization in a Partition

Beside the GETVIS space of 100K bytes the phase SDAID (called via EXEC SDAID) requires approximately 16K bytes of partition virtual storage. This is significantly less than the minimum VSE partition size. Therefore SDAID will run in any foreground or background partition.

Space Requirements for SDAID Execution

When the READY command is entered, SDAID allocates and fixes a certain number of pages in processor storage for SDAID execution. The amount of storage required for SDAID execution depends on the combination of trace operations that you request and on the size of the output buffer (specified in the OUTDEV command).

To execute SDAID, an area is allocated that is between:

- The supervisor area.
- The start-address of the SVA(24)

By default, the allocated SDAID area is 64K bytes. During IPL, you can increase the size of the SDAID area by using the SDSIZE parameter of the SYS command.

Here are some general guidelines for determining the required size of the SDAID area:

- For simple applications (where the OUTPUT parameter is not used), SDAID requires approximately 30K bytes.
- For more complex applications, SDAID requires approximately 60K bytes.

During execution, SDAID fixes a certain number of pages in processor storage. As a result, the number of page frames available to VSE for the execution of programs in virtual mode is reduced, which may affect overall system performance.

Space Requirements for the Buffer

The internal wrap-around buffer does not belong to the SDAID area. It is situated in the SVA (31) area, which avoids the problem that storage will be exhausted if a more complex SDAID trace and a BUFFER are specified. The buffer size is requested and prefixed as multiples of 4K bytes.

If you specify TAPE=CUU, the BUFFER parameter will not be used. SDAID will then allocate an internal default buffer of 8K bytes.

Space Requirements for SDAID Execution, Summary

Basic requirement for SDAID execution:	20K
Additional requirements	
Per specified trace:	2K
If BUFFER=nn is specified:	2K + buffer size
If OUTPUT is used and OUTDEV=Tape or Buffer	12K
If OUTPUT is used and OUTDEV=Printer	20K

Number of Traces per Session

In prompt and direct input mode

The number of TRACE commands that you can submit per session depends on the types of the specified traces and the requested trace options. For each TRACE command, SDAID builds at least one TRACE command control block; for some it builds two such blocks as shown below. The program can build (and use) a maximum of ten TRACE command control blocks per session. The number of blocks per TRACE command is:

Type of trace	PHase=phasename not specified	PHase=phasename specified
PGMLOAD	2	2
VTAMIO	3	3
all others	1	2

If the traces that you requested require more than ten trace control blocks, the program ignores the TRACE command that was submitted last and informs you about this action with a message.

Via procedures

Only one procedure statement is possible, but it can create one or more TRACE commands.

Chapter 8. SDAID General Description

This chapter contains a general description of all SDAID initialization definitions, and output examples for all trace types.

The format of the commands and definition examples are shown under the descriptions for the various initialization methods in Chapters 7, 8, and 9.

Defining the Output Device

The following output destinations can be defined when you initialize an SDAID trace:

- Printer
- Tape
- Wraparound buffer
- Wraparound buffer and printer
- Wraparound buffer and tape

You define the output destination for the event records in prompt and direct input mode via the 'OUTDEV' statement. If you use a procedure to initialize a trace, the 'BUFFER=', 'TAPE=', or 'PRINTER=' specifications are used to accomplish an OUTDEV definition.

Printer Defined as Output Destination

If a line printer is defined for the output, SDAID writes the event records on the printer at the time the particular events occur.

If any program in the system writes output directly to the printer, this output will be mixed with the SDAID output. You can avoid this by

- Collecting trace output on tape and printing it afterwards via DOSVSDMP, or
- Stopping the VSE/POWER controlled printer during tracing.

Tape Defined as Output Destination

If you define a magnetic tape as output device, the SDAID program moves the trace entries into an internal buffer. SDAID writes event records to the available tape volume whenever this buffer becomes full, or a 'STOPSD' or 'ENDSD' command is processed.

Every STOPSD or ENDS command writes a tapemark on the tape if there was any trace event. However, if there was no trace event since the last STARTSD command, the tape remains unchanged.

SDAID writes the buffers successively to the tape, and one trace entry can extend over several buffers. For this reason, if you change reels after end-of-volume, the second tape may start with an incomplete trace entry.

Buffer Defined as Output Destination

A buffer in storage may be defined to store the trace event records. During tracing, SDAID stores one event record after the other. When the buffer becomes full,

SDAID wraps around and continues to write event records at the beginning of the buffer overwriting previously stored records. A buffer used in this way is called a *wraparound buffer*.

A wraparound buffer is not automatically printed when it is full. SDAID writes the buffer contents on a printer or on a tape if you request this action explicitly:

- The contents of the buffer will be written to the output device if a buffer trace is specified (TRACE BUFFER). The buffer trace causes the buffer to be written whenever it is full.
- The contents of the buffer will be written to the output device if you specify OUTPut=BUffer on a TRACE statement. The specification OUTPut=BUffer causes the buffer to be written whenever the event specified in the TRACE statement occurs.
- The contents of the buffer will be written to the output device if you use the BUFFER and BUFFOUT keyword in an SDAID procedure.

The information contained in the buffer can also be retrieved with the attention routine DUMP command (see Figure 17 and “Printing an SDAID or DUMP Command Produced Tape” on page 30).

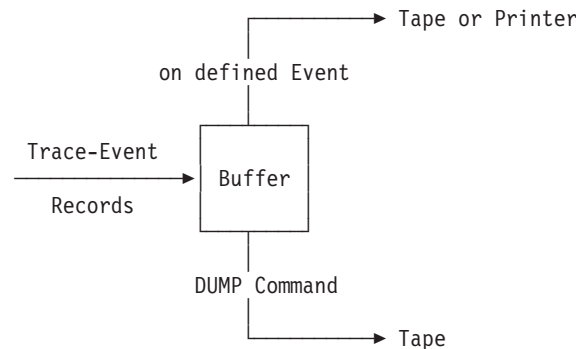


Figure 17. Overview, Tracing Events into a Buffer

You **define the size** of the buffer in number of blocks of 4K bytes.

The possible buffer size which you can define depends on the device type of the buffer output device.

If you define a tape device, the possible size may vary from 4K to 32K bytes.

You can request a buffer size from 4K to 256K bytes if a printer is defined as output device, or if no output device is defined.

Steps to Define a Wraparound Buffer

Assume, for example, that you want to:

- Collect the event records of an instruction trace in a 6K bytes buffer as long as no program check occurs;
- Write the contents of the buffer to a tape when a program check interruption occurs.

Perform the following steps:

1. Define the buffer and the tape device with the OUTDEV command (in direct input or prompt mode) or with the procedure statement in the form:
‘BUFFER=6 TAPE=280’;
2. Define the instruction trace;
3. Define the PGMCHECK trace with OUTPUT=BUFFER (in direct input or prompt mode), or define BUFFOUT=PGMC in the procedure statement.

For the formats of the output definitions, see the following sections according to the input mode used:

- “Defining the Output Device in Direct Input Mode” on page 86;
- “BUFFER=, PRINTER=, TAPE=Keyword Operands” on page 129;
- “Output Device Definition in Prompt Mode: OUTDEV Command” on page 142.

Exceptional Conditions on the Output Device

You define the SDAID output device via the OUTDEV statement. It is required that the specified printer or tape device is ready when you activate tracing via the STARTSD statement.

If an exceptional condition occurs on the output device, SDAID performs the following actions:

- Unrecoverable I/O error on tape or printer

SDAID resets the control registers to the previous state and stops trace data collection. The VSE system continues normal operation without tracing. On the next STOPSD statement SDAID presents an error message with an error code as shown in the manual *z/VSE Messages and Codes*. A final ENDSD statement releases all resources allocated to SDAID.

- End-of-volume condition on tape

SDAID writes two trailing tape marks to close the tape file, performs a ‘rewind-unload’ operation, and enters a soft wait state with the value X’00EEEEEE’ in the address part of the wait PSW. The operator should now perform the following actions:

- Mount a new tape reel
- Make the tape device ready
- Press the external interrupt key

SDAID will continue tracing onto the new tape reel.

- Intervention required on printer device (printer out of paper or stopped manually)

SDAID waits about two minutes to allow for paper refilling or other actions on the printer device. After this time has elapsed, SDAID enters a soft wait state with the value X’00EEEEEE’ in the address part of the wait PSW. The operator should now perform the following actions:

- Make the printer device ready
- Press the external interrupt key

SDAID continues tracing on the printer device.

- Intervention required on tape device

SDAID enters a soft wait state with the value X’00EEEEEE’ in the address part of the wait PSW. The operator should now perform the following actions:

- Make the tape device ready
- Press the external interrupt key

SDAID continues tracing to the tape device.

Note: If the system is in the soft wait state with X'00EEEEEE' in the address part of the PSW and the operator presses the interrupt key without making the SDAID output device ready, SDAID stops trace data collection. (SDAID reacts in the same way as for unrecoverable I/O errors).

Summary of TRACE Types

SDAID offers you various trace types so that you get the most suitable information for solving the problem in hand.

The following sections describe all SDAID trace types with their SDAID default values and shows trace event record examples. The shown output may be written to a buffer or to a tape or printer device, according to your output device specification.

Table 3 lists the commands which produce the different trace types, and summarizes the event traced in response to each command. The page references in this table help you to find the description and an output example of each trace type.

Table 3. Trace Type Summary

Trace Type	Provides a Trace of:	Page
BRANCH	Successfully executed branch instructions	56
BUFFER	The trace buffer when it is full	57
CANCEL	Program (main task) cancel or EOJ	57
EXTERNAL	External interrupts	58
EXTERNAL	External interrupts	58
GETVIS	GETVIS / FREEVIS requests	93
IO	I/O interrupts	60
LOCK	LOCK / UNLOCK requests	97
MONITORCALL	MC instructions	62
PGMCHECK	Program checks	63
PGMLOAD	Phase load requests, or actual load	64
SSCH	Start Subchannel instructions	65
STORAGE	Storage alterations	66
SVC	Executed supervisor calls	67
VTAMBU	Usage of VTAM buffers	68
VTAMIO	VTAM I/O operations	68
XPCC	Cross partition communication	108

BRANCH Trace

A branch-instruction trace provides an event record for every branch taken, if the branch target address falls into the defined area.

An example of the output is shown in Figure 18 on page 57.

BRANCH Trace Output Example

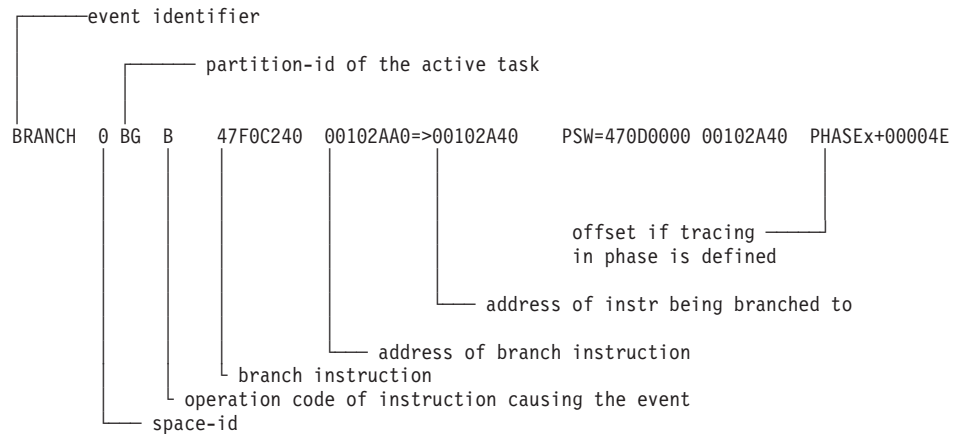


Figure 18. BRANCH Trace Event Record

BUFFER Trace

The buffer trace dumps the contents of the SDAID wraparound buffer to the output device (printer or tape) when the buffer is full.

The buffer trace can be used only if you have also 'Printer' or 'Tape' specified in the OUTDEV command or in a procedure.

The buffer trace output is the collection of all trace records contained in the buffer when a buffer overflow occurs. These trace records are written sequentially.

CANCEL Trace

This trace provides an event record when the main (or only) task of the traced partition is canceled or reaches EOJ.

You may use this trace type combined with additional output definitions to get more reasonable information at the time of a cancel or EOJ condition.

For example, use the cancel trace type to get the buffer or areas of interest together with the cancel event record recorded.

An example of the cancel trace output is shown in Figure 19.

CANCEL Trace Output Example

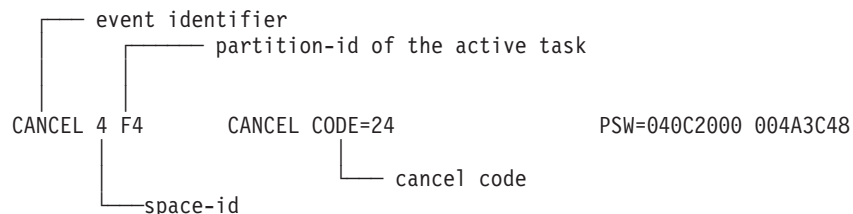


Figure 19. CANCEL Trace Event Record

EXTERNAL Trace

The external trace provides information concerning the occurrences of external interrupts such as pressing the external interrupt key. You may define one to eight of the following external interrupt types:

0040	Interrupt key	
1003	TOD-clock sync check	
1004	Clock comparator	
1005	CPU timer	
1200	Malfunction alert	
1201	Emergency signal	
1202	External call	
2401	Service signal	
2402	Logical device	* z/VM CP
2603	PFAULT handshaking	* z/VM CP
4000	IUCV, APPC	* z/VM CP
4001	VMCF	* z/VM CP

SDAID Default Value

If you do not define the type of interrupt, all external interrupts are traced.

The format of a printed external-interrupt trace event record is shown in Figure 20.

EXTERNAL Interrupt Trace Output Example

```
EXT 4 F4 INTERRUPT CODE=1004 CLOCK COMPARATOR PSW=070D0000 0041AC48
```

event identifier
Partition-id of the active task
space-id

Figure 20. EXTERNAL Interrupt Trace Event Record

GETVIS / FREEVIS Trace

A GETVIS / FREEVIS trace provides information about requests made to obtain or release virtual storage. These requests can be made using:

- SVC 3D
- SVC 3E
- An internal GETVIS call via SGETVIS and SFREEVIS macros.

The simple trace of the SVC's 3D and 3E only show the existence of SVCs at the point of invocation. However, the GETVIS / FREEVIS trace records the results of a virtual-storage request *after* it has been evaluated by the z/VSE GETVIS / FREEVIS routines.

You can limit the tracing of your GETVIS / FREEVIS requests to:

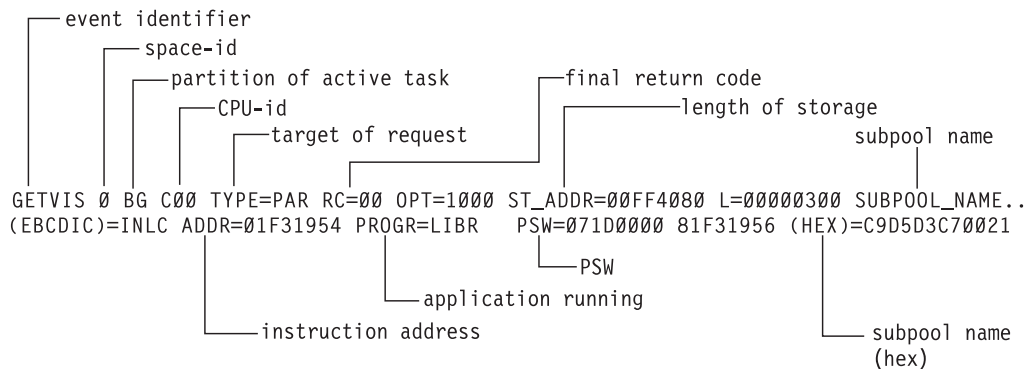
- A specific partition.
- The supervisor.
- A specific subpool name.
- A GETVIS location (24-bit or 31-bit area).

SDAID Default Value

If you do not define a partition or the supervisor, all tasks of the system are traced. All GETVIS / FREEVIS requests are traced if you do not specify any of the GETVIS trace definitions.

The format of a GETVIS / FREEVIS event record is shown in Figure 21.

GETVIS / FREEVIS Trace Output Example



```
FREEVIS 0 BG C00 TYPE=PAR RC=00 OPT=1000 ST_ADDR=00FF8000 L=00001000 SUBPOOL_NAME..
(EBCDIC)=INLC ADDR=000C997E PROGR=LIBR PSW=071D0000 800C9980 (HEX)=C9D5D3C70021
```

Figure 21. GETVIS / FREEVIS Trace Record

You can also display additional lines, as shown in Figure 22.

```
FREEVIS 0 BG C00 TYPE= RC=00 OPT=220C --ENTIRE STORAGE RELEASED--
ADDR=00609CB2 --JCL PHASE ACTIVE-- PSW=070D0000 00609CB4

FREEVIS 0 BG C00 TYPE=PAR RC=00 OPT=2002 --ENTIRE SUBPOOL RELEASED-- SUBPOOL_NAME..
(EBCDIC)=SP01 ADDR=006000A2 PROGR=GETV PSW=070D0000 006000A4 (HEX)=E2D7F0F14040
```

Figure 22. Additional Fields Displayed By GETVIS / FREEVIS Trace

INSTRUCTION Trace

An instruction trace provides information for instructions executed in the area defined for the trace. You may select certain types of instructions to be traced by defining the instruction operation codes or all types by defining an asterisk (*). Moreover, the trace can be defined to record all branch instructions, without regard to whether the branch is taken or not.

Note the difference from the branch trace, which only records the branches actually taken.

The format of a printed event record for an instruction trace is shown in Figure 23 on page 60.

INSTRUCTION Trace Output Example

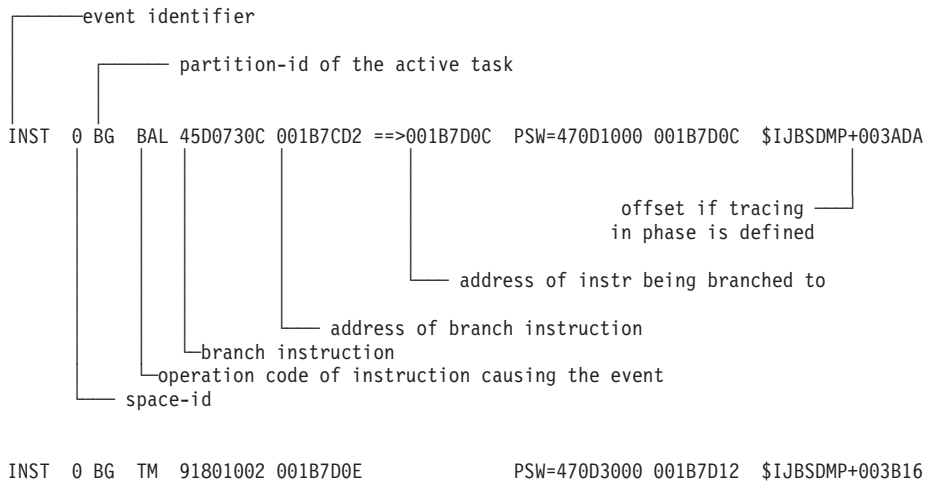


Figure 23. INSTRUCTION Execution Event Record

IO Trace (I/O Interrupt)

The IO trace collects information about I/O interrupts.

You may limit the I/O operations to be traced to a partition or to the supervisor. Another limitation is to define a particular unit, control unit, or channel to be traced.

SDAID Default Value

If you do not define a partition or the supervisor, all tasks of the system are traced. All I/O devices are traced if you do not specify any of the I/O definitions.

The format of an I/O-interrupt event record is shown in Figure 24 on page 61.

Note: Due to the fact that no tables are available in the supervisor for TP status modifier commands, the output of TP channel programs (e.g. VTAM) may be incomplete.

I/O Interrupt Trace Output Example

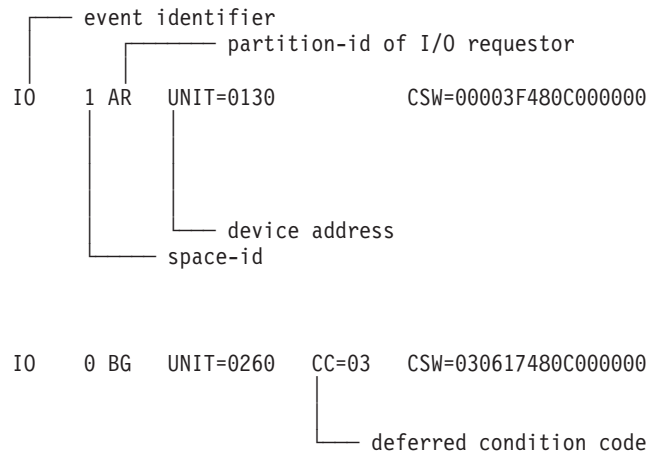


Figure 24. I/O-Interrupt Trace Event Record

LOCK / UNLOCK Trace

A LOCK / UNLOCK trace provides information about requests made to lock or unlock a resource.

You can limit the tracing of your LOCK / UNLOCK requests to:

- A specific partition.
- The supervisor.
- A specific resource name.
- A lock type.
- The scope of the lock request.
- A volume ID.
- A dedicated return code.

SDAID Default Value

If you do not define a partition or the supervisor, all tasks of the system are traced. All LOCK / UNLOCK requests are traced if you do not specify any of the LOCK trace definitions.

The format of a LOCK / UNLOCK event record is shown in Figure 25 on page 62.

LOCK / UNLOCK Trace Output Example

The three trace output lines below are listed separately, to enable explanations to be given.

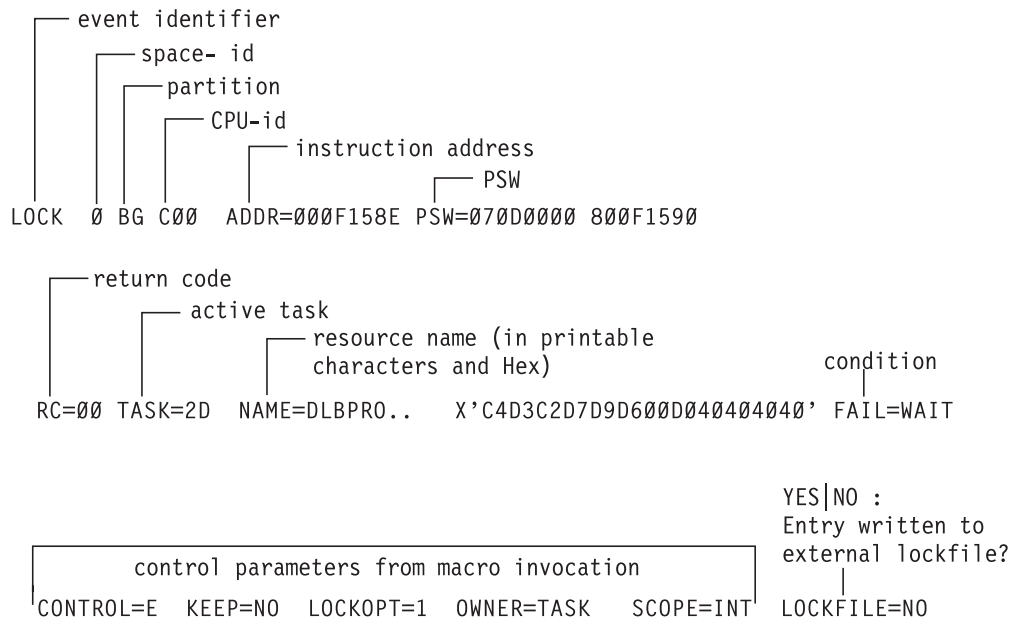


Figure 25. LOCK Trace Record

You can also display a line containing the contents of the LOCKTABLE entry, as shown in Figure 26.

```
---LTE---
000619B0 00000000 00000000 00000000 00000000 00000000 00000000 000619D0 00000000
```

Figure 26. Contents of LOCKTABLE

MONITORCALL Trace

The monitor call trace provides information about monitor call instruction executions.

You may define all (defined via an asterisk (*)) or up to eight mc (monitor classes) in hexadecimal notation of the MC instructions to be traced. An event record is provided when an executed MC instruction has a monitor class which matches any of the specified classes.

You may specify any valid monitor class; however, SDAID ignores a specification of class 2, because class 2 is used by SDAID to control tracing.

The format of an MC instruction trace event record is shown in Figure 27 on page 63.

MONITORCALL Trace Output Example

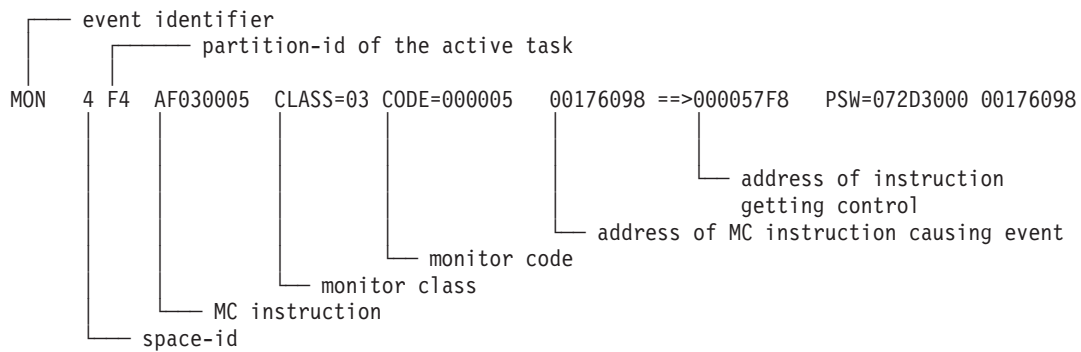


Figure 27. MONITORCALL Trace Event Record

PGMCheck Trace (Program Check)

The program check trace provides information on the occurrence of program check interrupts.

You may limit the trace operation by defining certain program interruption codes. Up to 16 program interruption codes of a value lower than X'40' in hexadecimal notation may be defined.

SDAID writes an event record only if the interrupt code returned by the system matches one of the specified interrupt codes.

If you do not want to limit the trace recording to a specific interrupt code, define an asterisk (*) to trace all program checks - except those page or segment translation exceptions which are caused by the temporary absence of a storage page. The specification `PGMC=(10 11)` traces all page or segment translation exceptions.

For a discussion of program interrupt codes, refer to the applicable *Principles of Operation* manual.

The format of a program-check event record is shown in Figure 28 on page 64.

PGMCHK Trace Output Example

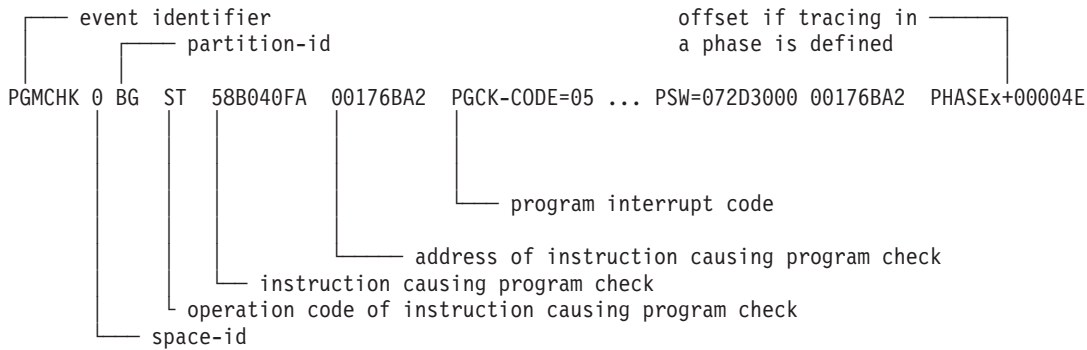


Figure 28. Program Check Trace Event Record

The program check trace also displays some additional fields, as shown in the following examples:

```
PGMCHK 0  BG  ST  58B040FA  00176BA2  PGCK-CODE=10  TEID=00581800  EAID=01  PSW=072D3000  00176BA2
          RADD=xxxxxxxx  STE=xxxxxxxx
PGMCHK 0  BG  ST  58B040FA  00176BA2  PGCK-CODE=10  TEID=00581800  EAID=01  PSW=072D3000  00176BA2
          RADD=xxxxxxxx
PGMCHK 0  BG  ST  58B040FA  00176BA2  PGCK-CODE=11  TEID=00581800  EAID=01  PSW=072D3000  00176BA2
          RADD=xxxxxxxx  PTE=xxxxxxxx
PGMCHK 0  BG  ST  58B040FA  00176BA2  PGCK-CODE=2A                EAID=01  PSW=072D3000  00176BA2
```

TEID=xxxxxxxx specifies the Translation Exception Identification. SDAID retrieves this information from low core location 144-147. **EAID=xx** shows the Exception Access Identification. SDAID retrieves this information from low core location 160. SDAID displays the fields TEID and/or EAID if the processor updates the low core locations 144-147 and/or 160 at interrupt time.

If the programming exception is a page or a segment translation exception, SDAID displays the field RADD and (if applicable) one of the fields PTE or STE. If the page or segment table portion of the virtual address points inside the page or segment table, **RADD=xxxxxxxx** shows the real address of the invalid page or segment table entry. The field **PTE=xxxxxxxx** shows the invalid page table entry. The field **STE=xxxxxxxx** shows the invalid segment table entry. If the page or segment table portion of the virtual address points outside the page or segment table, RADD shows the real address of the entry that would have been fetched if the length violation had not occurred. In this case SDAID displays no PTE or STE field.

PGMLOAD (Fetch/Load) Trace

The program load trace provides information about program load events.

Such a program load event can be one of the following:

- Phase fetch/load request
- Completion of fetch/load operation

You may limit the trace recording if you define that only the fetch/load request (REQ) or only the actual fetch/load completion (HDL) is to be traced.

You can limit the trace data collection to the load events of a certain phase.

You can limit the trace data collection to the load events occurring in an address range. This means that SDAID records

- The trace load SVCs issued within the specified address range,
- The trace load completion events if the phase is loaded into the specified address range.

SDAID Default Values

If you do not define the kind of the fetch/load request, both the request and its handling is traced. All phases are traced if you do not define a specific phase.

The format of a program load event record is shown in Figure 29.

PGMLOAD Trace Output Example

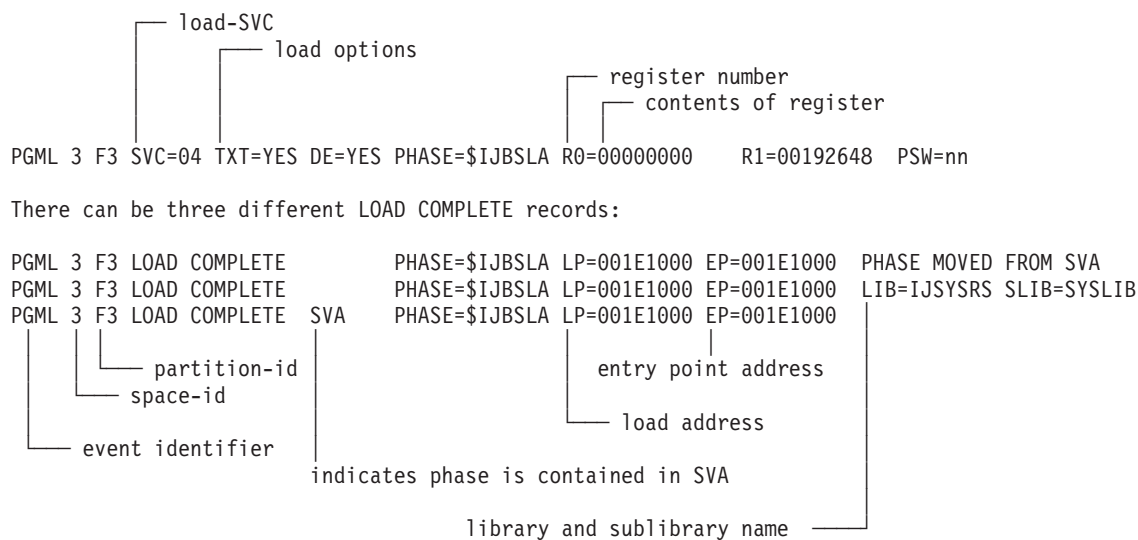


Figure 29. PGMLOAD Trace Event Records

SSCH Instruction Trace

The SSCH instruction trace provides a trace event record for every executed start subchannel (SSCH) instruction.

You may limit the trace operation by defining selected tasks or by defining a unit, a control unit, or a channel address. A partition or the supervisor may be defined as the traced tasks.

The SSCH trace may print two types of SSCH records. The SSCH-1 record shows the state before the SSCH has been performed. SDAID always displays a SSCH-1 record. The SSCH-2 record shows the condition code after the SSCH instruction. SDAID displays a SSCH-2 record if the condition code is different from zero; if the condition code is equal to zero, the SSCH-2 record is suppressed.

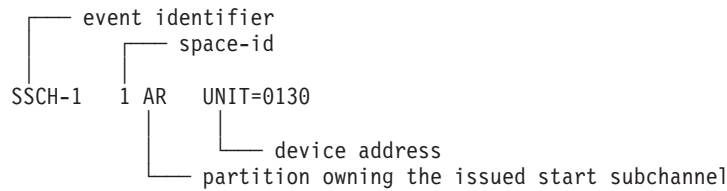
SDAID Default Value

If you do not define a partition or the supervisor, all tasks of the system are traced. All I/O devices are traced if you do not specify any of the I/O definitions.

The format of a printed standard SSCH instruction event record is shown in Figure 30.

Note: Due to the fact that no tables are available in the supervisor for telecommunications status modifier commands, the output of telecommunications channel programs may be incomplete.

Trace Event Record before SSCH Instruction



Trace Event Record after SSCH Instruction




Figure 30. SSCH (Start Subchannel) Trace Event Record

STORAGE Alteration Trace

The storage alteration trace writes an event record whenever a program alters the contents of a defined storage area. Storage alterations caused by I/O operations are not recorded.

With the AREA or JOBNAME specification you define the tasks you want to watch (the source space or partition). AREA=ALL causes all tasks of the system to be watched. With the STAREA, STJNAM, and STDSPN operands you define the altered storage area (the target area). If the source space (or source partition) is the same as the target space (or target partition), you may leave out the parameters STAREA, STJNAM and STJNUM. Via the ADDRESS definition you specify the storage range where the alteration occurs.

You may limit the trace operation by defining a certain storage *pattern*. If you define such a pattern, a trace event record is written only when a storage area is set to the specified value. Specify the pattern in hexadecimal notation. The pattern can be up to four bytes long.

If you specify an odd number of digits, a zero is inserted to the left of the first specified hexadecimal digit.

SDAID Default Value

You may omit the pattern definition. This causes each alteration of the defined storage interval to be traced.

An example of a storage alteration trace event record is shown in Figure 31 on page 67.

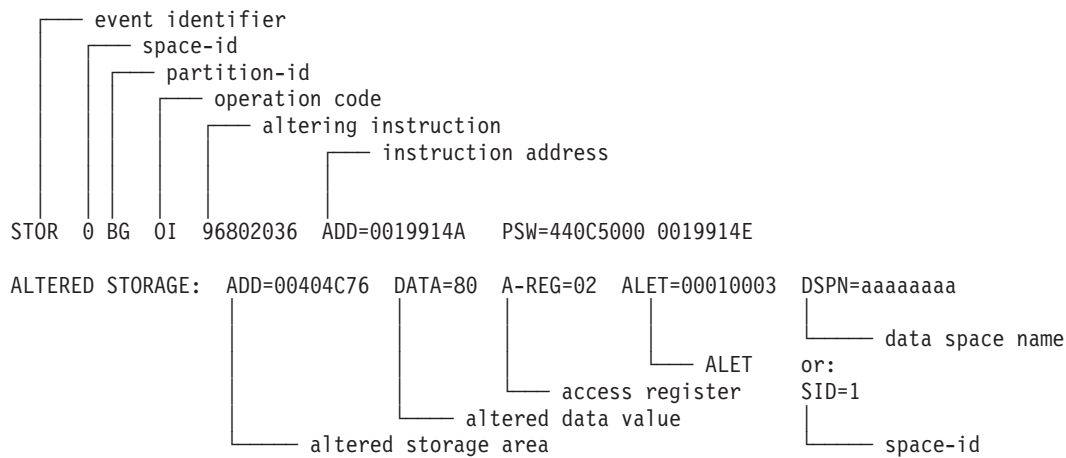


Figure 31. Storage-Alter Trace Event Record

The first line of the storage-alteration trace record displays information about the altering program. The record shows the space-id, the partition-id, the altering instruction and its storage address, and the PSW. The second line gives information about the altered storage area. The record shows the address of the altered storage area (ADD=) and the altered data value (DATA=).

If the traced instruction has accessed data within a different address space, the record shows the access register number (A-REG=), the access register contents (ALET=), and the space identifier (SID=).

If the traced instruction has accessed a storage location within a data space, the record shows the access register number (A-REG=), the access register contents (ALET=), and the name of the target data space name (DSPN=).

SVC Trace (Supervisor Call)

The supervisor call trace provides information about one, several, or all SVC instructions executed in the defined area.

You have to define at least one or may define up to 16 different SVC codes or an asterisk (*) specifying that all supervisor call instructions are to be traced. Specify the SVC code in hexadecimal notation.

A typical SVC trace event record is shown in Figure 32.

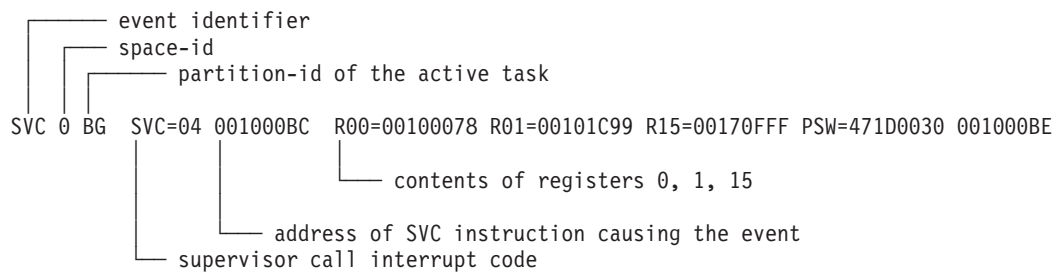


Figure 32. SVC Trace Event Record

VTAMBU Trace (VTAM Buffer)

The VTAM buffer trace provides an event record each time VTAM uses one of the buffers in its buffer pool.

The format of a printed standard VTAMBU event record is shown in Figure 33.

Note: For the VTAMBU trace you must condition VTAM either by entering a VTAM MODIFY command or by specifying the TRACE operation when you start VTAM. At the end of an SDAID session, end VTAM conditioning with the appropriate MODIFY command:

Start: F NET,TRACE,TYPE=SMS,ID=VTAMBUF

Stop: F NET,NOTRACE,TYPE=SMS,ID=VTAMBUF

For information on conditioning VTAM, see the *VTAM Diagnosis* publication.

```

  (A)  ┌───┐ partition-id
VT-I0  3 F3 UNIT=0130      CSW =000039A80C000000
        └───┘ space-id
  (B)  ┌───┐
VT-SVC 3 F3 SVC=31 0010012E  R00=00100A1C R01=00101F0F R15=00180FFF
  (C)  ┌───┐
SSCH-1 3 F3  UNIT=0130
        └───┘ sequence number of trace record
VTAM BUFFER POOL USE SEQ.NO = 00000001 DAY 097 TIME 19:47:10
VF IN USE=00005 MAX ALLOC=00005 MAX WAIT=0000 EXPAND=0000 AVAIL=0000 CUR AVAIL=0000
VP IN USE=00089 MAX ALLOC=00098 MAX WAIT=0000 EXPAND=0000 MAX AVAIL=0000 CUR AVAIL=0000
  (D)  (E)  (F)  (G)  (H)  (I)
└───┘ pool-ID
```

- (A) I/O Interrupt Trace (for description see Figure 24 on page 61).
- (B) SVC Trace (for description see Figure 32 on page 67).
- (C) SSCH Trace (for description see Figure 30 on page 66).
- (D) Number of buffers (pages for VF and VP) in use when buffer usage was recorded.
- (E) Maximum number of buffers (pages for VF and VP) at any point in time up to the point buffer usage was recorded.
- (F) Maximum number of requests for buffers (pages for VF and VP) that were queued at any point in time up to the point buffer usage was recorded.
- (G) Number of times the buffer (page) pool was expanded up to the point buffer usage was recorded.
- (H) Number of buffers (pages for VF and VP) that were in the pool when buffer usage was recorded.
- (I) Maximum number of buffers (pages for VF and VP) that were in the pool at any point in time up to the point buffer usage was recorded.

Figure 33. VTAMIO/VTAMBU Trace Record

VTAMIO Trace

The VTAMIO trace combines the following trace types:

- SVCs with codes X'31' and X'35'
- SSCH instructions
- I/O interrupts

Please find a description of the traces involved under the various trace type descriptions (SVC trace, SSCH trace and I/O trace).

SDAID Default Value

All I/O devices are traced if you do not specify any of the I/O definitions.

The format of the printed VTAMIO event records is shown in Figure 33 on page 68.

XPCC Trace

An XPCC trace provides information about connections between different applications (cross-partition communication). The information is gathered after the requested function has been processed and completed by the VSE cross-partition communication routine.

You can limit the tracing of your XPCC requests by using one or more of the XPCC trace definitions.

SDAID Default Value

All XPCC requests will be traced if you either:

- Do not specify one or more of the optional trace definitions.
- Code an asterisk (*) as parameter for the mandatory trace definitions.

The format of an XPCC event record is shown in Figure 34.

XPCC Trace Output Example

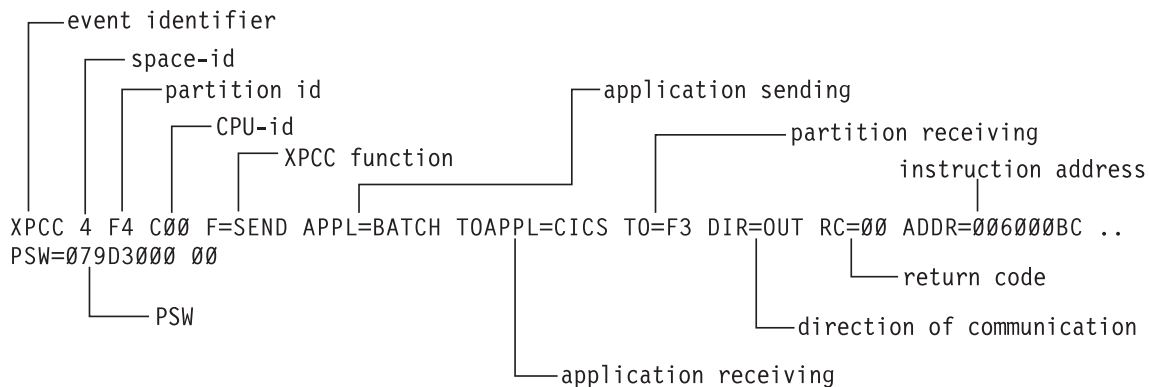


Figure 34. XPCC Trace Record

You can also display additional lines, as shown in Figure 35.


```
OUTPUT=XPCCB      XPCC control block
OUTPUT=XPDATABU  Data Transmit Buffer
```

Figure 35. Additional Fields Displayed By XPCC Trace

Notational Conventions

For a description of the syntax diagrams in this chapter please read "Understanding Syntax Diagrams" on page xv.

Defining the Area to be Traced: AREA Definition



You define the area for which processing is traced by using the AREA (or the JOBNAME) definition. You can specify either AREA or JOBNAME in the TRACE statement, but not both. For AREA, specify one of the following definitions:

- `partition_id` (static or dynamic)
- SUP
- ALL

Only one of these definitions is possible.

partition_id

Specifies the partition in which tasks are to be observed, like BG, F3, Y1.


If, for example, BG is specified, the steps executed by the BG main task and by all attached subtasks are traced. ARea=BG does not necessarily mean that tracing is restricted to the BG space. For a dynamic partition (Y1, for example), the JOBNAME definition is generally used (see below).

SUP Specifies that the activities of the supervisor are to be traced.

ALL Specifies that the activities of all tasks in the VSE system are to be traced.

Note: For the VTAMIO trace, only ARea=partition_id can be specified, where 'partition_id' denotes the partition where VTAM is active. In that case only the I/O activities of that partition are traced.

Defining the Job to be Traced: JOBNAME Definition



If you want to trace a job in a dynamic job class, the AREA specification is not sufficient, since you do not know in which partition the job may execute. Therefore, the keywords JOBNAME and JOBNUM have been introduced which allow to trace a VSE/POWER job in a dynamic or static partition. If VSE/POWER is not installed in the system, JOBNAME is not applicable. You can specify either AREA or JOBNAME in the TRACE statement, but not both.

jobname

Specifies the name of the VSE/POWER job to be traced. Tracing starts when VSE/POWER selects the specified job for execution. The specified job name must correspond to the JNM operand of the VSE/POWER JECL statement.

jobnumber

Specifies the VSE/POWER-defined job number. It may be used if more than one job with the specified job name is contained in the VSE/POWER reader queue.

Notes:

1. The JOBNAME operand cannot be used to trace a job in a VSE/POWER writer-only partition.
2. JOBNAME should not be used in a VSE/POWER-controlled partition started with the 'MT' option of the PSTART command.
3. SDAID does not accept POWER job names containing the character '-'. (The SDAID command language uses this character as continuation character.)

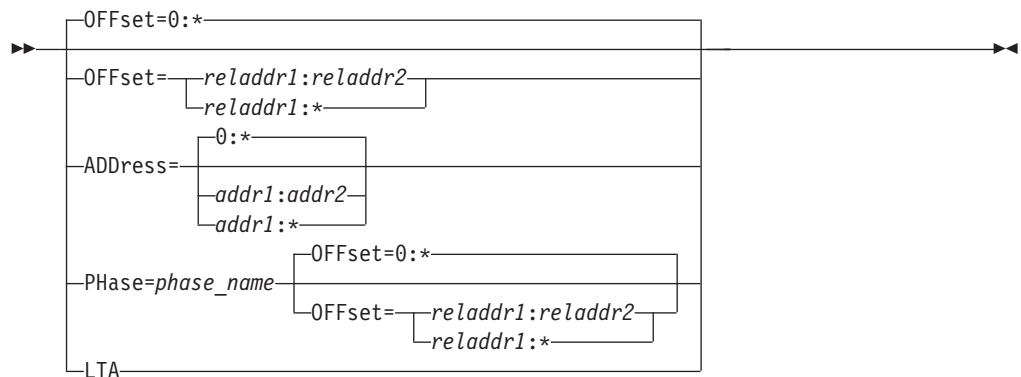
Defining the Storage to be Traced: OFFset, ADDRESS, PHase, LTA

Whereas the ARea and JOBNAME definitions allows you to specify the tasks that you want to observe, the storage definition determines the storage region that is to be investigated. The storage definition has to be specified in accordance to the task definition. For example, if you specified a partition-id for ARea, the definitions OFFset, ADDRESS, PHase, LTA can be used.

The following sections tell you the valid storage definitions in accordance with the various task specifications.

Storage Definition for AREA and JOBNAME

You may use one of the following storage area definitions for the AREA and JOBNAME definitions:



OFFset=

limits tracing to a certain address range via offsets relative to the partition start address. OFFset=0:* is the default specification when ARea is defined by a partition-id. The OFFSET specification is not valid for AREA=ALL.

reladdr1:reladdr2

defines a trace address range in hexadecimal notation.

'reladdr1' can be any relative address within the partition, supervisor or phase defined by the ARea operand.

'reladdr2' must be higher than or equal to reladdr1.

reladdr1:*

defines a trace address range starting with 'reladdr1' up to the end of the partition, the supervisor or the phase.

0:*

defines that the whole storage allocated to the partition, the supervisor or the phase is defined as trace storage area.

ADDRESS=

limits tracing to a certain address range within the storage allocated to VSE.

addr1:addr2

defines a trace address range in hexadecimal notation. 'addr1' and 'addr2' can be any address in the virtual storage defined for VSE.

'addr2' must be higher than or equal to 'addr1'.

addr1:*

defines a trace address range starting with 'addr1' up to the end of virtual storage.

0:*

defines that the whole storage allocated to VSE is defined as trace storage area.

PHase=

Limits tracing to a certain phase. If the named phase resides in the Shared Virtual Area (SVA), SDAID uses the start and end address of the named phase as tracing boundaries. If the named phase does not reside in the SVA, SDAID defers tracing until the named phase is loaded into a user partition. When a load request for the named phase is issued, SDAID activates the deferred trace and updates the trace start and end address with the start and end address of the phase just loaded. This means that only those phases can be traced via the PHase parameter which (1) reside in the SVA, or (2) which are loaded into the traced partition after the STARTSD statement has been issued.

If you specify the PHase= operand, the operand OFFset= can be used in addition. In this case, OFFset= defines a trace area within the phase. If you omit the OFFset= definition, the whole storage area of the phase is traced.

reladdr1:reladdr2

defines a trace address range in hexadecimal notation.

'reladdr1' can be any relative address within the phase.

'reladdr2' must be higher than or equal to reladdr1.

reladdr1:*

defines a trace address range starting with 'reladdr1' up to the end of the phase.

0:*

defines that the whole storage allocated to the phase is defined as trace storage area.

LTA

defines that the logical transient area is specified as tracing range.

SDAID Default Values

If you do not specify OFFset, ADDRESS, PHase or LTA, the following defaults will be set:

- If AREA=*partition* | *supervisor*, the complete storage area that is allocated to the partition or supervisor will be traced. OFFset=0:* is the definition that results in direct input-mode notation.
- If AREA=ALL, or if the parameter is omitted, the complete storage area that is allocated to z/VSE will be traced.

Note: If you wish to trace in a particular partition (parameter Area=*partition*) and if parts of your program belong to the SVA (24-bit or 31-bit), to include this event you must explicitly increase the range of the trace using the ADDR=0:* parameter.

Defining Additional Trace Output: OUTPUT Definition



You may specify additional trace output with the OUTPUT definition. This additional trace information is recorded together with the trace event records. For example, you may define that a dump of defined control blocks or address ranges be recorded in addition to the trace event record.

You may select one or more definitions for a specific trace type.

For a summary of all OUTPUT definitions, see Table 4. The page numbers in the third column refer you to the description of the OUTPUT definitions in this chapter.

Table 4. OUTPUT Definition Summary

Definition	Records/prints in addition:	Page
BUffer	Contents of SDAID output buffer	74
CCB	CCB or IORB (TRACE=IO, SSCH, or VTAMIO only)	74
CCW	CCWs, IRB (TRACE=IO, SSCH, or VTAMIO only)	74
CCWD=nnnn	CCWs plus nnnn bytes of data, IRB (TRACE=IO, SSCH, or VTAMIO only)	74
COMReg	Partition communication region	75
CReg	Control registers	75
DUMP	Virtual storage	75
FReg	Floating point registers	77
GReg	General purpose and access registers	78
IOTab	PUB, LUB, ERBLOC, ERRQ, CHANQ	78
LOCKTE	Lockable entry (LOCK trace only)	61
LOWcore	Processor storage from zero to X'2FF'.	79
LTA	Logical transient area	79
PTA	Physical transient area	79
PTAB	Partition related control blocks: PCB, PIB, PIB2	79
SUP	Supervisor plus GREGs and CREGs	80
SYSCom	System communication region	80
TOD	Time-of-Day clock	80
TTAB	Task related control blocks: TIB, TCB, PCB, PIB, PIB2	80
XPCCB	XPCC control block (XPCC trace only)	69
XPDATABU	Buffer for data to be transmitted (XPCC trace only)	69

Writing the Trace Buffer

▶▶—OUTPut=BUffer—▶▶

BUffer

writes the contents of the SDAID buffer to the device specified with the OUTDEV statement. The buffer is written immediately after the associated event has been recorded in the buffer.

Note: Specifying BUffer as one of a number of output options may result in the original event record(s) getting lost due to the wraparound recording technique used.

Recording the CCB or IORB

▶▶—OUTPut=CCB—▶▶

CCB records or prints the contents of either the CCB or the IORB (input/output request block) plus the TOD (time-of-day clock). This output option is meaningful only with an IO or SSCH trace request.

Recording the CCW

▶▶—OUTPut=—
└─CCW─┐
└─(CCWD=nnnn)—┘▶▶

CCW records or prints the available channel program (Channel Command Word chain) when the trace type is SSCH. In case of an IO trace, only the CCWs which refer to transferred data are recorded or printed.

The output contains also the first 16 bytes of the associated data, the CCB, the IRB (interruption request block) and the TOD (time of day clock).

Specifying this output option for an event other than IO or SSCH is not meaningful.

CCWD=nnnnn

(CCW plus data) records or prints up to a maximum of nnnn bytes of the transferred data, the CCB, the IRB (interruption request block) and the TOD clock in addition to the information processed with the CCW specification. The number nnnn may be any (decimal) number between 1 and 65535.

The most meaningful trace type to be combined with this output option is the IO trace.

For an example of the output produced with this option, see Figure 36 on page 75.

```

I/O   S BG  UNIT=0110          CSW =0000D610 0C000000
TOD = 95.016 17.53.54.333
CCB= 0000D2E0 00000004 0000880B 0000D5E8 0000D610

```

```

CCW= 0000D5E8 0700D47A 40000006 DATA= 0000000B 000A          *.....      *
CCW= 0000D5F0 2300D481 40000001 DATA= 9C                    *.          *
CCW= 0000D5F8 3100D47C 60000005 DATA= 000B000A 17          *.....      *
CCW= 0000D600 0800D5F8 00000000
CCW= 0000D608 8600F400 20000400

```

```

---CCW DATA---
0000F400 D7C8C1E2 C5404040 00805BD1 D6C2C3E3 D3C90020 00000000 14480048 00000000 *PHASE ..$JOBCTLI.....*
0000F420 171C0480 00060000 00010000 16E80001 F9F4F1F1 F0F90213 717CF9F4 F1F1F1F2 *.....Y..941109...@941112*

```

...

```

----IRB-----
0001B940 00004007 0000D610 0C000000 00400000 00000000 00000000 00000000 00000000 *.. ...0.....*
0001B960 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 *.....*

```

Note: For a CCW with a data length less than 17 bytes data is displayed in the CCW line, else in a separate block.

Figure 36. Output of `OUTPut=(CCWD=256)`

Recording the Partition Communication Region

▶▶ `OUTPut=COMReg` ◀◀

COMReg

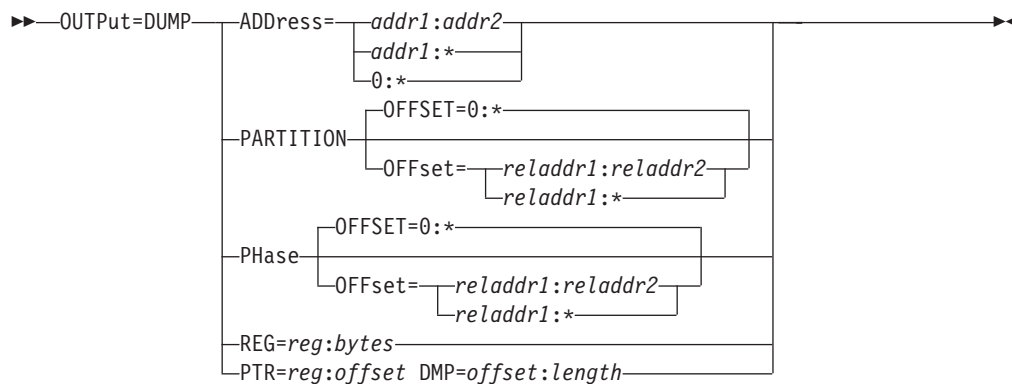
records or prints the contents of the partition communication region.

Recording the Control Registers

▶▶ `OUTPut=CReg` ◀◀

`CReg` records or prints the contents of all control registers.

Dumping Virtual Storage



Note: SDAID dumps only those areas which are in real storage, that is, accessible without page fault.

DUMP

records or prints the contents of virtual storage.

You may request up to ten different dumps.

You have to specify one or more of the dump area specifications as shown below.

PARTITION

allows to dump the partition for which the trace is active. This operand is valid for the static as well as for the dynamic partitions (AREA or JOBNAME). The OFFset operand may be defined in addition to PARTITION to limit the dump area in the partition:

reladdr1:reladdr2

defines a trace address range in hexadecimal notation.

'reladdr1' can be any relative address within the partition or the phase defined by the ARea operand.

'reladdr2' must be higher than or equal to 'reladdr1'.

reladdr1:*

defines a trace address range starting with 'reladdr1' up to the end of the partition or the phase.

0:* defines that the whole storage allocated to the partition or the phase is defined as trace storage area.

PHase limits the dump to the phase that is specified in the applicable area definition. The OFFset operand may be defined with the PHase operand to limit the dump area in the phase. (See above, under "PARTITION".)

ADDress

defines a dump range by a pair of addresses. For example, if you want to dump the contents of four bytes starting at storage location 0080 (hexadecimal). The definition in direct input mode would look like this:

```
ADD=0080:0083
```

addr1:addr2

defines a trace address range in hexadecimal notation. 'addr1' and 'addr2' can be any address in the virtual storage defined for VSE.

'addr2' must be higher than or equal to addr1.

addr1:*

defines a trace address range starting with 'addr1' up to the end of virtual storage.

0:* defines that the whole storage allocated to VSE is defined as trace storage area.

REG=reg:bytes

if the starting address of the dump is specified by a register, and the number of bytes to be dumped is specified by a hex value.

The maximum number of bytes that can be specified is 1000.

PTR=reg:offset DMP=offset:length

if the dump area is located via a register plus an offset which addresses a pointer. The dump area itself is determined by the definition

```
DMP=offset:length
```

which defines the dump start address relative to that pointer and the dump length.

The offsets and the dump length are specified in hexadecimal notation. The maximum value for offset and length is 1000. The following example and Figure 37 explain such a dump area definition.

For example:

The contents of register E is used to locate a control block. The fullword at relative offset X'10' in this control block is used as a pointer to another control block or data area.

Starting at offset X'200' an area of X'100' bytes is dumped.

Direct Input Mode Format

OUTPut=(DUMP PTR=E:10 DMP=200:100)

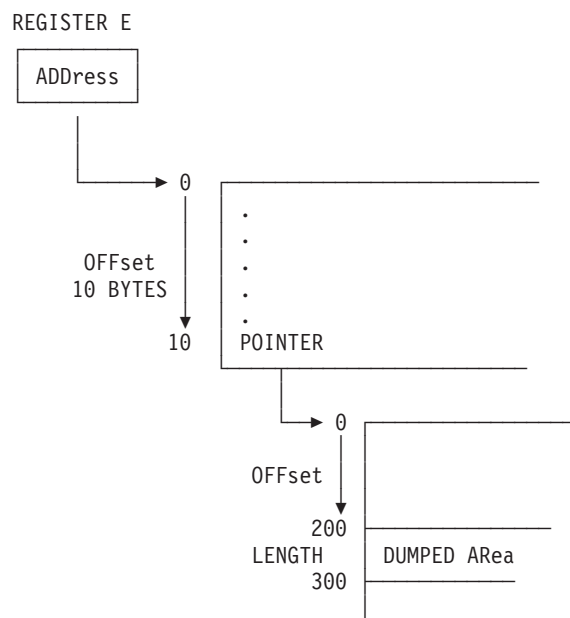


Figure 37. Overview: Defining the Area to be Dumped

In the operands REG and PTR, you can specify a general register for addressing the dump area. The PSW of the traced program determines whether the contents of the general register is interpreted as a 24-bit address or as a 31-bit address. If access registers are active, SDAID uses the general register X and the access register X to address the dumped interval.

Recording Floating-Point Registers

►►—OUTPut=FReg—◄◄

FReg records or prints the contents of all floating-point registers. Figure 38 on page 78 shows an example of a printout of the contents of these registers.

Trace Output Example with OUTPUT=FReg

```

      |-----instruction that triggered the output
INST  0 BG  L  5870817C  001010B2  PSW=470D5000 001010B6
FP-REGS 00000000 00000000 00000000 00000000 00000000 00000000 64E2D100 010300FA
```

Figure 38. Printout of Floating-Point Registers

Recording General-Purpose and Access Registers

▶▶—OUTPUT=GReg—◀◀

GReg records or prints the contents of all general-purpose and access registers (if available).

Figure 39 shows an example of a printout of the contents of the general-purpose registers, Figure 40 of general-purpose and access registers.

Trace Output Example of General-Purpose Registers

```

      |-----instruction that triggered output
INST  0 BG  BALR  0577  001010B6  PSW=470D1000 0010A6D4
GR  0-7 80400C00 02021B40 80000000 FFFFFFFF 00000000 00000000 00000000 00000000
      8-F 00000000 0000000F 00100000 001A76FF 00000000 00000000 8F000000 00000200
```

Figure 39. Printout of General-Purpose Registers

Trace Output Example with Access Registers

```

      |-----instruction that triggered output
INST  0 BG  BALR  0577  001010B6  PSW=470D5000 0010A6D4
GR  0-7 80400C00 02021B40 80000000 FFFFFFFF 00000000 00000000 00000000 00000000
      8-F 00000000 0000000F 00100000 001A76FF 00000000 00000000 8F000000 00000200
AR  0-7 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00020004
      8-F 00020004 00000000 00000000 00000000 00000000 00000000 00000000 00000000
```

Figure 40. Printout of General-Purpose and Access Registers

Recording PUB, LUB, ERBLOC, CHANQ

▶▶—OUTPUT=IOTab—◀◀

IOTab records the contents of the following I/O tables: PUB, LUB, ERBLOC, CHANQ.

Note: If the attention routine is traced, no LUB table is recorded (only partitions have LUB tables).

Dumping Processor Storage from X'00'to X'2FF'

▶▶—OUTPut=LOWcore—◀◀

LOWcore

(Lowcore contents) records or prints the contents of the first 768 bytes of processor storage (X'00' to X'2FF').

See Figure 41 for a sample output.

Trace Output Example with OUTPut=LOWcore

address of COMREG for active partition

```
PGMCHK  0 BG  ????  0000  001012B8  PGCK-CODE=01  PSW=471D0000 001012BA
—LOWCORE—
000000 400FA248 00000000 00001000 000FB1FA 00000000 00000330 ...
000020 47100000 00100228 471D0000 001012BA 00000000 00000000 ...
000040 00000000 00000000 000488B0 00000000 FEE67800 00000000 ...
...
```

Figure 41. Printout of Low Address Storage

Recording the Locktable Area

▶▶—OUTPut=LOCKTE—◀◀

LOCKTE

records or prints the Locktable entry of the current LOCK or UNLOCK event.

Recording the Logical Transient Area

▶▶—OUTPut=LTA—◀◀

LTA records the contents of the LTA on occurrence of the associated event.

Recording the Physical Transient Area

▶▶—OUTPut=PTA—◀◀

PTA records or prints the contents of the physical transient area on occurrence of the associated event.

Recording Partition-Related Control Blocks

▶▶—OUTPut=PTAB—◀◀

PTAB records or prints the contents of the Partition Control Block (PCB) and the Program Information Block (PIB and PIB2) of the active partition.

Recording the Supervisor Area

▶▶—OUTPut=SUP—◀◀

SUP records or prints the contents of the storage area used by the supervisor.

Recording the System Communication Region

▶▶—OUTPut=SYSCom—◀◀

SYSCom
records or prints the contents of the system communication region.

Recording the Time-of-Day Clock

▶▶—OUTPut=TOD—◀◀

TOD records or prints the setting of the time-of-day clock each time the associated event occurs.

Figure 42 shows an example of this option.

Trace Output Example with OUTPut=TOD

```
PGMCHK 0 BG ST 58B04057 001012BA PGCK-CODE=05 PSW=471D0000 001012BE
TOD = 90.101 13.54.02.763
```

instruction that triggered output

Figure 42. Program-Check Event with Time of Day

Recording Task-Related Control Blocks

▶▶—OUTPut=TTAB—◀◀

TTAB records or prints the contents of the Task Information Blocks (TIBs), the Task Control Blocks (TCBs) of the tasks belonging to the active partition, and the partition-related control blocks PCB, PIB, PIB2.

Recording the XPCC Communication Control Block

▶▶—OUTPut=XPCCB—◀◀

XPCCB
records or prints the contents of the XPCC communication control block.


Recording the XPCC Data Buffer

▶▶—OUTPut=XPDATABU—◀◀

XPDatabu

records or prints the contents of the XPCC data buffer which is used to transfer data.

Defining the Trace Options: OPTion Definition



SUPervisor

traces a code segment within the supervisor. It allows to trace supervisor routines while they are working for a user partition (by activating the PER bit in the PSW of the supervisor routines). The specification OPTion=SUPervisor is applicable only for the branch trace, the instruction trace, and the storage alteration trace. You may use OPTion=SUPervisor if you specify AREA=partition-id or JOBNAME=jobname. Examples:

```
TRACE INST=* AREA=F3 ADD=5000:5200 OPTION=SUP
```

This statement traces the instructions in the storage interval between 5000 and 5200 if this supervisor code is executed for the F3 partition (with the PIK of the F3 partition).

```
TRACE BRANCH JOBNAME=TESTPROG ADD=0:* OUTPUT=GREG OPTION=SUP
```

This statement traces the effective branch instructions of the job named TESTPROG. ADD=0:* means that all storage within and outside the private user partition is to be traced. OPTION=SUPervisor causes supervisor routines to be traced while they are servicing the traced partition.

Halt

stops processing of the system when the defined trace event occurs. SDAID puts the system into a wait state with address X'00EEEE' in the address portion of the wait PSW. This wait on event enables you to perform some debugging work, for example to display registers or selected storage areas.

How to Get out of this WAIT

- Give an external interrupt.

The trace remains initialized and the system stops at the next trace event occurrence.

- Alter storage location 0 to a value of X'FF', then give an external interrupt.

The trace remains active but the Halt option is canceled.

NOJCL

suppresses tracing of Job Control Phases. If the keyword NOJCL is omitted, the user program and the job control statements are also traced.

NOSource

Do not record if a space switch is in effect, and the current event is within the SOURCE space.

NOTarget

Do not record if a space switch is in effect, and the current event is within the TARGET space.


Terminate

allows you to terminate SDAID output at the occurrence of the specified event. You may start the trace output again if you issue the STOPSD and the STARTSD attention routine commands. Note, that you should issue the attention routine command ENDS if you want to end the trace operation and release all system resources which SDAID has used.

OCcurrence=occ1:occ2

specifies the number of associated events to be traced. For example, the specification 1:20 defines that tracing starts with the first occurrence of the specified trace event and ends with occurrence 20.

Defining the Traced I/O Devices

**Notes:**

- 1 You can specify up to eight device addresses.
- 2 You can specify up to 16 control unit addresses.
- 3 You can specify up to 16 channel addresses.
- 4 You can specify up to eight device addresses.
- 5 You can specify up to 16 control unit addresses.
- 6 You can specify up to 16 channel addresses.

UNit=cuu

specifies one or up to 8 device addresses.

CU=cu

specifies one or up to 16 control unit addresses.

CHannel=c

specifies one channel address or a list of up to 16 addresses.

Notes:

1. You may define these operands with the TRACE statements for the IO trace, the SSCH trace, and the VTAMIO trace.
2. The parameters UNit, CHannel, and CU are optional. If you do not specify the UNit statement, all devices are traced.
3. The parameters UNit, CHannel, and CU are mutually exclusive (in the same TRACE command).

Chapter 9. Initialize an SDAID Trace in Direct Input Mode

This chapter describes how you can initialize a SDAID trace via statements read in from a SYSIN device or the attention routine (AR).

Direct input means that you enter the trace specifications directly via control statements. Beside the trace initialization in prompt mode or via procedures the direct input mode can be used to set up SDAID traces. You can use the direct input mode of SDAID to enter complete SDAID commands avoiding the time-consuming prompt input mode.

This chapter consists of these main sections:

- “Initializing an SDAID Trace in Direct Input Mode”
- “Starting the SDAID Trace Initialization” on page 85
- “Ending the SDAID Trace Initialization” on page 86
- “Defining the Output Device in Direct Input Mode” on page 86
- “The TRACE Statement” on page 88
- “BRanch Trace” on page 89
- “BUffer Trace” on page 90
- “CAnceL Trace” on page 90
- “EXTernal Trace” on page 91
- “GETVIS Trace” on page 93
- “INSTRUction Trace” on page 95
- “I/O Interrupt Trace” on page 96
- “LOCK Trace” on page 97
- “MONitor Call Trace” on page 99
- “PGMCheck Trace” on page 100
- “Program Load Trace (Fetch/Load Trace)” on page 101
- “SSCH Instruction Trace” on page 103
- “Storage Alteration Trace” on page 103
- “Supervisor Call Trace” on page 105
- “VTAM BUffer Trace” on page 107
- “VTAMIO Trace” on page 107
- “XPCC Trace” on page 108
- “Additional Definitions” on page 110

Initializing an SDAID Trace in Direct Input Mode

SDAID trace initialization in direct input mode uses the following statements:

Table 5. Input Statement Summary

SDAID Statement	Page
EXEC SDAID or SDAID	85
OUTDEV specification	86
TRACE type and additions	88

Table 5. Input Statement Summary (continued)

SDAID Statement	Page
/*, READY or ENTER	86

Selecting the SDAID Input Mode

In direct mode, SDAID commands can be entered either via the attention routine or via job control.

Commands Entered Via the Attention Routine:

You may initialize traces via the attention routine either in direct input mode or in prompt mode. The SDAID program may switch from one input mode to the other according to your command input. However, the statements described in this section are used for direct input mode only. The **SDAID program is started** by the attention routine command

```
SDAID
```

You **determine the direct input mode** by defining at least one keyword operand together with either the TRACE or the OUTDEV command, for example:

```
OUTDEV P=E
```

The SDAID program **switches to prompt input mode** when you omit all possible keyword operands, or if you use a prompt-input-mode statement. An example of prompt input is a statement with a question mark (?), for example:

```
TRACE INST=* AR=BG OUTP=?
```

Commands Entered Via Job Control

You may **invoke the SDAID program** in a partition using the job control statement:

```
// EXEC SDAID
```

This statement and the SDAID trace initialization commands may be entered from a console or from a SYSIN device.

You **specify the output device** of the trace with the OUTDEV statement. Only one output device specification is possible at one time in the system. Any subsequent OUTDEV statement overrides the existing one.

The TRACE statement contains the **definition of the trace type** and additional trace keyword operands. You can enter up to ten TRACE statements.

You **end the trace initialization** with the READY statement. If you entered the statements via SYSIN, the end of data (/*), or ENTER in case of console input are treated as READY statements and end the initialization process. When the READY statement has been read in, no further OUTDEV or TRACE statement can be entered.

You can **cancel the SDAID setup** during initialization with the ENDS or CANCEL statement. Two initialization examples are shown in Figure 43 on page 85. One entered directly via the attention routine, the other read in via a SYSIN device.

In both examples the same SDAID trace is initialized.

Direct Input Setup in the Attention Routine

```
→ sdaid
  AR 015 4C05I PROCESSING OF 'SDAID' COMMAND SUCCESSFUL.
  AR 015 1I40I READY
→ outdev t=280
  AR 015 4C05I PROCESSING OF 'OUTDEV' COMMAND SUCCESSFUL.
  AR 015 1I40I READY
→ trace branch ar=bg off=0:200 outp=(greg dump add=40000:40100)
  AR 015 4C05I PROCESSING OF 'TRACE' COMMAND SUCCESSFUL.
  AR 015 1I40I READY
→ trace inst=(d207 95 9103) ar=bg add=0:* outp=greg
  AR 015 4C05I PROCESSING OF 'TRACE' COMMAND SUCCESSFUL.
  AR 015 1I40I READY
→ ready
  AR 015 4C05I PROCESSING OF 'READY' COMMAND SUCCESSFUL.
  AR 015 1I40I READY
```

Direct Input Setup under Control of a Partition

```
// EXEC SDAID
OUTDEV T=280
TRACE BRANCH AR=BG OFF=0:200 -
      OUTP=(GREG DUMP ADD=40000:40100)
TRACE INST=(D207 95 9103) AR=BG ADD=0:* OUTP=GREG
/*
```

Figure 43. Trace Initialization Examples (Direct Input Mode)

Notational Conventions

- The various operands may be separated by at least one blank or by a comma.
- Enter all operands in the specified order. Examples are shown in chapter Command Input Paths starting on page Command Input Paths on page 135.
- For a description of the syntax diagrams please read “Understanding Syntax Diagrams” on page xv.
- Command continuation is allowed. It is specified by a trailing hyphen (-).
- Comments may be specified via SYSIN together with SDAID statements or as separate comment lines.

A /* sign specifies the begin of a comment. All text from the /* sign up to the end of the line is treated as a comment. /* must not start in column 1.

Starting the SDAID Trace Initialization



You start the initialization process as follows:

- If you want to setup the trace from the AR, type only 'sdaid'.
- If you want to initialize the trace in a partition via SYSLOG, enter 'exec sdaid' in that particular partition. Submit 'exec sdaid' or '// exec sdaid' if you use SYSRDR as input device.

Ending the SDAID Trace Initialization

READY In the Attention Routine
/* In a Partition from SYSIPT
ENTER In a Partition from the Console


You end the initialization process with:

- The statement 'READY' or
- An end-of-file indication, that is:
 - /* from SYSIPT;
 - A blank line from the console.

Now you can start the initialized trace with the AR command STARTSD.

You find more information about starting and stopping a trace in Chapter 12, "Start/Stop and End the Trace," on page 163.

Defining the Output Device in Direct Input Mode



You define the destination of your trace output with the OUTDEV statement.

Enter the OUTDEV statement with an operand to get into the direct input mode. Otherwise, if you entered the SDAID statement from the attention routine, the SDAID program would prompt you for the necessary information (prompt mode).

Note: An OUTDEV statement in a partition must contain at least one operand (prompt mode is not possible in a partition).

Defining the Output Device

You can define the following output destinations with the OUTDEV statement:

- Printer
- Tape
- Wraparound buffer
- Wraparound buffer and printer
- Wraparound buffer and tape

For the appropriate OUTDEV and TRACE statement for your trace initialization, see Table 6.

Table 6. OUTDEV Summary

Device	Buffer	Output when:	SDAID Statements	Note
Printer	no	immediately	OUTDEV P=cuu	1
Tape	yes	buffer full	OUTDEV T=cuu	2
-	yes	-	OUTDEV BU=nn	3

Table 6. OUTDEV Summary (continued)

Device	Buffer	Output when:	SDAID Statements	Note
Printer	yes	certain event occurs	OUTDEV BU=nn P=cuu TRACE type OUTP=BU	4
Tape	yes	certain event occurs	OUTDEV BU=nn T=cuu TRACE type OUTP=BU	5

nn..... stands for the wraparound buffer size in units of K bytes.

If you do not define an output device or you use a printer device, nn may have a value of 4 up to the maximum of 256.

If you use a tape device, nn may have a value of 3 up to the maximum of 32.

cuu ... stands for the unit address.

You may abbreviate it in the following way:

00E, 0E, E

type .. represents the type of trace event which forces the output operation.

Notes:

1. No buffer is allocated. The event records are printed on the printer with the device address cuu.
2. The event records are written into an internal buffer. This buffer is written to a tape mounted on the device cuu when it is full or when an ENDS or STOP command is issued.
3. A trace defined with this OUTDEV statement writes the event records into a wraparound buffer. You can retrieve the trace records only with the attention routine command:
DUMP BUFFER,cuu
DOSVSDMP can be used to print the tape. For further information, see "Printing an SDAID or DUMP Command Produced Tape" on page 30.
4. A trace defined with this OUTDEV and TRACE statement prints the contents of the buffer when the trace event (type), defined with the TRACE statement, occurs.
5. The defined wraparound buffer is written to tape (cuu) when the trace event (type) occurs.

Note: From z/VSE 3.1 onwards, the buffer is allocated in the 31-bit SVA. It is no longer part of the SDAID area. The buffer file is rounded up to a multiple of 4K bytes.

SDAID Trace Initialization Example

Assume you want to use SDAID to trace instructions executed in the BG partition. The buffer should be 6K to allow for a reasonable number of trace event records, and should be written to tape when end-of-job is reached or the partition is canceled. Figure 44 on page 88 shows the statements necessary to initialize this trace. The SDAID statements are entered in direct input mode via SYSIN.

```
// EXEC SDAID
OUTDEV BU=8 T=280
TRACE INST=* AR=BG
TRACE CA AREA=BG OUTP=BU,OPT=TERM
/*
```

Figure 44. Example: Initializing an SDAID Trace

For the trace initialized in Figure 44:

- The buffer size is 8K bytes (OUTDEV BU=8 ...);
- The output tape is mounted on unit address 280 (OUTDEV ... T=280);
- SDAID traces instructions (TRACE INST=* ...) executed in the BG partition (... AR=BG) into the buffer;
- The buffer is written to a tape when a CANCEL or an EOJ condition is encountered (TRACE CA OUTP=BU ...);
- Tracing stops after the buffer is written (... OPT=TERM).

The TRACE Statement

The TRACE statement is used to define the trace types you need to get the most reasonable information about errors in your computing environment. For the appropriate trace type, see Table 7.

This section describes the format of all SDAID trace type initializations in direct input mode and shows trace statement and trace initialization examples.

The possible abbreviations are shown through lowercase letters.

Trace statements can be entered up to column 72 of the input line, and may be continued on the following line or lines. To continue a statement, enter at least one blank character and a hyphen (-) after the last operand in the first line, and continue the statement between columns 1 and 72 of the following line.

Enter all operands in the specified order. Refer to the examples in chapter Command Input Paths starting on page Command Input Paths on page 135.

Summary of Trace Types

Table 7. Trace Type Summary

Trace Type	Provides a Trace of:	Page
BRanch	Successfully executed branch instructions	89
BUffer	The trace buffer when it is full	90
CAncel	Program (main task) cancel or EOJ	90
EXternal	External interrupts	91
GETVIS	Getvis / Freevis requests	93
INSTruction	Selected or all instruction(s) execution	95
IO	I/O interrupts	96
LOCK	Lock / unlock requests of resources	97
MONitorcall	MC instructions	99
PGMCheck	Program checks	100
PGMLoad	Phase load requests, or actual load	101


Table 7. Trace Type Summary (continued)

Trace Type	Provides a Trace of:	Page
SSCH	Start Subchannel instructions	103
STorage	Storage alterations	103
SVC	Executed supervisor calls	105
VTAMBU	Usage of VTAM buffers	107
VTAMIO	VTAM I/O operations	107
XPCC	XPCC communication requests	108

Besides the type of the trace and some definitions which belong to the trace type, other keyword operands like AREA, OUTPUT or OPTION may be defined to limit the trace or to produce additional trace output.

These other operands are grouped together and their format is shown under “Additional Definitions” on page 110.

BRanch Trace



Note: For performance reasons, ARea=all requires a specification of a limited address range via the ADDRESS parameter.

See for area_definition page 111, jobname definition page “ARea or JOBNAME Definition” on page 111, output_definition page 113, option_definition page 116.

For a description of the trace and an example of the output, see “BRANCH Trace” on page 56.

Initialization Example


```

// EXEC SDAID
OUTDEV T=280
TRACE BR AR=BG ADDR=0:* -
          OPT=(NOJCL,SUP)
/*
  
```

The following items are covered by the trace set up as shown in the example.

- Use the SYSRDR device to set up the trace.
- Direct the output of the trace data to the tape at device address 280.
- The branch instructions under the control of the BG partition are to be traced.
- The whole system is to be observed to record the instructions executed for the BG partition, including supervisor routines working for BG.
- Do not trace the job control branch instructions.

BUffer Trace



For option_definition see page 116.

For a description of the trace, see “BUFFER Trace” on page 57.

The BUffer trace output is the collection of all trace records contained in the buffer when a buffer overflow occurs.


Initialization Example

```
// EXEC SDAID
OUTDEV BU=6 T=280
TRACE BR AR=BG ADDR=0:* -
      OPT=NOJCL
TRACE BU
/*
```

The following items are covered by the trace set up as shown in the example.

- Use the SYSRDR device to set up the trace.
- Trace all successfully executed branch instructions.
- The branch instructions of the BG tasks are to be traced.
- Observe the whole storage.
- Collect the trace data in a 6K byte buffer.
- Output the buffer whenever it is full.
- Write the output to the tape on device address 280.
- Do not trace the job control branch instructions.

CAnceL Trace



Note: The parameters ADDRESS, OFFSET, PHASE and LTA are not applicable for the CANCEL trace.

See for area_definition page 111, jobname definition page 111, output_definition page 113, option_definition page 116.

For a description of the trace and an example of the output, see “CANCEL Trace” on page 57.

Statement Example

```
TRACE CA AR=BG -  
      OUTP=BU
```

The trace statement shown in the example initializes a Cancel trace which writes an event record and the buffer whenever a cancel or EOJ condition occurs in the BG partition.


Initialization Example

```
// EXEC SDAID  
OUTDEV P=E  
TRACE CA AR=BG -  
      OUTP=(DUMP PART OFF=78:1000)  
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.
- Trace cancel and EOJ conditions.
- Print the area between BG relative address X'78' and X'1000' together with the event record.
- Output the trace record on the printer at device address 00E.

EXTERNAL Trace



Notes:

1 Up to eight codes may be specified.

See for output_definition page 113, option_definition page 116.

'code' may be one to 8 of the following:

0040	Interrupt key	
1003	TOD-clock sync check	
1004	Clock comparator	
1005	CPU timer	
1200	Malfunction alert	
1201	Emergency signal	
1202	External call	
2401	Service signal	
2402	Logical device	* z/VM CP
2603	PFAULT handshaking	* z/VM CP
4000	IUCV, APPC	* z/VM CP
4001	VMCF	* z/VM CP

EXTERNAL=* traces all external interruptions.

If you specify more than one external interrupt type, separate them by one or more blanks or by a comma (with or without blanks) and enclose them in brackets.

For a description of the trace and an example of the output, see “EXTERNAL Trace” on page 58.

Statement Example

```
TRACE EXT=0040 -  
      OUTPUT=(TOD,BU)
```

The example shows an external interrupt trace. 0040 is defined as external interrupt type. This interrupt is used to have the wraparound buffer and the TOD clock recorded or printed together with the external trace event record.

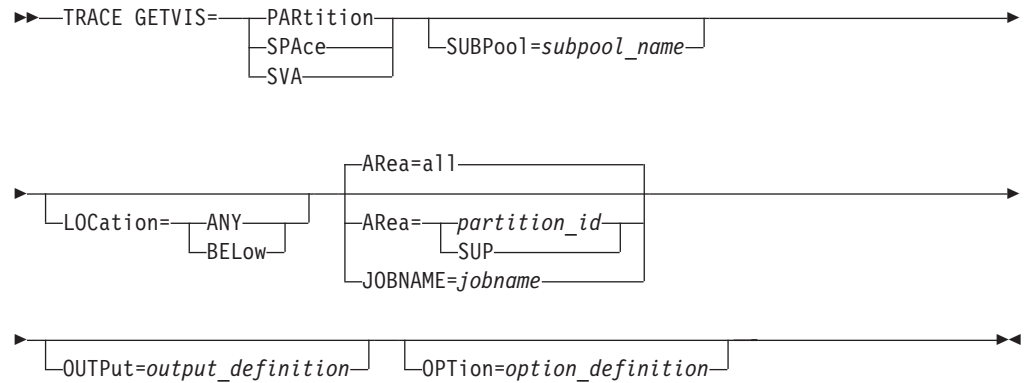
Initialization Example

```
// EXEC SDAID  
OUTDEV T=280  
TRACE EXT=(2401 1005) -  
      OUTPUT=TOD  
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.
- Trace all signal and timer interrupts.
- Output the trace data to the tape on device address 280.
- Add the TOD clock to the trace data output.

GETVIS Trace



See for jobname definition page 111, output_definition page 113, option_definition page 116.

GETVis=PARTition | SPAcE | SVA

GETVis=PARtition traces all Getvis or Freevis requests within the Partition Getvis area.

GETVis=SPAcE traces all Getvis or Freevis requests within the Space Getvis area of a dynamic partition.

Note: If you issue a Space Getvis request for a *static* partition, the required space will be allocated in the SVA.

GETVis=SVA traces all Getvis or Freevis requests within the Shared Virtual Area (SVA).

SUBPool=subpool_name | nnn* | nnn<hex> | DEFAULT

A subpool_name may consist of up to six characters.

- If you specify a subpool name, Getvis or Freevis requests will be traced that are within the specified subpool.
- If you do not specify a subpool name, *all* Getvis or Freevis requests will be traced.

You can enter a subpool as one of the following:

subpool_name Printable characters (EBCDIC) which must follow the naming conventions described in the Getvis / Freevis macros.

*nnn** If a character string is followed by an asterisk (*), all requests will be traced to those subpools whose names begin with the character string.

nnn<hex> A subpool name may be a mixture of printable and hexadecimal characters. Characters enclosed in angled brackets < and > will be taken as hexadecimal characters. All requests will be traced to those subpools whose names have this format. Here are some examples of mixed printable and hexadecimal characters:

- SUBP=INLC<21> refers to a subpool with an internal representation of C9D5D3C321.
- SUBP=INLC21 refers to a subpool with an internal representation of C9D5D3C3F2F1.

DEFAULT Causes all requests to the common default subpool to be

traced. This default subpool is used if you do not specify a subpool_name in the GETVIS invocation macro.

Note: DEFAULT is accepted although it has a length that is greater than six characters. This is because DEFAULT is treated as a keyword and not as a subpool name.

LOCation=BELOW | ANY

LOCation can take one of two values:

LOCation=BELOW

Causes only those Getvis / Freevis requests to be traced that are within the 24-bit Getvis area or within the 24-bit SVA.

LOCation=ANY

Causes *all* Getvis / Freevis requests to be traced that are within the 24-bit or 31-bit areas.

ARea=*partition_id* | SUP | ALL

Causes Getvis / Freevis requests to be traced for tasks running within the specified area. If you omit the ARea operand, all Getvis / Freevis requests will be traced that are executed within the system.

For a description of the trace and an example of the output, see “GETVIS / FREEVIS Trace” on page 58.

Statement Examples

```
TRACE GETVIS=PAR SUBP=MYPOOL AREA=BG
TRACE GETVIS=SVA LOC=BEL
TRACE GETVIS=SVA SUBP=DEFAULT AREA=BG
TRACE GETVIS=PAR SUBP=INLC<00>
```

In the above four examples, the TRACE statements shown:

1. Trace all BG requests within the subpool MYPOOL in the BG partition Getvis area.
2. Trace all system-wide requests within the 24-bit SVA.
3. Trace all BG requests within the common system default subpool.
4. Trace all system-wide requests within the subpool INLC00, where “00” is treated as a hex character.

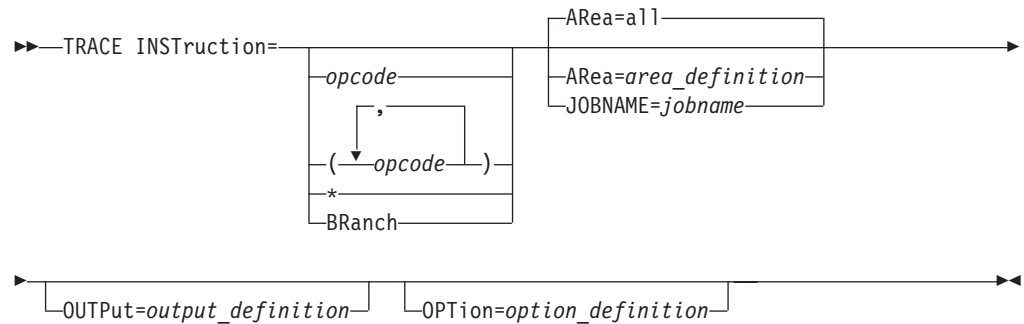
Initialization Example

```
// EXEC SDAID
OUTDEV T=280
TRACE GETVIS=PAR AR=BG OPT=NOJCL
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.
- Trace all partition Getvis requests for BG.
- Output the trace data to the tape on device address 280.
- Do not trace the job control Getvis requests.

INSTRUCTION TRACE



Note: For performance reasons, ARea=all requires a specification of a limited address range via the ADDRESS parameter.

See for area_definition page 111, jobname definition page 111, output_definition page 113, option_definition page 116.

opcode

(one to eight) entered as either one-byte or two-byte hexadecimal instruction codes.

If you specify more than one operation code, separate them by one or more blanks or by a comma (with or without blanks) and enclose them in brackets.

* (asterisk)

requests a trace of *all* executed instructions.

BRanch

requests that all types of branch instructions be recorded.

For a description of the trace and an example of the output, see "INSTRUCTION Trace" on page 59.

Statement Examples

```
TRACE INST=D7 AR=BG
TRACE INSTR=* AR=BG ADD=0:*
TRACE INST=BR AR=BG
TRACE INST=(92 D204) AR=BG
```

The statements shown:

- Trace CLC instructions in BG partition
- Trace all BG task instructions
- Trace branch type instructions in BG partition
- Trace selected instructions in BG partition.


Initialization Example

```
// EXEC SDAID
OUTDEV BU=4 T=280
TRACE INST=D2 AR=BG -
      OPT=NOJCL
TRACE PGMC=* AR=BG ADD=0:* -
      OUTP=BU
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.
- Collect all trace data into a wraparound buffer.
- The MVC instructions of BG tasks are to be traced.
- Write the buffer to the tape on device address 280 when a program check occurs in the BG partition.
- Do not trace the job control MVC instructions.

I/O Interrupt Trace



See for area_definition page 111, jobname definition page 111, for cuu, cu and c page 113, output_definition page 113, option_definition page 116.

Note that for the area and jobname definitions, the parameters ADDRESS, OFFSET, PHASE, and LTA are not applicable.

For a description of the trace and an example of the output, see "IO Trace (I/O Interrupt)" on page 60.

Statement Examples

```
TRACE IO UNIT=(130 133) -
      OUTP=(CCWD=512)
```

```
TRACE IO CU=28 OUTP=CCW
```

The TRACE statements shown in the example define additional trace output (CCWD=512 and CCW) to the normal IO trace event records.

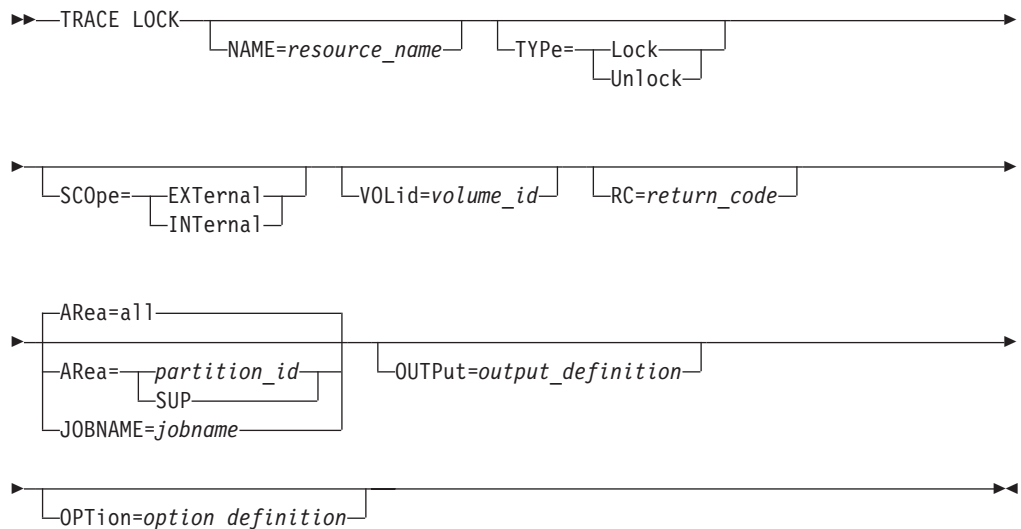
Initialization Example

```
// EXEC SDAID
OUTDEV T=280
TRACE IO -
      UN=320 -
      OUTP=CCWD=512
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.
- Trace all I/O interrupts from device with the address 320.
- Output the trace data to the tape on device address 280.
- Add the CCW plus up to 512 bytes of transferred data to the trace event record.

LOCK Trace



See for jobname definition page 111, output_definition page 113, option_definition page 116.

NAME=resource_name | *nnn** | *nnn<hex>*

A resource name may consist of up to twelve characters.

- If you specify a resource name, lock or unlock requests will be traced that are related to the specified resource.
- If you do not specify a resource name, *all* lock or unlock requests will be traced.

You can enter a resource name as one of the following:

resource_name Printable characters (EBCDIC) which must follow the naming conventions described in the DTL invocation macros.

*nnn** If a character string is followed by an asterisk (*), all requests will be traced to those resources whose names begin with the character string.

nnn<hex> A resource name may be a mixture of printable and hexadecimal characters. Characters enclosed in angled

brackets < and > will be taken as hexadecimal characters. All requests will be traced to those resources whose names have this format. Here are some examples of mixed printable and hexadecimal characters:

- NAME=INLC<21> refers to a resource with an internal representation of C9D5D3C321.
- NAME=INLC<21> refers to a resource with an internal representation of C9D5D3C3F2F1.

TYPe=Lock | Unlock

Type can take one of two values:

TYPe=Lock Causes only the *locking* of one or more resources to be traced.

TYPe=Unlock Causes only the *unlocking* of one or more resources to be traced.

If this parameter is omitted, both locking and unlocking of one or more resources will be traced.

SCOpe=EXTeRnal | INTeRnal

Scope can take one of two values:

SCOpe=EXTeRnal Causes only *external* locks to be traced.

SCOpe=INTeRnal Causes only *internal* locks to be traced.

If this parameter is omitted, both external and internal locks will be traced.

VOLid=volume_id

A six-character volume ID. Only events related to the specified volume_id will be traced. If the parameter is omitted, all volume IDs will be traced.

Note: The parameters scope and volume ID are mutually exclusive. This means, you can only specify one of these parameters.

RC=nn | >nn | <nn | (nn ...)

Return code can take one of four values:

RC=nn Causes all events with a final return code equal to the specified return code to be traced.

RC=>nn Causes all events with a final return code greater than the specified return code to be traced.

RC=<nn Causes all events with a final return code lower than the specified return code to be traced.

RC=(nn ...) You can specify up to sixteen return codes within parentheses. This causes all events with a final return code equal to one of the specified return codes to be traced.

ARea=partition_id | SUP | ALL

Causes lock / unlock requests to be traced for tasks running within the specified area. If you omit the ARea operand, *all* lock / unlock requests will be traced that are executed within the system.

For a description of the trace and an example of the output, see “LOCK / UNLOCK Trace” on page 61.

Statement Examples

```
TRACE LOCK
TRACE LOCK NAME=MYRESOURCE TYPE=LOCK
TRACE LOCK RC=>0 SCOPE=INT AREA=BG
TRACE LOCK VOLID=SHARE1
```

In the above four examples, the TRACE statements shown:

1. Trace all lock / unlock events.
2. Trace all lock events for resource MYRESOURCE.
3. Trace all internal BG lock / unlock events with a final return code greater than zero.
4. Trace all lock / unlock events related to VOLID=SHARE1.

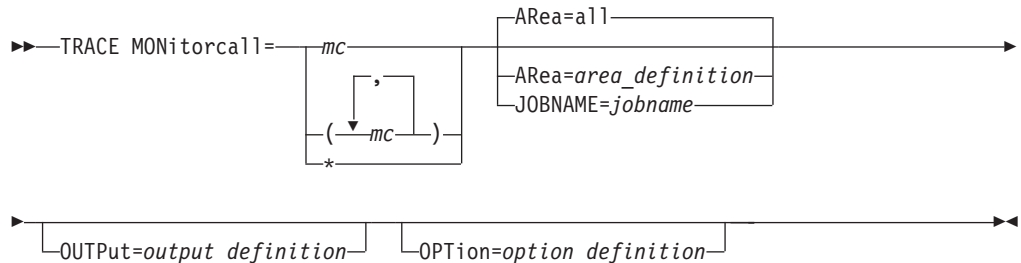
Initialization Example

```
// EXEC SDAID
OUTDEV T=280
TRACE LOCK AR=BG OPT=NOJCL
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.
- Trace all lock / unlock events for BG.
- Output the trace data to the tape on device address 280.
- Do not trace the job control Getvis requests.

MONitor Call Trace



See for `area_definition` page 111, `jobname` definition page 111, `output_definition` page 113, `option_definition` page 116.

mc (monitor classes)

defines the monitor class of the MC instructions to be traced. Only the monitor call instructions with the defined class are traced. Up to eight classes may be defined.

The monitor classes must be specified as one-digit hexadecimal values. If you specify two or more classes, they must be enclosed in brackets and separated by one or more blanks, or by a comma with or without one or more blanks.

You may specify any valid monitor class; however, SDAID ignores a specification of class 2, because class 2 is used by SDAID to control tracing.

*** (asterisk)**

provides an event record for any execution of an MC instruction (except an MC instruction with class 2 specified) within the range of the trace operation.

For a description of the trace and an example of the output, see "MONITORCALL Trace" on page 62.

Statement Examples

```
TRACE MON=3 AREA=ALL
```

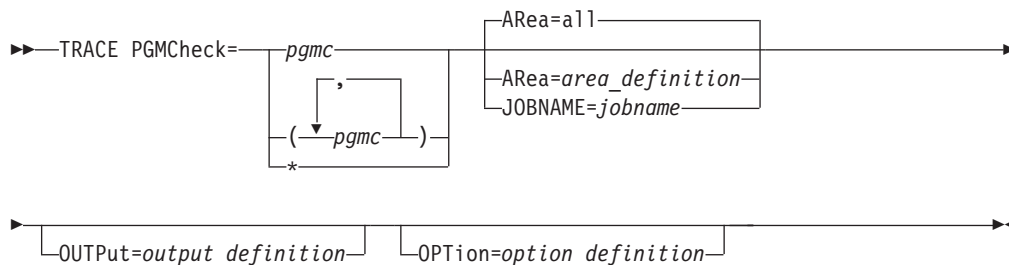
```
TRACE MON=* AREA=ALL
```

```
TRACE MON=(3 4) AREA=ALL
```

The statements shown:

- Trace all class 3 monitor calls;
- Trace all monitor calls;
- Trace all class 3 and 4 MCs.

PGMCheck Trace



See for area_definition page 111, jobname definition page 111, output_definition page 113, option_definition page 116.

pgmc At least one program interruption code (up to 16) must be specified in hexadecimal notation; leading zeros may be omitted.

If you specify more than one program interruption code, they must be enclosed in parentheses and separated by one or more blanks, or by a comma with one or more blanks.

*** (asterisk)**

requests a trace of all valid program interrupt codes with a value less than X'40', except those page or segment translation exceptions which are caused by the temporary absence of a storage page. The specification PGMCheck=(10 11) traces all page or segment translation exceptions.

For a description of the trace and an example of the output, see "PGMCheck Trace (Program Check)" on page 63.

Statement Examples

```
TRACE PGMC=5 AR=BG
TRACE PGMC=* AR=BG ADD=0:*
TRACE PGMC=(1 A 11) AR=BG
```

The statements shown:

- Trace program check addressing exceptions in BG partition;
- Trace all program checks of BG tasks;
- Trace the program checks of BG partition with interruption codes:
 - 1 ... operation exception;
 - A ... decimal overflow exception;
 - 11 ... page translation exception.


Initialization Example

```
// EXEC SDAID
OUTDEV BU=4 T=280
TRACE INST=D2 AR=F2 -
      OPT=NOJCL
TRACE PGMC=1 AR=F2 ADD=0:* -
      OUTP=BU
/*
```

The following items are covered by the trace set up as shown in the example:

- Use the SYSRDR device to set up the trace;
- Direct the output to the tape on device address 280;
- Trace the MVC instructions (D2) of the F2 partition;
- Collect the trace event records in a 4K bytes wraparound buffer.
- Write the buffer to the output tape when a program check interrupt with interruption code 'operation exception (0001)' occurs in the F2 partition.
- Do not trace the instructions executed during job control processing.

Program Load Trace (Fetch/Load Trace)



See for `area_definition` page 111, `jobname` definition page 111, `output_definition` page 113, `option_definition` page 116.

REQ defines, that an event record for each request for fetching/loading a phase is to be written.

HDL defines that an event record is to be written each time a phase fetch/load request is handled; that is, when a requested phase is actually loaded into storage for execution.

ALL combines REQ and HDL. This is the default.

PH=phase

defines the phase whose program load events should be traced.

If the ARea definition is included, only the following ADDRESS definition is allowed (without OFFset, PHase, LTA):

ADDRESS=addr1:addr2 | addr1:* | 0:*

defines that only program load events occurring within the specified address range are to be traced.

SDAID records the SVCs issued within the address range and those load completion events that occur if the phase is loaded into the specified address range.

For a description of the trace and an example of the output, see “PGMLOAD (Fetch/Load) Trace” on page 64.

Statement Examples

```
TRACE PGML AR=BG
TRACE PGML PH=PROGR1 AR=F2
```

The statements shown:

- Trace all program load events for BG tasks;
- Trace all F2 task program load events for the phase PROGR1;


Initialization Example

```
// EXEC SDAID
OUTDEV BU=6 T=280
TRACE PGML HDL AR=BG -
      OUTPUT=(DUMP PART OFF=1000:2000) -
      OPT=NOJCL
TRACE PGMC=1 AR=BG ADD=0:* -
      OUTPUT=(DUMP PART OFF=1000:2000 -
      BUFFER) -
      OPT=NOJCL
/*
```

The following items are covered by the trace setup as shown in the example:

- Use the SYSRDR device to set up the trace;
- Trace all fetch/load executions of the BG partition;
- Record the trace data in a 6K bytes wraparound buffer;
- Write the trace data to the output tape at device address 280 when a program check operation exception (interrupt code 1) occurs in the BG partition;
- Add a dump of the BG area between relative address 1000 to 2000 to both event records;
- Do not trace the job control activities.

SSCH Instruction Trace



For jobname see page 111, UNit=page 113, output_definition page 113, option_definition page 116.

AREa=partition-id | SUP | ALL

causes SSCH instructions of tasks running in the specified area to be traced. Only the specifications shown above are possible. If you omit the AREa operand, all SSCH instructions executed in the system will be traced.

UNit,CU,CHannel

limits the trace to SSCH instructions to a certain unit, control unit or channel. If you omit these operands, no device address limitation is used.

For a description of the trace and an example of the output, see “SSCH Instruction Trace” on page 65.


Statement Examples

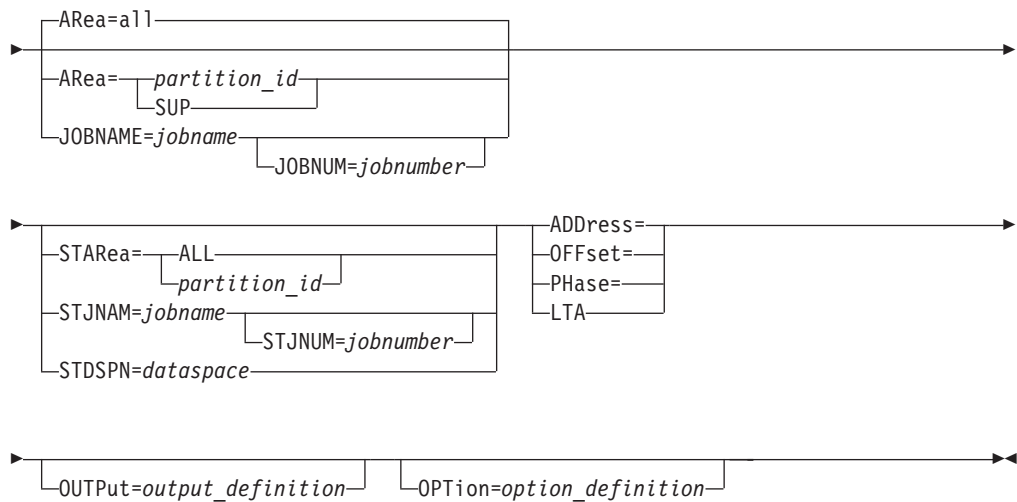
- (1) TRACE SSCH AR=F4 UNIT=009 -
OUTPUT=CCB
- (2) TRACE SSCH CHANNEL=(2 3) -
OUTP=TOD

The TRACE statements in the example define the following functions:

- (1) Trace each SSCH instruction of F4 tasks for the device with the device address 009. Add a dump of the CCB to each SSCH trace event record.
- (2) Trace all SSCH instructions which concern the channels 2 and 3. Add the time of day entry (TOD) to each SSCH trace event record. You will find an example of such a TOD entry under “Trace Output Example with OUTPut=TOD” on page 80.

Storage Alteration Trace





Note: For performance reasons, ARea=all requires a specification of a limited address range via the ADDRESS parameter.

See for area_definition page 111, jobname definition page 111, address-, offset-, and phase-definition page 111, output_definition page 113, option_definition page 116.

The storage alteration trace monitors instructions which alter a specified storage location. The altering program (source of alteration) and the storage area to be altered (target of alteration) may be in the same space or in a different space. A program running in the primary address space A can alter a storage area in the primary space A, or in an address space B, or in a data space C. The keywords AREA=partition-id|SUP|ALL and JOBNAME specify the tasks which alter a storage location (source of alteration). The keywords STAREA=partition-id|ALL, STJNAM, and STDSPN specify the target space where storage alteration is to be monitored. The keywords ADDRESS, OFFSET, PHASE, and LTA specify the target address.

PATtern=xxxxxxx

defines a hexadecimal storage pattern up to four bytes long.

If you specify an odd number of digits, a zero is inserted to the left of the first specified hexadecimal digit.

ARea=

JOBNAME=

defines those tasks whose alteration activities you want to trace. If you do not know which task does the alteration, specify ARea=ALL to have all tasks of the system watched.

STAREa=ALL

specifies a storage alteration within any address space or data space.

STAREa=partition-id

specifies a storage alteration in the private address space where the named partition is allocated.

STJNAM=jobname[STJNUM=jobnumber]

specifies a storage alteration in the private address space where the named POWER job (with the named job number) executes. Note that SDAID does not

accept POWER job names containing the character '-'. (The SDAID command language uses this character as a continuation character.)

STDSPN=dataspace

specifies a storage alteration within the specified data space.

If none of the keywords STARea, STJNAM, STJNUM, and STDSPN is specified, the corresponding ARea, JOBNAME, and JOBNUM keywords apply.

ADDRESS= | OFFSET= | PHASE= | LTA

specifies the address (or offset, phase, LTA) of the **target** area (where the alteration takes place).

Do not use the definition OFFSET if STARea=ALL or STDSPN= is specified. Do not use the definition PHASE or LTA if STDSPN= is specified.

For a description of the trace and an example of the output, see "STORAGE Alteration Trace" on page 66.

Statement Example

```
TRACE ST PATT=D205 -  
        AR=ALL -  
        ADD=65674:65675
```

The example records all instructions which alter the contents of two bytes starting with storage location X'65674' to the pattern X'D205'.

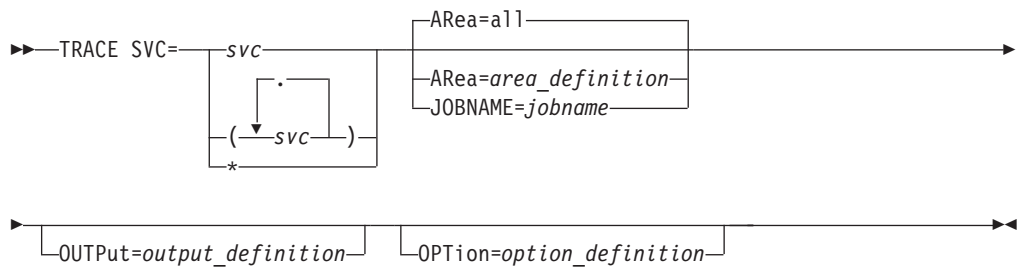
Initialization Example

```
// EXEC SDAID  
OUTDEV BU=3 T=280  
TRACE INST=* AR=BG -  
        OUTP=GREG -  
        OPT=NOJCL  
TRACE ST PATT=FFFF -  
        AR=ALL -  
        ADD=40100:40101 -  
        OUTP=BU  
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.
- Trace all BG instructions excluding the job control instructions.
- Record the event records in a 3K bytes wraparound buffer.
- Write the buffer together with a storage alter trace event record to the tape on device address 280 when the storage area with the address 40100 to 40101 is altered to X'FFFF'.
- Observe all tasks of your system, in respect to altering the storage X'40100'-X'40101' to X'FFFF'.

Supervisor Call Trace



See for area_definition page 111, jobname definition page 111, output_definition page 113, option_definition page 116.

svc defines a certain Supervisor Call Code. You may define up to 16 different SVC codes. Specify the SVC code in hexadecimal notation.

If you specify more than one SVC code, the codes must be enclosed in parentheses and separated by one or more blanks, or by a comma with or without one or more blanks.

***** (asterisk) defines that all SVC instructions are to be traced.

For a description of the trace and an example of the output, see “SVC Trace (Supervisor Call)” on page 67.

Statement Examples

```
TRACE SVC=* AR=BG ADD=0:*
```

```
TRACE SVC=A AR=BG
```

```
TRACE SVC=(1D 25) AR=BG
```

The statements shown:

- Trace all BG SVCs in BG partition and in system areas;
- Trace set timer SVCs (X'A') in BG partition;
- Trace BG partition WAITM and STXIT AB SVCs.


Initialization Examples

```
// EXEC SDAID
OUTDEV T=280
TRACE SVC=* AR=BG ADD=0:* -
      OPT=NOJCL
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.
- Trace all supervisor call instructions.
- The SVC instructions from the BG partition are to be traced.
- Output the trace data to the tape on device address 280.
- Do not trace the job control branch instructions.

VTAM BUffer Trace



See for output_definition page 113, option_definition page 116.


For a description of the trace and an example of the output, see “VTAMBU Trace (VTAM Buffer)” on page 68.

Statement Example

```
TRACE VTAMBU
```

The example defines a VTAM buffer trace.

VTAMIO Trace



See for area_definition page 111, jobname definition page 111, for cuu, cu and c page 113, output_definition page 113, option_definition page 116.

A VTAMIO trace requires an area definition. Define the operands as shown above.

For a description of the trace and an example of the output, see “VTAMIO Trace” on page 68.

Statement Example

```
TRACE VTAMIO AREA=F3
```

The example defines a VTAMIO trace for the F3 tasks.

Initialization Example

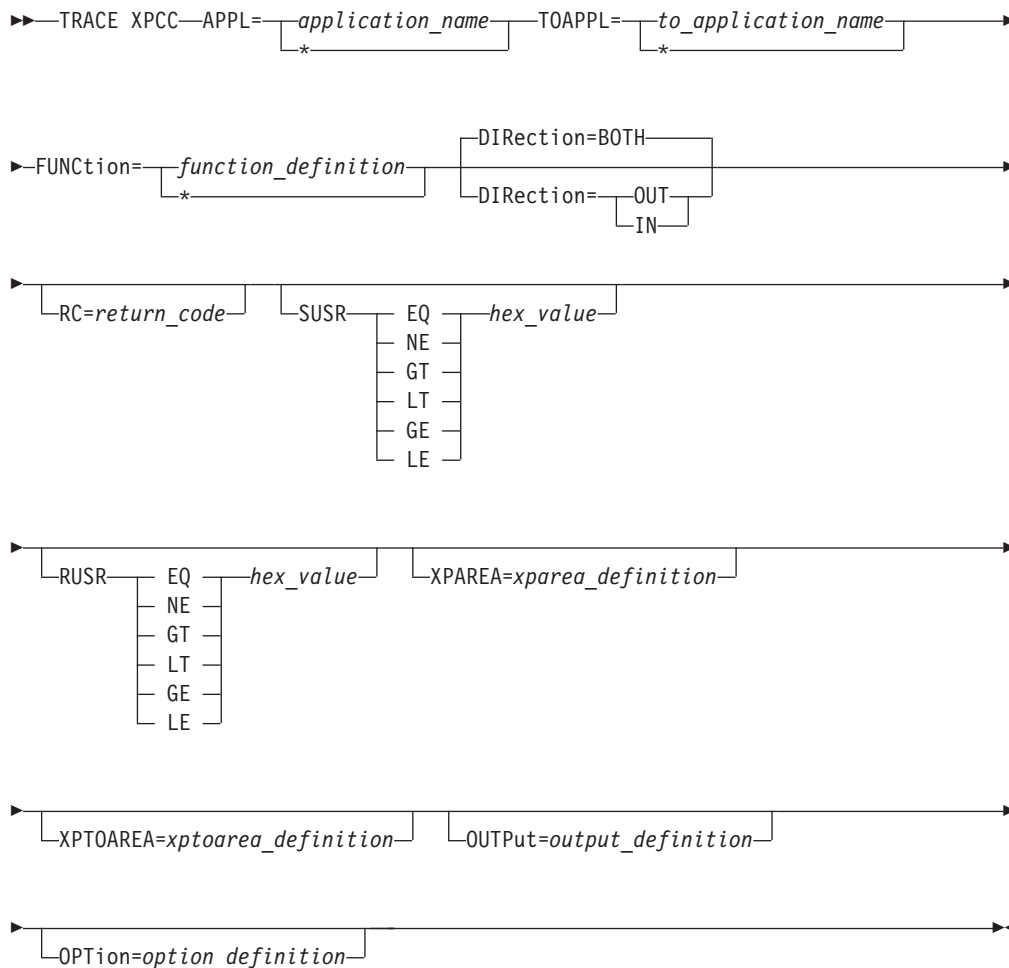
```
// EXEC SDAID  
OUTDEV T=280  
TRACE VTAMIO AR=F3 -  
                UN=020  
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.

- Trace all VTAM I/O operations concerning unit at address 020 and F3 tasks.
- Output the trace data to the tape on device address 280.

XPCC Trace



See for output_definition page 113, option_definition page 116.

APPL=*application_name* | *nnn** | *

An application name may consist of up to eight characters which must follow the naming conventions described in the XPCC invocation macro.

You can enter an application name as one of the following:

application_name

Only XPCC requests will be traced that have the specified name.

*nnn**

If a character string is followed by an asterisk (*), all requests will be traced to those XPCC requests whose names begin with the character string.

*

If you specify an asterisk (*), all requests will be traced for all XPCC requests.

TOAPPL=*to_application_name* | *nnn** | *

The rules for using TOAPPL are the same as for APPL (above).

FUNCTION=TERMQsce | TERMPrg | TERMIn | DISCAI | DISCPrg | DISCOnn | SENDI | SENDR | SEND | RECeive | REPlY | IDent | COnnect | CLear | PUrge | *

FUNCTION can take any of the functions listed above, or asterisk (*). If you enter one of the functions, only XPCC events of the specified function will be traced. If you enter an asterisk (*), XPCC events for *all* functions will be traced.

DIRection=IN | OUT | Both

You can enter a DIRection as one of the following:

DIR=IN Only incoming XPCC events will be traced.
DIR=OUT Only outgoing XPCC events will be traced.
DIR=Both Both incoming and outgoing XPCC events will be traced.

If this parameter is omitted, DIR=Both is used as default.

RC=return_code

If you specify a value for *return_code*, only events with a final return code equal to this value will be traced. If this parameter is omitted, events with any final return code will be traced.

SUSR | RUSR comparator hex_value

Compares a specified string of hexadecimal characters with the contents of the corresponding field contained within the XPCC control block.

- SUSR corresponds to field IJBXSUSR in the XPCC control block.
- RUSR corresponds to field IJBXRUSR in the XPCC control block.
- *comparator* can be one of: EQ, NE, GT, LT, GE, or LE.
- *hex_value* is a string of up to sixteen hexadecimal characters. Characters within this string that should *not* be compared can be substituted by a dot (.).

Here are some examples of the use of SUSR:

```
SUSR EQ FFFFFFFFFFFFFFFF (check all hex characters in IJBXSUSR)
SUSR NE 1020              (check only first four hex characters)
SUSR LT .....FFD         (check only characters 7 to 9)
SUSR GT FF...FF..        (THIS IS AN INVALID STATEMENT!)
```

If the SUSR | RUSR parameter is omitted, no checking will be performed.

XPAREA=syslog_id

Defines the partition where the application that is defined using APPL=*application_name* is running, in a pair of two interacting partitions. If the parameter is omitted, all XPAREAS are assumed.

XPTOAREA=syslog_id

Defines the partition where the application that is defined using TOAPPL=*application_name* is running, in a pair of two interacting partitions. If the parameter is omitted, all XPTOAREAS are assumed.

Note: Using the OUTPUT=*output_definition* parameter (described on page 113), you can request this type of output:

```
OUTPUT=XPCCB (prints the XPCC Control Block)
OUTPUT=XPDATABU (prints the contents of the Transmit Data Buffer)
```

For a description of the trace and an example of the output, see “XPCC Trace” on page 69.

Statement Examples

```
TRACE XPCC APPL=* TOAPPL=* FUNC=* OUTPUT=XPCCB
```

```
TRACE XPCC APPL=RESI TOAPPL=* FUNC=SEND DIR=OUT XPAREA=BG  
XPTOAREA=F7 OUTPUT=XPDATABU
```

In the above two examples, the TRACE statements shown:

1. Trace all XPCC events.
2. Trace outgoing XPCC SEND events from application RESI to any partner application, but only if RESI runs in BG and the partner application is in F7.

Initialization Example

```
// EXEC SDAID  
OUTDEV T=280  
TRACE XPCC APPL=* TOAPPL=* FUNC=*  
/*
```

The following items are covered by the trace setup shown:

- Use the SYSRDR device to set up the trace.
- Trace all XPCC events.
- Output the trace data to the tape on device address 280.

Additional Definitions

```
ARea=, JOBNAME=, ADDRess=, OFFset=, PHase=, OPTion=, OUTPut=,  
UNit=, CHannel=, CU=
```

The following section describes definitions which may follow the trace type specification in the TRACE statement.

Table 8 shows a list of all additional definitions, a summary of their function, and a page reference to their format description and examples in this chapter. The various definitions are described in detail under:

- “Defining the Area to be Traced: AREA Definition” on page 70.
- “Defining the Job to be Traced: JOBNAME Definition” on page 70.
- “Defining the Storage to be Traced: OFFset, ADDRess, PHase, LTA” on page 71.
- “Defining Additional Trace Output: OUTPut Definition” on page 73.
- “Defining the Trace Options: OPTion Definition” on page 81.
- “Defining the Traced I/O Devices” on page 82.

Table 8. Additional Definitions Summary

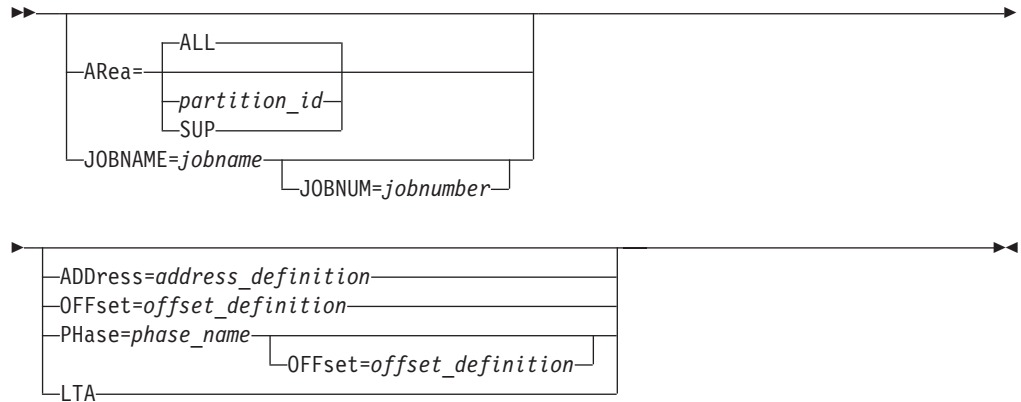
Operand	Function	Page
ARea	Limit tracing to a certain system area	111
JOBNAME	Limit tracing to a certain VSE/POWER job	111
ADDRess	Limit tracing to a certain address range	111
OFFset	Limit tracing in a partition or phase area	112
PHase	Limit tracing to a certain phase	112
UNit	Define the device address	113
CHannel	Define the channel address	113

Table 8. Additional Definitions Summary (continued)

Operand	Function	Page
CU	Define the control unit address	113
OUTPut	Define additional trace output	113
OPTion	Define additional trace options	116

ARea or JOBNAME Definition

The possible storage area definitions together with ARea or JOBNAME are:



ADDRESS=

See "Address Definition."

OFFset=

See "Offset Definition" on page 112 (not for ARea=ALL).

PHase=

See "Phase Definition" on page 112.

LTA Defines the Logical Transient Area as tracing range.

Default Value

If you use ARea=partition-id | SUP or JOBNAME without an additional specification, OFFset=0:* is assumed. OFFset=0:* defines the whole partition (or the area between zero and end-of-supervisor) as trace area.

If you use ARea=ALL without an additional specification, ADDR=0:* is assumed (that is, all virtual storage).

ADDRESS Definition



You can limit the trace to a certain address range within the storage allocated to VSE with the ADDRESS definition.

addr1:addr2

Defines a trace address range in hexadecimal notation in any virtual storage defined to VSE.

For example:

ADD=500000:*

Default Value

If you omit the ADDRESS specification, the default trace address range depends upon the specifications in the AREa= or JOBNAME= parameters.

If you use AREa=*partition_id*, the complete storage allocated to the partition is assumed to be trace address range.

If you use JOBNAME=, the complete storage allocated to the specified JOB is assumed to be trace address range.

If you use AREa=ALL, then ADDR=0:* is assumed (that is, the complete storage range is allocated to VSE).

Note: If you use AREa=*partition_id* and parts of your program are located in the SVA, you must specify ADDR=0:* to include these parts of your program in the trace.

OFFset Definition


You can limit the trace to a certain address range via offsets relative to the defined partition, supervisor or phase with the OFFset definition.

For example:

OFF=200:*

Default Value

If you omit the OFFset definition, 0:* is assumed.

PHase Definition

With the PHase definition the traced storage area is defined by the area occupied by that phase.

phase-name

For example:

PH=PROGRAM1

Default Value


You may limit the traced storage area within the defined phase with the OFFset keyword operand. See “OFFset Definition” on page 112.

For example:

```
PH=PROG OFF=20:400
```

The example above initializes a trace which is active only in the phase with the name PROG between program address X'20' to X'400'.

I/O Device Definition



UNit=cuu

For example:

```
UN=280
```

```
UN=(280 310) Use the parentheses if you  
specify more than one address.
```

```
UN=e Same as 00e.
```

CU=cu

For example:

```
CU=28
```

```
CU=00
```

```
CU=(28 31) Use the parentheses if you  
specify more than one address.
```

CHannel=c

For example:

```
CH=2
```


```
CH=(2 3) Use the parentheses if you  
specify more than one address.
```

Only one of the parameters UNit, CHannel, or CU can be specified in the same TRACE command.

Default Value

If none of the I/O parameters is specified, all devices are traced.

OUTPut Definition



You may take one or more definitions together with one TRACE statement. If you enter more than one OUTPut definition in direct input mode, enclose them in parentheses.

A summary of all definitions which you can specify with OUTPut= is given in Table 9 on page 114. This table contains the format and a short description of the data which is recorded together with the trace event record. For those output

definitions which allow additional definitions, a page reference to the information contained in this chapter is shown.

Table 9. OUTPut Definition Summary

Definition	What it records/prints in addition:	Page
BUffer	Contents of SDAID output buffer	-
CCB	CCB or IORB (TRACE=IO, SSCH, or VTAMIO only)	-
CCW	CCWs, IRB (TRACE=IO, SSCH, or VTAMIO only)	114
CCWD=nnnn	CCWs plus nnnn bytes of data, IRB (TRACE=IO, SSCH, or VTAMIO only)	114
COMReg	Partition communication region	-
CReg	Control registers	-
DUMP	Virtual storage	115
FReg	Floating point registers	-
GReg	General purpose and access registers	-
IOTab	PUB, LUB, ERBLOC, ERRQ, CHANQ	-
LOCKTE	Lock table entry (LOCK trace only)	61
LOWcore	Processor storage from zero to X'2FF'	-
LTA	Logical transient area	-
PTA	Physical transient area	-
PTAB	Partition related control blocks: PCB, PIB, PIB2	-
SUPvr	Supervisor plus GREG and CREG	-
SYSCom	System communication region	-
TOD	Time-of-Day clock	-
TTAB	Task related control blocks: TIB, TCB, PCB, PIB, PIB2	-
XPCCB	XPCC control block (XPCC trace only)	69
XPDATABU	Buffer for data to be transmitted	69

Note: A description of all output definitions is given under “Defining Additional Trace Output: OUTPut Definition” on page 73.

Recording CCW



CCW (channel command word) records/prints the available channel program (CCW chain) plus the CCB and the TOD clock when the trace type is SSCH.

In case of an IO trace only the CCWs which refer to transferred data are recorded or printed.

Specifying this output option for an event other than IO or SSCH is not meaningful.

CCWD=nnnn

(CCW plus data) records/prints up to a maximum of nnnn bytes of the transferred data, the CCB and the TOD clock in addition to the information

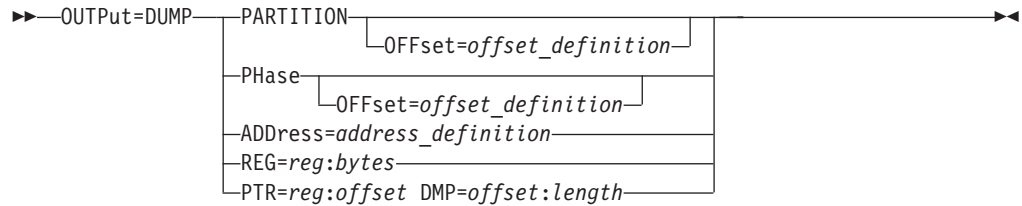
processed with the CCW specification. The number nnnn may be any (decimal) number between 1 and 65535.

The most meaningful trace type to be combined with this output option is the IO trace.

For an example of the output produced with this option, see Figure 36 on page 75.

You may define either CCW or CCWD=nnnn.

Dumping Virtual Storage



DUMP

records or prints the contents of virtual storage.

You may request up to ten different dumps.

You have to specify one or more of the dump area specifications as shown below.

PARTITION

For example, dump the storage beginning with offset X'0' up to X'78' of the partition for which the trace is active.

OUTP=(DUMP PARTITION OFF=0:78)

PHase For example, dump the area starting with relative address X'40' up to relative address X'60' in the phase defined via the 'PHase=' keyword operand.

OUTP=(DUMP PH OFF=40:60)

ADDRESS

For example, dump the contents of two bytes starting on storage location 0080 (hexadecimal). The definition in direct input mode looks like this:

OUTP=(DUMP ADD=80:81)

REG=reg:bytes

For example, dump 16 bytes of storage pointed to by register 15.

OUTP=(DUMP REG=F:10)

PTR=reg:offset DMP=offset:length

For example, dump a four-byte field which is located in a table with an offset of X'20' bytes. The table address is stored in storage pointed to by register 6 plus displacement X'100':

OUTP=(DUMP PTR=6:100 DMP=20:4)


Trace Statement Example: Dump an Area in a Phase

```
TRACE PGMC=* -  
      AR=BG -  
      PH=PHASE1 -  
      OUTP=(GREG DUMP PH -  
            OFF=0:400 LOWC)
```

The following items are covered by the trace setup shown in the example:

- Trace program check interrupts
- Traced tasks: BG partition main and subtasks
- Traced storage area: phase1 storage area
- Additional trace output:
 - general registers (GREG)
 - dump of X'400' bytes of phase1 area starting at relative address 0 (DUMP PH OFF=0:400)
 - low-core (LOWC).

OPTION Definition



For a description of the OPTION definitions, see “Defining the Trace Options: OPTION Definition” on page 81.

OCCurrence Examples

- | | |
|----------------|--|
| OPT=OCC=1:1 | Trace only the first occurrence of the event |
| OPT=OCC=1:* | Trace all occurrences of the event (this is the default value) |
| OPT=OCCUR=5:12 | Trace selected occurrences (5 to 12) of the specified event |

Chapter 10. Initialize an SDAID Trace via a Procedure

This chapter describes how you initialize SDAID traces by using just one job control (JCL) EXEC PROC statement. VSE/Advanced Functions offers a set of predefined JCL procedures to initialize SDAID traces under control of a partition. These procedures are included in the system sublibrary IJSYSRS.SYSLIB.

The most frequently used SDAID functions are covered by these JCL procedures. The JCL procedures contain reasonable default values to ease the SDAID trace initialization process. You may define your own procedures tailored to the requirements of your installation or to a special debugging problem.

Introduction

Besides the direct input mode and prompt mode trace initialization a third initialization method is available under VSE, the initialization via cataloged procedures. These cataloged procedures contain direct input mode command skeletons. You activate the initialization via the job control EXEC PROC statement.

```
// EXEC PROC=trace-type,specification,specification,...
```

The specifications in the EXEC PROC statement are translated to SDAID direct input mode command operands.

Each EXEC PROC statement contains the name of the procedure (trace-type) plus additional specifications. You may define the specifications in any order. A continuation sign has to follow the comma if you use the console for input. If you use SYSIN to enter the procedure statement, the continuation sign has to be in column 72 and the continuation line must start in column 16.

Notational Conventions

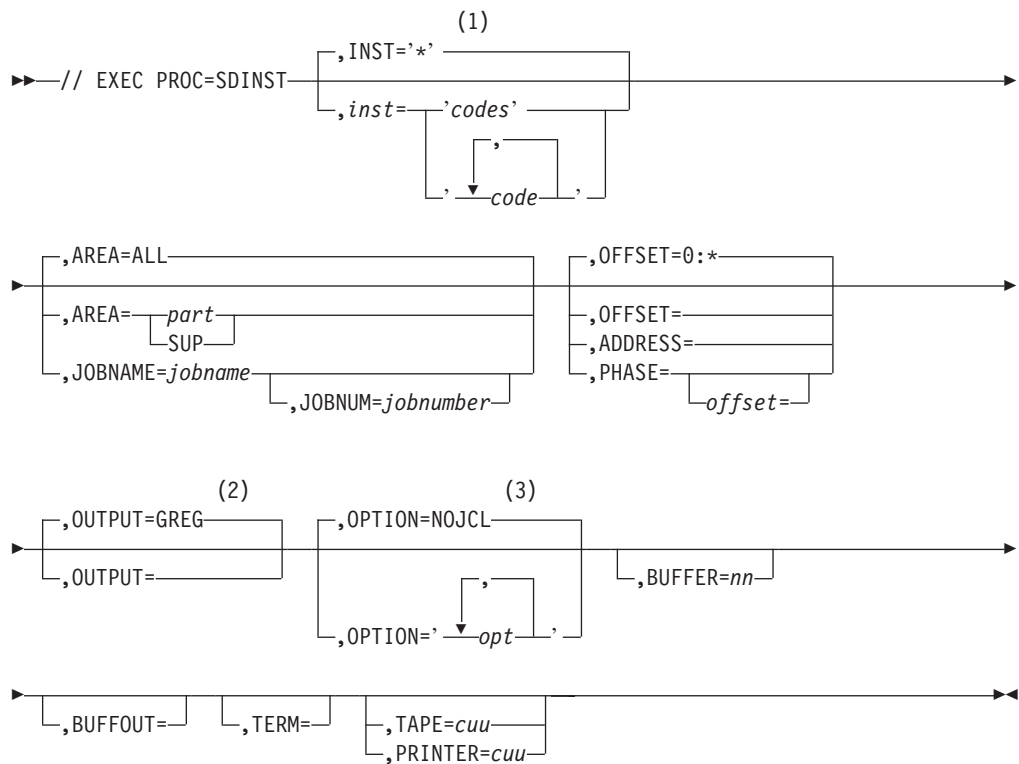
In this chapter, the EXEC PROC statements needed to initialize the various SDAID traces are described in detail. The notation of each EXEC PROC statement description shows you:

- Which operands are optional, and which are mandatory;
- What the default values of the operands are.

The syntax of the EXEC PROC statement follows the conventions for job control statements as described in *z/VSE System Control Statements*.

The SDAID command defaults are not shown in this chapter. The defaults for each trace type are given in “Summary of TRACE Types” on page 56.

Figure 45 on page 118 shows a sample description of the EXEC PROC statement used to call a trace. The handling of the operands and their default values is described in the section following the figure.



Notes:

- 1 The default values are defined in the cataloged procedure SDINST.
- 2 The default values are defined in the cataloged procedure SDINST.
- 3 The default values are defined in the cataloged procedure SDINST.

Figure 45. SDINST Sample Procedure

Default Value Considerations

In an EXEC PROC statement, you can specify the SDAID trace operands in three ways:

1. Using the default value defined in the procedure;
2. Using the default value of the SDAID trace command itself;
3. By specifying a value of your choice in the EXEC PROC statement.

To use a **default defined in the procedure**, simply omit the appropriate operand from the EXEC PROC statement. If the procedure has no default for the operand, this will cause the SDAID default to be used.

To use the **SDAID default value**, nullify the operand in the procedure by coding 'keyword=' in the EXEC PROC statement. For example:

```
EXEC PROC=SDINST,OUTPUT=
```

would cause the SDAID default value for OUTPUT in the trace command to be used.

An operand which has no procedure-defined default value does not have to be nullified. Simply omit the operand from the EXEC PROC statement.

To **specify a value of your choice**, include the appropriate keyword and value in the EXEC PROC statement. This overrides the procedure-defined default, if any, and the SDAID default value. For example:

```
EXEC PROC=SDINST,AREA=BG,OUTPUT=PTAB
```

overrides the procedure defined OUTPUT value GREG. The trace runs as if OUTPUT=PTAB had been specified in the SDAID trace command.

Writing Cataloged Procedures

You can create and catalog your own procedures for particular problem-determination situations.

For example, you can define additional default values, or you can create a procedure for a trace type for which no procedure has been cataloged.

When you write a procedure, consider that you have to follow the correct command-input sequence. For example, the TRACE= definition for some trace types has to be followed by the AREA or JOBNAME specification. You can use the figures shown under “Command Input Path Example” on page 135 to establish the correct command-input sequence.

The Statements of a Cataloged Procedure

The result of the execution of each EXEC PROC statement is a complete direct-input trace initialization. The direct input mode statements are cataloged as:

- Fixed definitions;
- Placeholder definitions;
- Placeholder definitions with default values.

Fixed Definitions

are those definitions in the cataloged procedure which are always active. They cannot be altered or overridden by values specified in the EXEC PROC statement. Code them as you would in direct-mode trace initialization.

Placeholder Definitions

can be replaced by a value which you specify in the EXEC PROC statement. These are handled as follows:

- A placeholder, beginning with an ampersand (&), takes the place of the value after the equals sign (=) in the cataloged trace command (for example: UNIT=&UNIT);
- The placeholder name, without the ampersand, is used in the EXEC PROC statement to provide a definition at execution time (for example: EXEC PROC=SDIO,UNIT=280).

The statement in the cataloged procedure:

```
UNIT=&UNIT
```

Your definition in the EXEC PROC statement:

```
UNIT=280
```

The created direct-input-mode statement:

```
UNIT=280
```

If the operand of a trace statement can have a list of values after the equals sign, one placeholder is still enough. In the EXEC PROC statement, the list must be enclosed in single quotes (this is a requirement of job control). SDAID replaces these quotes with parentheses in the trace initialization statement which the procedures produces. For example:

The statement in the cataloged procedure:

```
UNIT=&UNIT
```

Your definition in the EXEC PROC statement:

```
UNIT='280 281'
```

The created direct-input-mode statement:

```
UNIT=(280 281)
```

Placeholder with Default Value Definitions

You can define default values for placeholders in cataloged procedures. The default value must follow the placeholder and be enclosed in “less-than” (<) and “greater-than” (>) signs (for example: UNIT=&UNIT<280>). In the IO trace procedure in Figure 46, the default output value is specified as follows:

```
OUTPUT=&output<CCWD=256>
```

If you omit OUTPUT=value definition from the EXEC PROC statement, the default value is inserted in the direct input statement, which is generated as:

```
OUTPUT=CCWD=256
```

If you do not want to provide any definition, and also want to avoid the procedure default, you must code:

```
OUTPUT=
```

(with no value) in the EXEC PROC statement. The operand in the procedure is nullified. No OUTPUT definition is inserted in the created direct input statement.

Figure 46 shows an example of a fixed definition, a placeholder definition, and a placeholder with a default value definition. The cataloged procedure in this example is called by the member name under which you cataloged it.

```
// EXEC SDAID
TRACE SSCH AREA=&area           -
          JOBNAME=&jobname      -
          UNIT=&unit            -
          OUTPUT=TODO          -   Fixed Definition
          OPTION=&option
TRACE IO  AREA=&area           -   Placeholder Definition
          JOBNAME=&jobname      -
          UNIT=&unit            -
          OUTPUT=&output<CCWD=256> -   Placeholder with Default
                                     Value Definition
          OPTION=&option
```

Figure 46. Example: Cataloged Procedure

In the two trace types which are initialized:

- The same AREA, UNIT and OPTION values are used for both traces. These values are specified in the EXEC PROC statement;
- If you do not specify the OUTPUT operand, the default OUTPUT=CCWD=256 is defined for the IO trace;
- The SSCH trace event record always contains the time-of-day clock.

Procedures to Initialize SDAID Traces

This section describes the trace procedures available with VSE/Advanced Functions to initialize SDAID traces. The additional keyword operands which you find in the trace procedure statements are described under “Additional Keyword Operands in Trace Procedure Statements” on page 128.

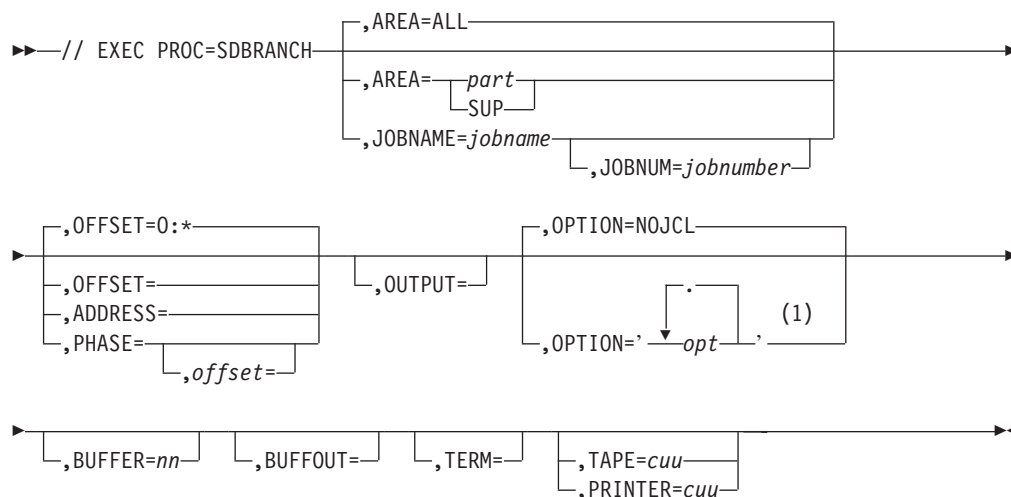
Choose the appropriate procedure from Table 10.

Summary of Trace Procedures

Table 10. Trace Procedures Summary

Procedure	Provides Information on:	Page
SDBRANCH	Successfully executed branch instructions	121
SDINST	Selected or all instruction(s) execution	122
SDIO	I/O interrupts and SSCH instructions	123
SDLOAD	Phase load requests, or actual load	124
SDPGMC	Program check interruptions	125
SDSTOR	Storage alterations	126
SDSVC	Executed supervisor calls	127

Branch Trace Initialization



Notes:

- 1 Up to 6 options may be specified.

See the “Additional Keyword Operands in Trace Procedure Statements” on page 128.

The procedure SDBRANCH initializes traces for all branch instructions which actually caused a branch.

Find the description of the trace type and an example of the output under “BRANCH Trace” on page 56.

Defaults Set in the Procedure

OPTION=NOJCL is active if you omit OPTION=.

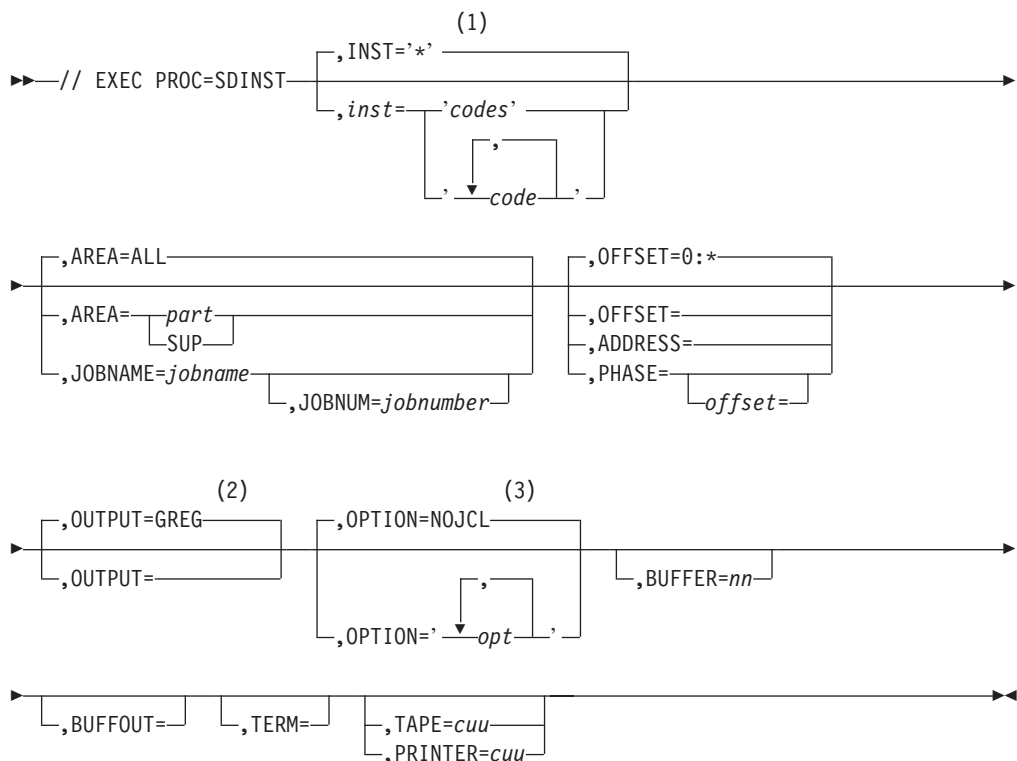
Statement Example

Here are the items of the trace setup shown below:

- Trace type: BRANCH
- Area for which events are collected: storage address 80010 up to address 80100
- Traced tasks: F4 main task and its subtasks
- Output destination: Tape with device address 280
- Avoid the tracing of JCL instructions (default)

```
// EXEC PROC=SDBRANCH,AREA=F4,ADDRESS='80010:80100',TAPE=280
```

Instruction Trace



Notes:

- 1 The default values are defined in the cataloged procedure SDINST.
- 2 The default values are defined in the cataloged procedure SDINST.

3 The default values are defined in the cataloged procedure SDINST.

See the “Additional Keyword Operands in Trace Procedure Statements” on page 128.

The procedure SDINST initializes traces for all instructions or for selected instructions executed within a specified area. Find the description of the trace type and an example of the output under “INSTRUCTION Trace” on page 59.

Defaults Set in the Procedure

If you omit INST, all instructions are traced (INST='*' is the default).
OPTION=NOJCL and OUTPUT=GREG are assumed if you omit both these operands.

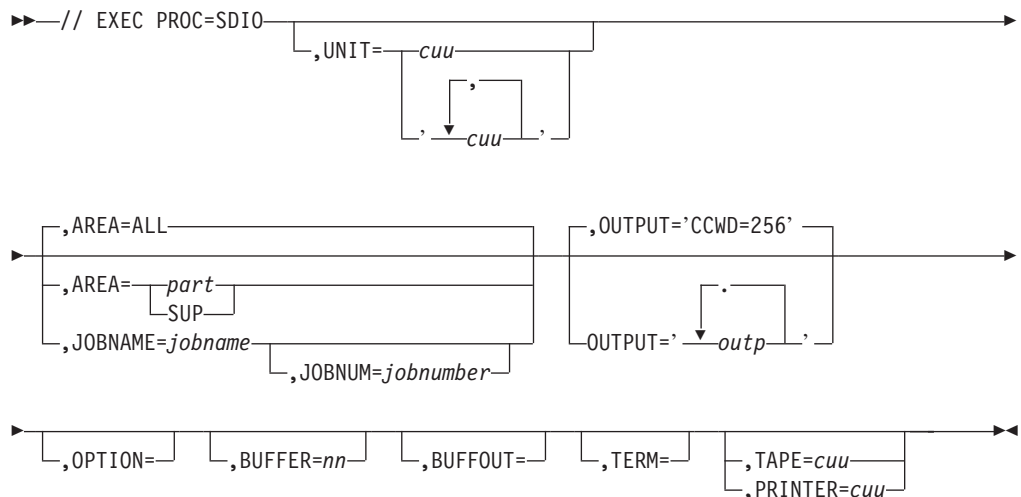
Statement Example

Here are the items of the trace setup shown below:

- Trace type: INSTRUCTION
- Trace all instructions (default)
- Area for which events are collected: storage address 40328 up to address 40350
- Traced tasks: BG main and subtasks
- Additional Output: default GREG output
- Output destination: 16K bytes buffer
- Output device for buffer: tape with device address 281
- Event to write the buffer to tape: program check in BG partition
- Avoid the tracing of JCL instructions (default)
- Note that the continuation sign has to follow the comma if you use the console for input. If you use SYSIN to enter the procedure statement the continuation sign has to be in column 72.

```
// EXEC PROC=SDINST,AREA=BG,ADDRESS='40328:40350',-  
                BUFFER=16,BUFFOUT=PGMC,T=281
```

SSCH and I/O Interrupt Trace



See the “Additional Keyword Operands in Trace Procedure Statements” on page 128.

The procedure SDIO initializes the SSCH instructions and I/O interruptions trace.

Note that the TOD clock entry is added to each SSCH instruction event record.

Find the description of the trace types and examples of the output under “IO Trace (I/O Interrupt)” on page 60 and “SSCH Instruction Trace” on page 65.

Default Set in the Procedure

If you do not define UNIT, all devices are traced.

If you omit the AREA definition, all tasks in the system are traced.

These are both SDAID defaults.

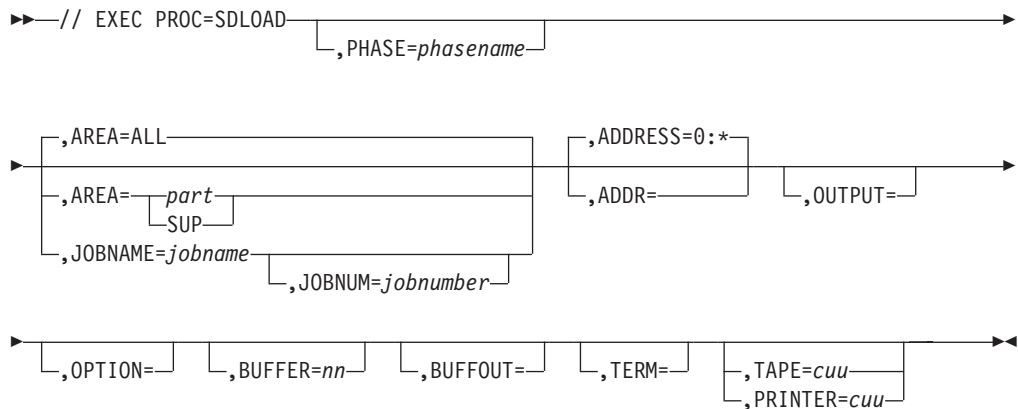
OUTPUT='CCWD=256' is the default definition for the I/O interrupt trace.

Statement Example

- Trace types: IO, SSCH
- Traced tasks: SDAID default value used (ALL)
- Traced unit: 281
- Additional Output: procedure default CCWD=256
- Output destination: printer with device address 00E

```
// EXEC PROC=SDIO,UNIT=281,PRINTER=00E
```

Fetch/Load Trace



See the “Additional Keyword Operands in Trace Procedure Statements” on page 128. The procedure SDLOAD initializes traces for all phase load requests and phase load operations.

For the description of the trace type and an example of the output, see “PGMLOAD (Fetch/Load) Trace” on page 64.

Defaults Set in the Procedure

ADDRESS='0:*' is defined if you omit ADDRESS=.

Statement Example

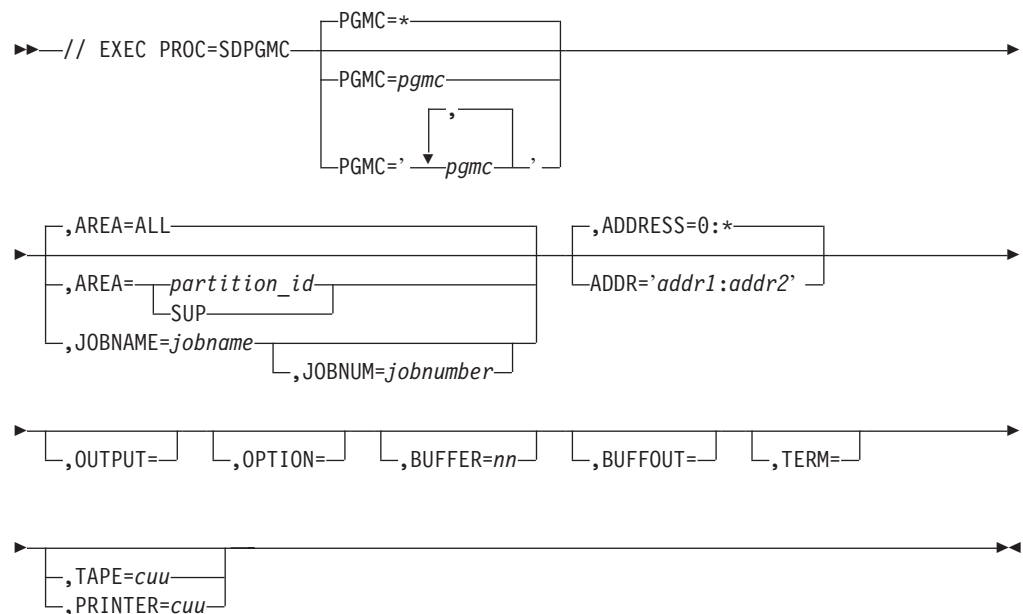
Here are the items of the trace setup shown below:

- Trace type: PGMLOAD

- Traced tasks: all tasks of the BG partition
- Traced storage area: whole VSE/Advanced Functions storage (default)
- Phase whose fetch/load operation is to be traced: MYPHASE
- Additional Output: dump of the storage contents with the address 0 to X'3000', relative to the BG partition start address on occurrence of the PGMLOAD trace event.
- Output destination: tape with device address 280
- Note that the continuation sign has to follow the comma if you use the console for input. If you use SYSIN to enter the procedure statement the continuation sign has to be in column 72.

```
// EXEC PROC=SDLOAD,PHASE=MYPHASE,AREA=BG,OUTPUT='DUMP PART OFF=0:3000',-
      TAPE=280
```

Program Check Trace



See the “Additional Keyword Operands in Trace Procedure Statements” on page 128.

The procedure SDPGMC initializes traces for program check interruptions.

For a description of the trace type and an example of the output, see “PGMCheck Trace (Program Check)” on page 63.

Defaults Set in the Procedure

ADDRESS='0:*' is defined if you omit ADDRESS=.

All program check interruptions are traced if you omit PGMC=.

Statement Example

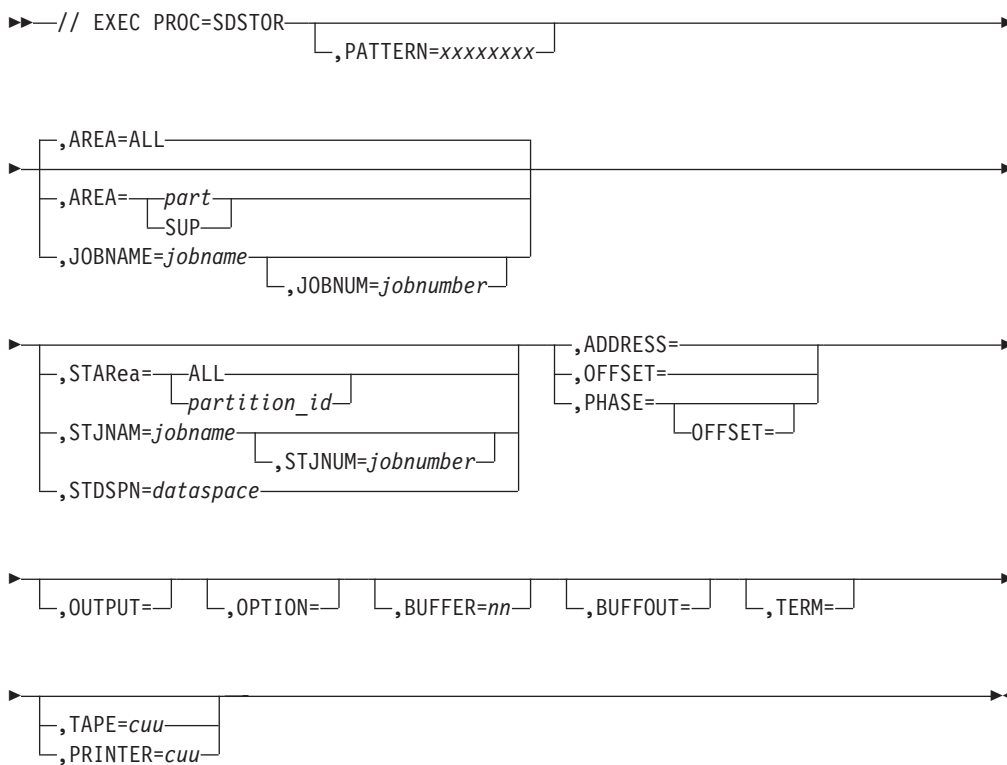
Here are the items of the trace setup shown below:

- Trace type: PGMCHECK
- Traced tasks: all tasks of the BG partition

- Traced storage area: BG partition area (OFF=0:* defined by SDAID defaults)
- Additional Output: dump of the storage contents with the address 0 to X'5000', relative to the BG partition start address on occurrence of the PGMCHECK trace event.
- Output destination: tape with device address 280

```
// EXEC PROC=SDPGMC,AREA=BG,OUTPUT='DUMP PART OFFSET=0:5000',TAPE=280
```

Storage Alteration Trace



See the “Additional Keyword Operands in Trace Procedure Statements” on page 128.

The procedure SDSTOR initializes traces for storage alterations.

You use this trace type as a tool to find those instructions which modify a certain storage area. In most cases you do not know which phase in your system alters this area. For this, define AREA=ALL to watch all tasks operating in your system. The observed storage area is defined via the ADDRESS= keyword.

The optional keyword 'PATTERN=' restricts monitoring to those instructions which change the storage contents into the defined pattern. The specified storage interval which you define with the ADDRESS= keyword should have the same length as the specified pattern (if any).

For a description of the trace type and an example of the output, see “STORAGE Alteration Trace” on page 66.

Statement Example

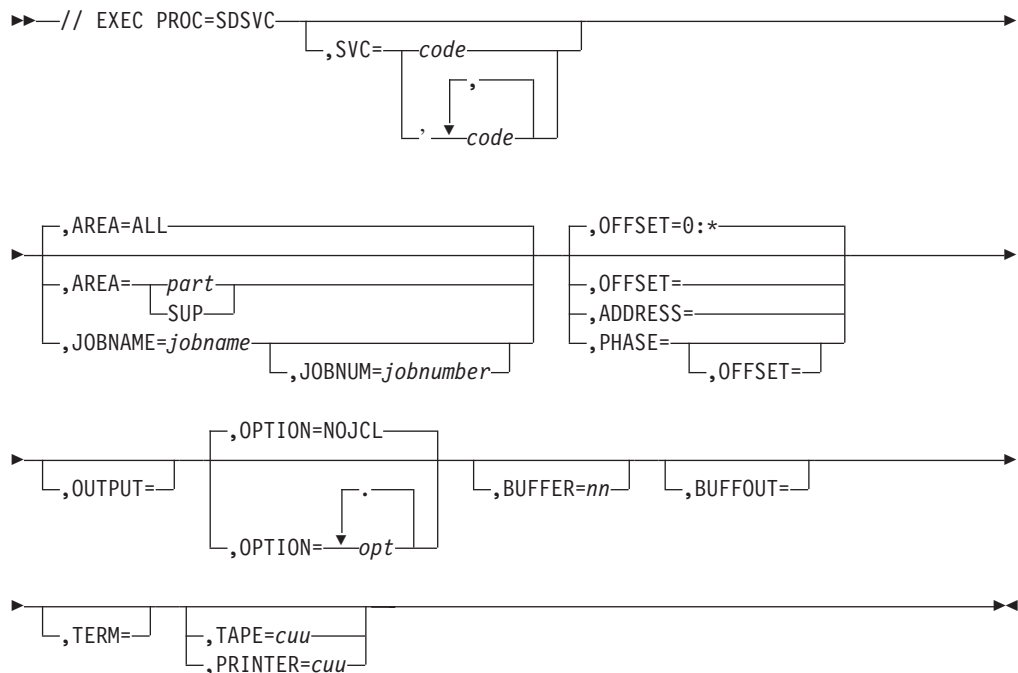
The example finds the instruction which alters a specified storage location into the pattern (FEFE); it puts the system into a wait state when this event occurred.

Here are the items of the trace setup shown below:

- Trace type: STORAGE
- Traced tasks: all tasks in the VSE system
- Address area whose contents alteration is traced: X'074754' up to X'074755'
- Alteration value to be traced: FEFE
- Additional trace definition: put the system into a wait state on the event of the defined storage alteration.
- Output destination: tape with device address 280
- Note that the continuation sign has to follow the comma if you use the console for input. If you use SYSIN to enter the procedure statement, the continuation sign has to be in column 72.

```
// EXEC PROC=SDSTOR,PATTERN=FEFE,AREA=ALL,ADDRESS='74754:74755',-
      OPTION=HALT,TAPE=280
```

SVC Trace



See the “Additional Keyword Operands in Trace Procedure Statements” on page 128. The procedure SDSVC initializes traces which provides event records for all or specified SVC instructions. Define the SVC code in hexadecimal form.

For a description of the trace type and an example of the output, see “SVC Trace (Supervisor Call)” on page 67.

Defaults Set in the Procedure

If you omit SVC=, all SVC instructions are traced.

Statement Example

Here are the items of the trace setup shown below:

- Trace type: SVC
- SVC instruction defined by the SVC code: 3F
- Traced tasks: all tasks of the BG partition
- Traced storage area: BG partition area (SDAID default)
- Output destination: 10K bytes buffer
- Output device for buffer: tape with device address 280
- Event to write the buffer to tape: cancel or EOJ condition in the BG partition

```
// EXEC PROC=SDSVC,AREA=BG,SVC=3F,TAPE=280,BUFFER=10,BUFFOUT=CANCEL
```

Additional Keyword Operands in Trace Procedure Statements

When you initialize a trace using a procedure, the trace type which SDAID actually calls corresponds to the procedure name. The additional trace operands, for example the specification of the output device, correspond to the operands specified or defaulted in the procedure operands.

The additional operands which are specific for the trace initialization via procedures are described in this section. The other additional operands have been described in Chapter 8, "SDAID General Description," on page 53.


The table shows all additional keyword operands in the format accepted in the EXEC PROC statement, a short description for each and a page reference to their detailed description.

Table 11. Additional Keywords, Summary

Operand	Function	Page
ADDRESS	Limit tracing to a certain address range	71
AREA	Limit tracing to a certain system area	70
JOBNAME[JOBNUM]	Limit tracing to a certain VSE/POWER job	70
OFFSET	Limit tracing to a partition or phase area	71
PHASE	Limit tracing to a certain phase	71
OPTION	Define additional trace options	81
OUTPUT	Define additional trace output	73
UNIT	Define the device address	82
BUFFER BU	Define the size of the output buffer	129
BUFFOUT	Define the event to write the buffer	129
TERM	Define the event which terminates the trace	129
PRINTER P	Define the printer device address	129
TAPE T	Define the tape device address	129

Define the Output Device in a Procedure Statement

BUFFER=, PRINTER=, TAPE=Keyword Operands



You define the output destination of the event trace records via the keyword operands BUFFER=nn, TAPE=cuu, or PRINTER=cuu.


Note: The abbreviations BU=nn, T=cuu, or P=cuu may be used.

BUFFER=nn

Defines the size of a wraparound buffer to collect the trace event records.

Note: The definition of a large wraparound buffer may cause a lack of SDAID storage. For information, see “Space Requirements for SDAID Execution” on page 51.

BUFFOUT=Keyword Operand



Via the BUFFOUT= keyword operand you define the condition which forces the write buffer operation.

BUFFOUT=CANCEL

Defines that the contents of the wraparound buffer is to be written to the output device (Printer or Tape) when a cancel or EOJ condition occurs.

BUFFOUT=PGMC

Defines that the contents of the wraparound buffer is to be written to the output device (Printer or Tape) on any program check interruption (except page faults).

Note: If you specify BUFFOUT=CANCEL or BUFFOUT=PGMC, you must also specify the keyword operand AREA=partition-id or JOBNAME=.

BUFFOUT=FULL

Defines that the buffer is to be written to the output device whenever it is full.

BUFFOUT=EXT

Defines that the buffer is to be written to the output device whenever the external interrupt key is pressed.

TERM=Keyword Operand



TERM=CANCEL

defines that tracing is to be terminated as soon as a cancel condition occurs in the traced partition.

TERM=PGMC

defines that tracing is to be terminated as soon as a program check occurs in the traced partition.

TERM=EXT

defines that tracing is to be terminated as soon as the external interrupt key is pressed.

Note: If TERM=CANCEL or TERM=PGMC is specified, AREA=partition-id or JOBNAME has to be specified, too.

Chapter 11. Initialize a Trace in Prompt Input Mode

This chapter describes how you initialize an SDAID trace in prompt input mode. The prompt input mode works only in the attention routine.

Overview

You can set up SDAID traces in prompt mode, which operates in the attention routine.

You invoke the SDAID program in prompt mode by entering the attention routine (AR) command **SDAID** without another specification. Prompt mode is also activated if you process direct input mode commands in the attention routine with at least one prompt mode statement, like the question mark (?).

The trace output device is defined via prompts after you enter the **OUTDEV** command.

SDAID prompts you for the trace type when you enter the **TRACE** command in the AR.

You end the initialization process with the **READY** command.

Once you have initialized the SDAID trace, attention routine commands are used to start the trace execution (**STARTSD**), suspend it (**STOPSD**), and end it (**ENDSD**).

The trace output, an **event record**, is supplied for each occurrence of a traced event, according to your instructions.

You may request the event records to be written to a line printer, onto a magnetic tape, or into a wraparound buffer. The definition of the output device is given via the prompts following the **OUTDEV** command.

The prompts and the possible replies are shown in "Command Input Path Example" on page 135.

How to Initialize an SDAID Trace in Prompt Mode

SDAID trace initialization in prompt mode requires the commands shown in Table 12:

Table 12. Input Command Summary

Command	Description	Page
SDAID	Attention routine command to invoke the SDAID program.	-
OUTDEV	Defines output device for the trace (printer, tape, or buffer).	142
TRACE	Defines the event(s) to be traced. At least one TRACE command is required; up to ten may be submitted.	143

Table 12. Input Command Summary (continued)

Command	Description	Page
READY	Ends input of initialization commands OUTDEV and TRACE.	-

The Various SDAID Commands

SDAID prompts you for the output device of the trace when you enter **OUTDEV**.

One OUTDEV definition can be active in the system at one time. Any newly entered OUTDEV command overwrites the existing one.

Enter **TRACE** to be prompted by SDAID for the type(s) of traces you want. Up to ten 'TRACE' commands may be entered in one session.

You end the trace initialization in the attention routine with the **READY** command. When the READY command has been processed, no further OUTDEV or TRACE command can be entered.

Sample SDAID Trace Initialization

Figure 47 on page 133 shows a typical trace initialization session.

The session starts with the AR command 'SDAID'. With the command 'OUTDEV' the output device is defined and the command 'TRACE' is entered to specify the trace type. The initialization process ends with the READY command.

Prompt Setup via the Attention Routine

```
→ sdaid ␣
4C05I PROCESSING OF 'SDAID' COMMAND SUCCESSFUL.
→ outdev ␣
4C08D SPECIFY OUTPUT DEVICE.+
→ tape ␣
4C08D SPECIFY PHYSICAL ADDRESS OF PRINTER/TAPE.+
→ 281 ␣
4C05I PROCESSING OF 'OUTDEV' COMMAND SUCCESSFUL.
→ trace ␣
4C08D SPECIFY TRACE TYPE.+
→ inst ␣
4C08D SPECIFY OP-CODE(S) OR '*' OR 'BRANCH'.+
→ * ␣
4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+
→ area ␣
4C08D SPECIFY TRACE AREA.+
→ ? ␣
ENTER SYSLOG-ID (LIKE BG OR F1)
FOR TRACING A PARTITION, OR
'SUP' FOR THE SUPERVISOR, OR
'ALL' FOR TRACING ENTIRE SYSTEM
→ SUP ␣
4C08D SPECIFY TYPE OF LIMITS.+
→ add ␣
4C08D SPECIFY ADDRESS RANGE.+
→ 4000:7A00 ␣
4C08D SPECIFY OUTPUT.+
→ ␣
4C08D SPECIFY OPTIONS.+
→ nojcl ␣
4C08D SPECIFY OPTIONS.+
→ ␣
4C05I PROCESSING OF 'TRACE' COMMAND SUCCESSFUL.
1I40I READY.
→ ready ␣
4C05I PROCESSING OF 'READY' COMMAND SUCCESSFUL.
1I40I READY.
→ startsd ␣
4C05I PROCESSING OF 'STARTSD' COMMAND SUCCESSFUL.
1I40I READY.
```

Figure 47. Example: Prompt Mode Trace Initialization. The arrows on the left indicate your input.

Notational Conventions

- SDAID messages (or help information) are shown in uppercase with a message number.
- Responses or commands for you to enter are shown in mixed case. In most responses a short form of the command is also allowed, and this is shown in uppercase. The non-mandatory part of the response is in lowercase. For example, the BRanch trace type specification can be abbreviated in the following way:
BR BRa BRan BRanc BRanc
- The symbol ␣ indicates you are to press the ENTER key (generally after entering any response or command).

How to Use Help and Cancel in Prompt Mode

- Messages for which you can request additional help information are indicated by a plus sign (+) at the end of the message.
- Request additional help by entering a question mark (?).
- You can cancel data entered for the current command by entering two question marks (??).

Figure 48 shows how you can request help information and how the initialization process can be canceled.


```

→ trace □
  4C08D SPECIFY TRACE TYPE.+
→ ? □
  ENTER ONE OF THE FOLLOWING KEYWORDS:
  SVC          PGMCHECK   MONITOR   CANCEL
  INSTR        STORAGE   BRANCH
  PGMLOAD      EXTERNAL   BUFFER
  IO           SSCH       VTAMIO     VTAMBU
→ ?? □
  4D03I COMMAND CANCELED DUE TO USER REQUEST
  
```


Figure 48. Example: Help and Cancel Initialization

How to Read the Following Prompting Mode Syntax Diagrams


The following diagrams use a solid line, or a number of solid lines in parallel, as a specification path. Follow the line of the option that you select for your SDAID execution.




Indicates where you get a prompting message.



Three options, the dotted line indicating a default (in this case from 0 to end); Press ENTER for the default.



Same as above, but END/ENTER indicates no pattern instead of a default.



The comma (,) indicates that the optional operands may be entered more than once and in any sequence. Press ENTER when you have finished the input or none of the options is desired.



On-page connector.



Off-page connector (to a following page).



Command input completed.

Command Input Path Example

This section shows the prompt messages and the possible replies in the sequence of their processing.

Figure 49 shows an example of the trace statement path. You can find the possible input in accordance to the prompt message 'SPECIFY TRACE TYPE' (BR, CA, ..). The example also indicates the prompt message after the reply 'inst' (SPECIFY OP * OR BR).

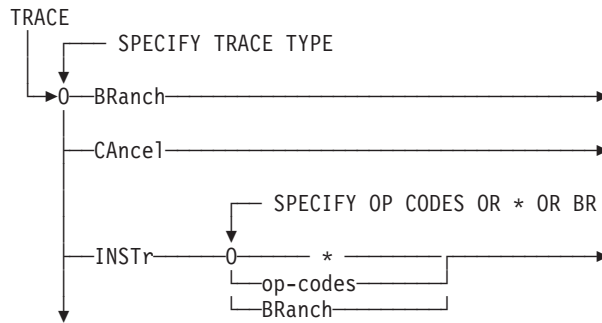


Figure 49. Sample Command Input Path

Command Input Paths

OUTDEV Command Input Path

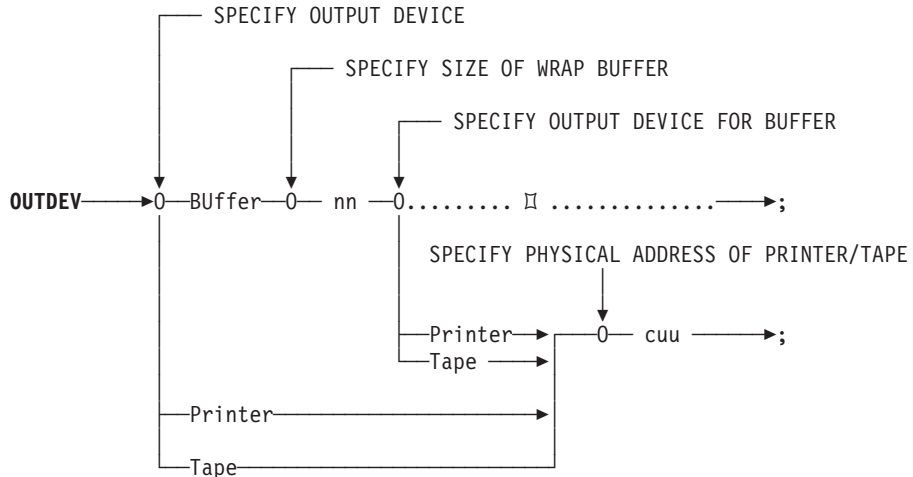


Figure 50. OUTDEV Command: Syntax Diagram

TRACE Command Input Path

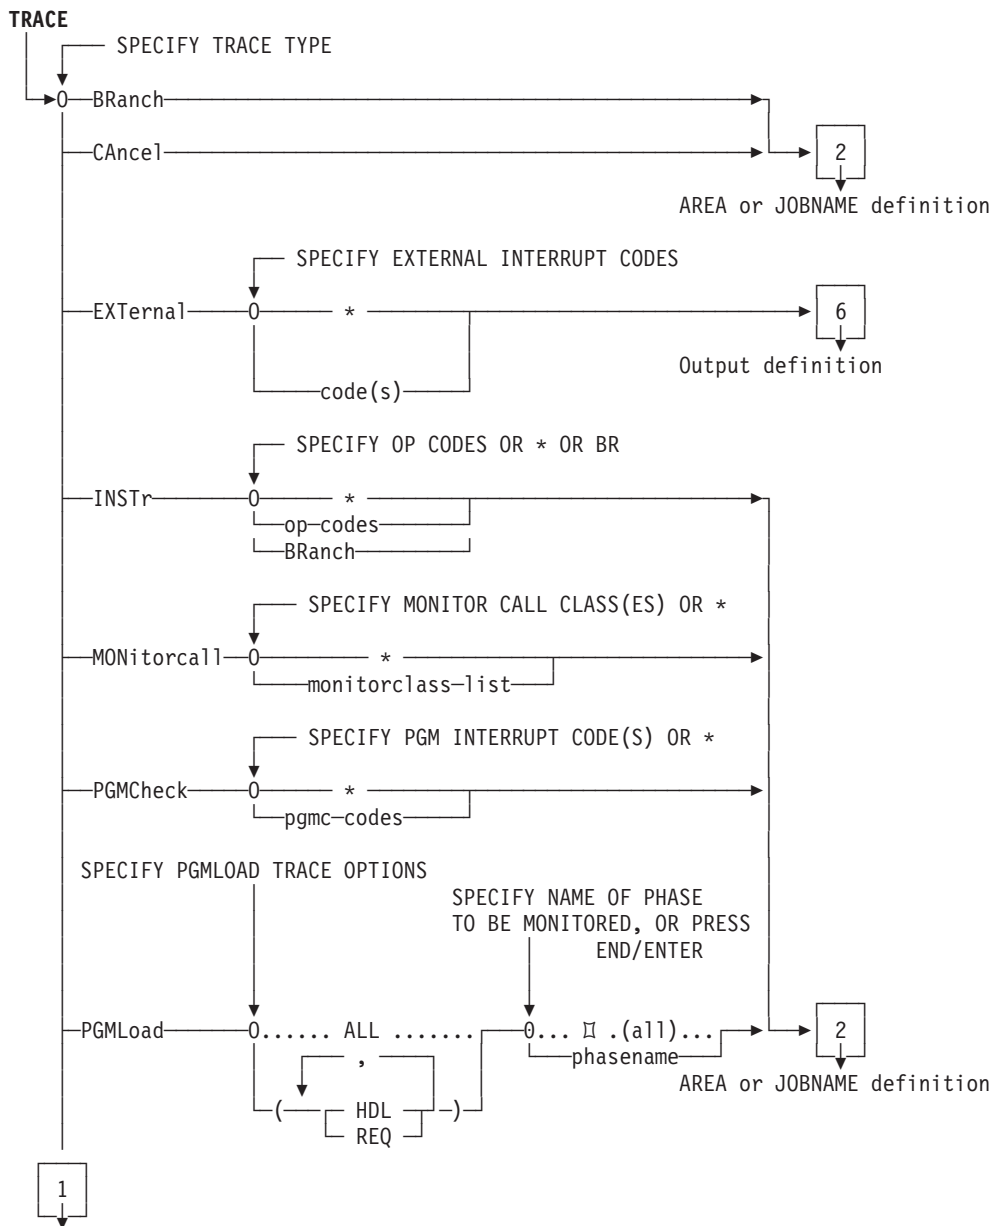


Figure 51. TRACE Command: Syntax Diagram (Part 1 of 7)

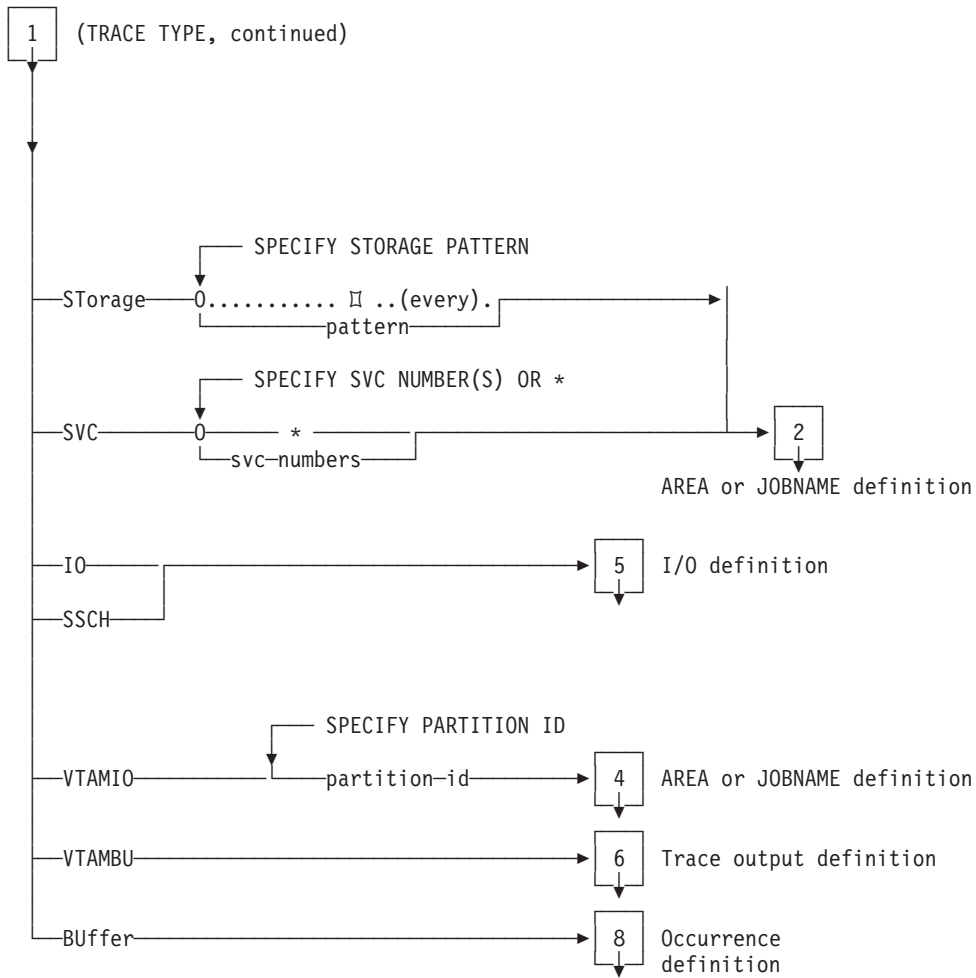


Figure 51. TRACE Command: Syntax Diagram (Part 2 of 7)

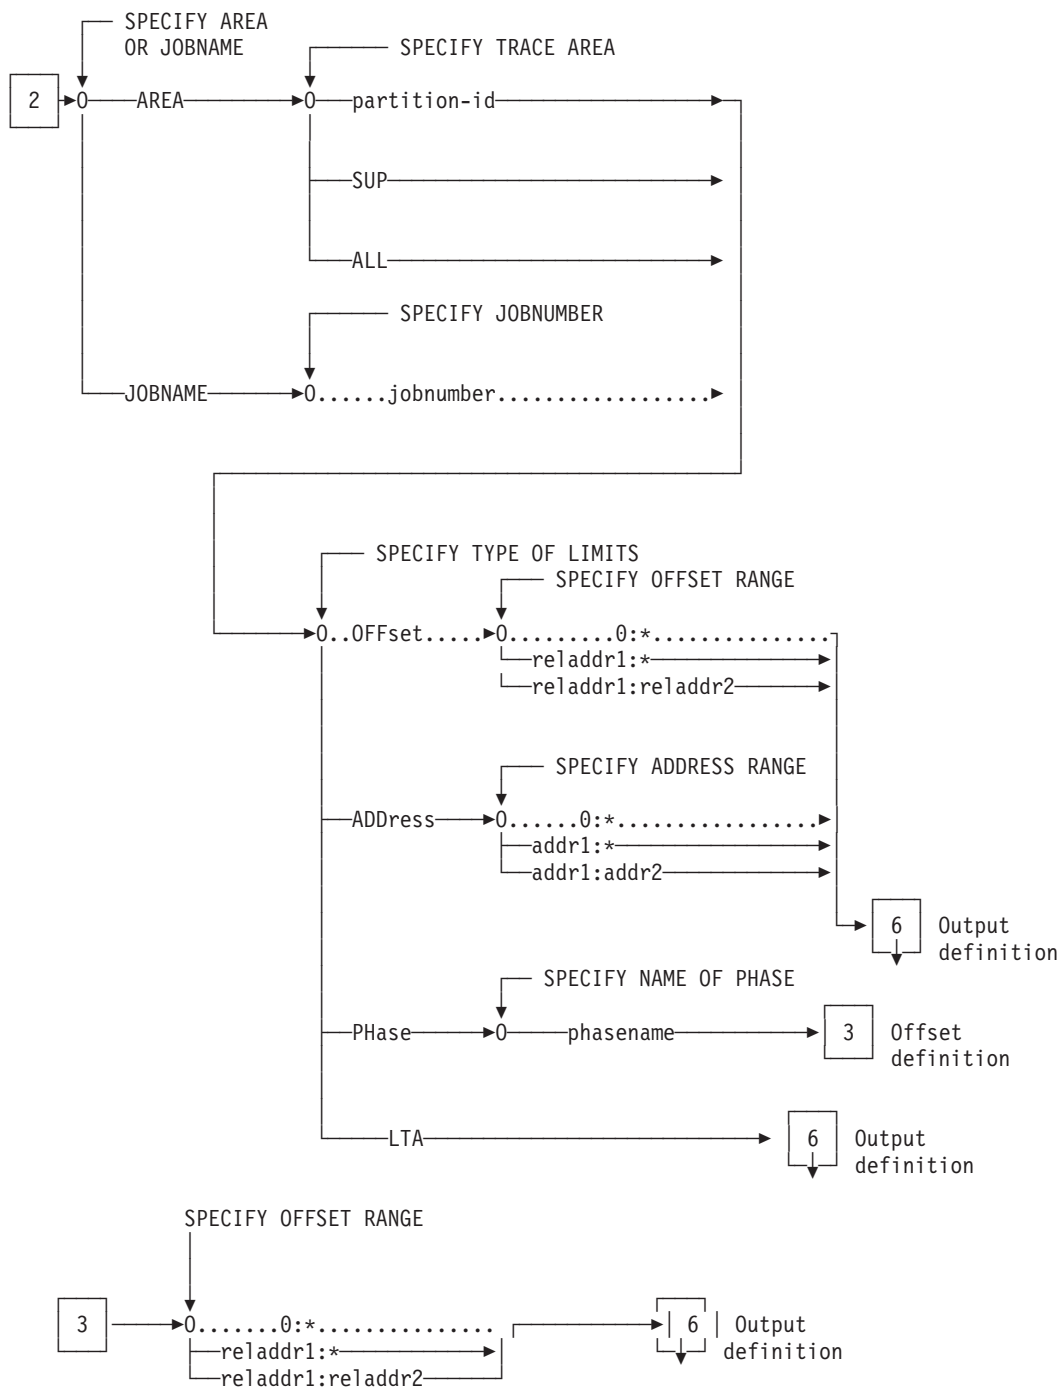


Figure 51. TRACE Command: Syntax Diagram (Part 3 of 7)

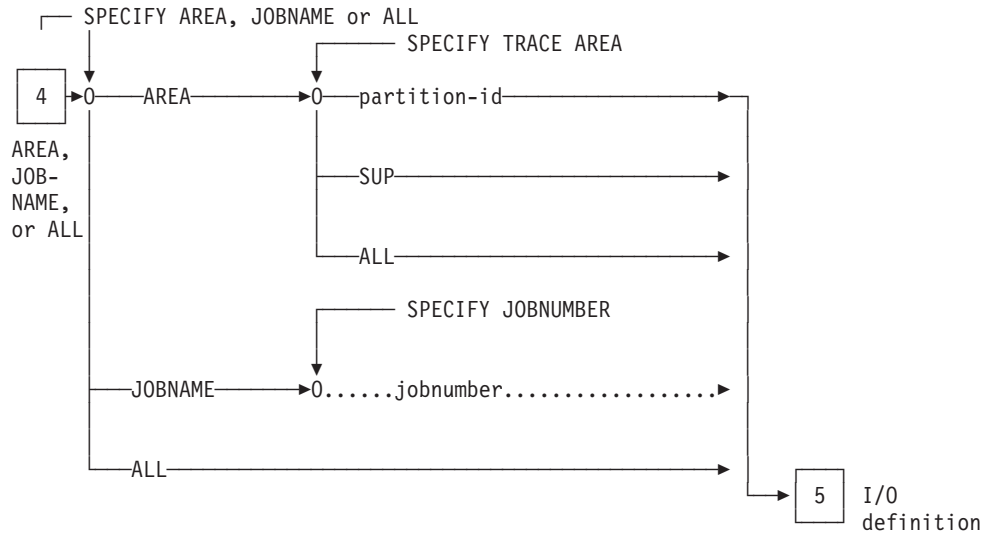


Figure 51. TRACE Command: Syntax Diagram (Part 4 of 7)

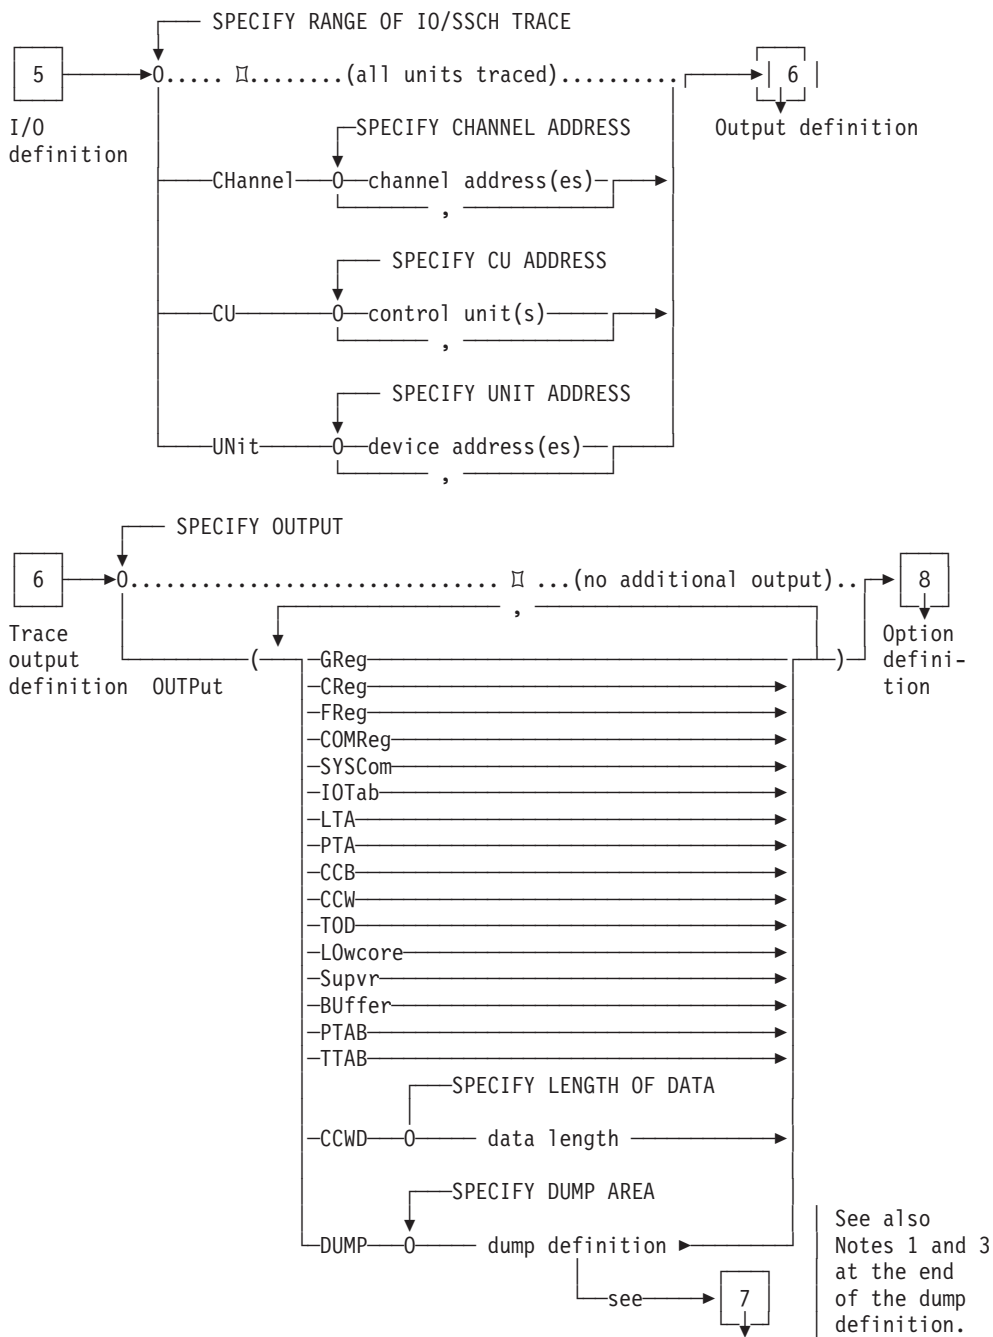
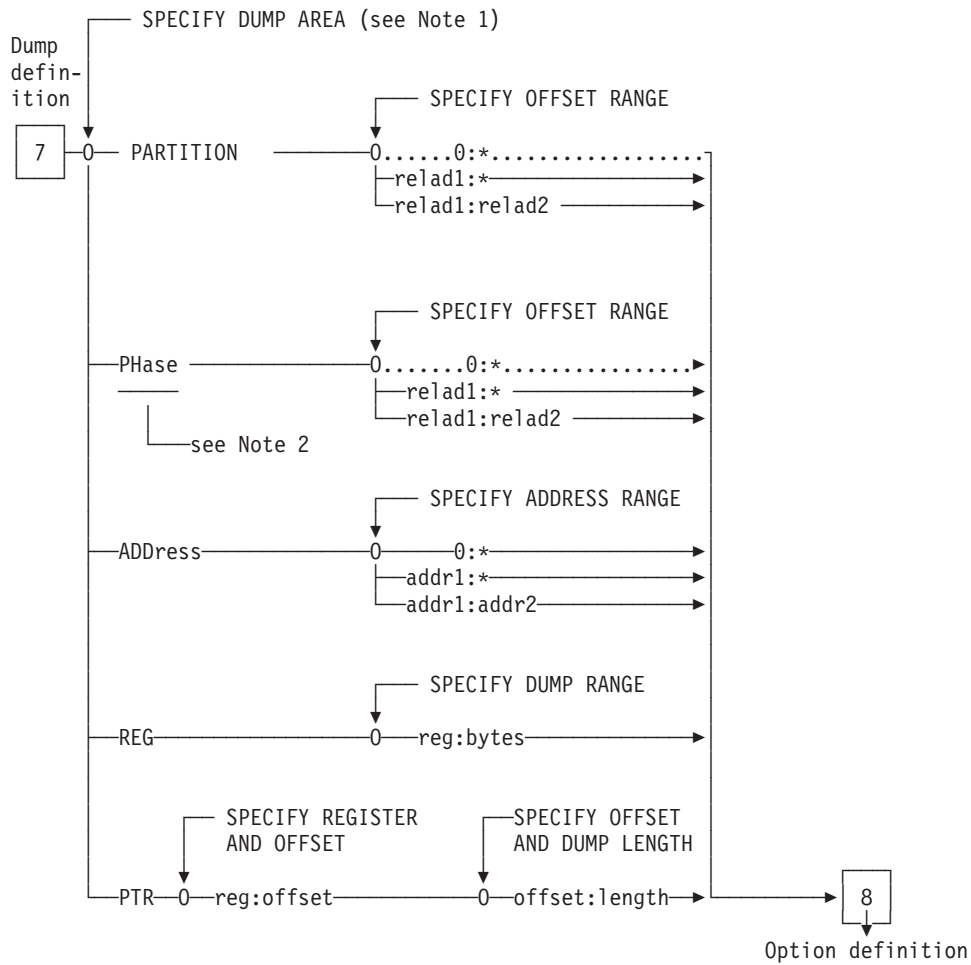


Figure 51. TRACE Command: Syntax Diagram (Part 5 of 7)



Notes:

1. Up to ten different areas may be specified with DUMP.
2. Can be specified only if a phase was previously defined in the area definition of the TRACE.
3. You need not specify the word OUTPUT in prompt mode. SDAID prompts you for the definition of the additional output.

Figure 51. TRACE Command: Syntax Diagram (Part 6 of 7)

Enter the desired size of the buffer in number of blocks of 1K byte.

Possible Buffer Sizes

The possible buffer sizes depend on the output device for the buffer which is defined next.

Table 13. Buffer Sizes

Buffer to Printer or no output device	4K - 256K
Buffer to Tape	4K - 32K

Now, SDAID prompts you as follows:

```
4C08D SPECIFY OUTPUT DEVICE FOR BUFFER.+
```

Respond with either **Printer**, **Tape**, or **END/ENTER**.

Pressing END/ENTER causes no output device being defined.

Specifying the Trace: TRACE Command

Once you enter the command: **TRACE**, SDAID prompts you for the following control information:

Trace-type definition:

The type of event to be traced. See “Defining the Trace Type.”

Area definition:

The range of the trace in storage. See “AREA Definition” on page 156.

I/O definition:

Limits a trace operation to one or more channels, control units, or devices. See “I/O Definition” on page 158.

Output definition:

Additional (optional) trace information that is required to analyze the particular problem. See “Additional Output Definition” on page 160.

Option definition:

An option to:

- Stop system execution when the specified trace event occurs.
- Discontinue tracing when the specified trace event occurs.
- Avoid tracing of Job Control phases.
- Discontinue tracing when a defined number of events has been exceeded.
- Include supervisor routines into a partition trace.

See “Option Definition” on page 160.

You will find sample event records and a description of most of the trace types under “Summary of TRACE Types” on page 56.

Defining the Trace Type

This section, and the descriptions of the various trace types that follow, show the promptings and the possible replies when defining an SDAID trace.

Detailed information about the various trace types is given under “Summary of TRACE Types” on page 56.

You start the definition of your trace with the trace command in the attention routine.

```
trace
4C08D SPECIFY TRACE TYPE.+
```

Respond to the prompting message with any of the available trace types. For example if you want to initialize a branch trace, the response would look like this:

```
trace
4C08D SPECIFY TRACE TYPE.+
branch ← Your response
```

The SDAID then prompts you for additional information.

Please find a summary of the SDAID trace types in Table 14.

Summary of Trace Types

Table 14 gives the following information:

- The trace types shown in the format they can be entered. Note, that the uppercase letters indicate the shortest possible abbreviation.
- A short description of the trace type. All trace types are described in more detail under “Summary of TRACE Types” on page 56.
- A page reference to the format description of the trace type.

Table 14. Trace Type Summary

Trace Type	Provides a Trace of:	Page
BRanch	Successfully executed branch instructions	145
BUffer	The trace buffer when it is full	145
CAncel	Program (main task) cancel or EOJ	145
EXTernal	External interrupts	145
GETVIS	Getvis / Freevis requests	146
INSTruction	Selected or all instruction(s) execution	147
IO	I/O interrupts	148
LOCK	Lock / Unlock requests of resources	148
MONitorcall	MC instructions	150
PGMCheck	Program checks	150
PGMLoad	Phase load requests, or actual load	151
SSCH	Start Subchannel instructions	152
STorage	Storage alterations	153
SVC	Executed supervisor calls	153
VTAMBU	Usage of VTAM buffers	154
VTAMIO	VTAM I/O operations	154

Table 14. Trace Type Summary (continued)

Trace Type	Provides a Trace of:	Page
XPCC	XPCC communication actions	154

BRanch Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
branch ← Your response
```

SDAID then prompts you for the definition of the trace area or the job name:

```
4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+
```

See “AREA Definition” on page 156 and “JOBNAME Definition” on page 157.

BUffer Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
buffer ← Your response
```

The SDAID prompts you for the OCCurrence definition as next. For the format of these definitions. see “Option Definition” on page 160.

CAnceL Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
cancel ← Your response
```

SDAID then prompts you for the definition of the trace area or the job name:

```
4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+
```

See “AREA Definition” on page 156 and “JOBNAME Definition” on page 157.

EXternal (External Interrupt) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
external ← Your response
```

The SDAID prompts you for additional control information until you press END/ENTER, as follows:

4C08D SPECIFY TYPE OF EXTERNAL INTERRUPT OR '*.*+

Your response may be one to 8 of the following:

*	To trace all types of external interrupts.
0040	To trace only key interrupts.
1003	To trace TOD-clock sync check.
1004	To trace clock comparator.
1005	To trace CPU timer.
1200	To trace malfunction alert.
1201	To trace emergency signal.
1202	To trace external call.
2401	To trace service signal.
2402	To trace logical device.* z/VM CP
2603	To trace PFAULT handshaking.* z/VM CP
4000	To trace IUCV, APPC.* z/VM CP
4001	To trace VMCF.* z/VM CP
END/ENTER	To continue.

The SDAID now prompts you for the definition of the OUTPUT.

GETVis (Getvis / Freevis Request) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
getv      ←————— Your response
```

The SDAID prompts you for additional control information, as follows:

4C08D SPECIFY GETVIS/FREEVIS STORAGE REGION.+

Your response may be one of the following:

PARTition

Writes an event record if the requested or released space is within the partition Getvis area.

SPACE Writes an event record if the requested or released space is within the dynamic partition Getvis area.

SVA Writes an event record if the requested or released space is within the SVA.

SDAID then prompts you for a specific subpool name.

```
4C08D SPECIFY SUBPOOL NAME.+
```

Your response may be one of the following:

Press ENTER

Writes an event record for all subpools.

subpool_name

Writes an event record for the specific subpool only.

See “GETVIS / FREEVIS Trace” on page 58 for details of a subpool name format.

SDAID then prompts you for the LOCation of a Getvis / Freevis request:

```
4C08D SPECIFY LOCATION FOR GETVIS/FREEVIS REQUESTS.+
```

Your response may be one of the following:

Press ENTER

Writes an event record for all Getvis/Freevis requests.

BELOW

Writes an event record for requests within the 24-bit Getvis area.

ANY Writes an event record for requests within the 24-bit and 31-bit Getvis area.

SDAID then prompts you for the definition of the trace area or the job name:

```
4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+
```

See “AREA Definition” on page 156 and “JOBNAME Definition” on page 157.

INSTRUCTION (Instruction Execution) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
instr ←————— Your response
```

The SDAID prompts you for additional control information, as follows:

```
4C08D SPECIFY OP-CODE(S) OR *.+
```

Your response may be one of the following:

op code(s)

(one to eight) Entered as either one-byte or two-byte hexadecimal values. If you specify more than one operation code, separate them by one or more blanks or by a comma (with or without blanks).

asterisk (*)

Defines all op codes.

You can also specify the following:

BRanch

Defines that all branch instructions have to be traced regardless whether the branch has been taken or not.

Sample responses:

```
d2
18 41,58 40, 50 9608
*
BRanch
```

SDAID then prompts you for the definition of the trace area or the job name:

```
4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+
```

See "AREA Definition" on page 156 and "JOBNAME Definition" on page 157.

IO (I/O Interrupt) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
io ← Your response
```

SDAID then prompts you for the definition of the trace area or the job name:

```
4C08D SPECIFY ONE OF THE KEYWORDS AREA, JOBNAME OR ALL.+
```

See "AREA Definition" on page 156 and "JOBNAME Definition" on page 157. ALL means all tasks of the system.

SDAID then prompts you for the definition of the specific I/O channel(s), control unit(s), or unit(s), as described under "I/O Definition" on page 158.

LOCK (Lock / Unlock of Resources) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
getv ← Your response
```

The SDAID prompts you for a specific resource name:

```
4C08D SPECIFY RESOURCE NAME OR PRESS END/ENTER.+
```

Your response may be one of the following:

Press ENTER

Writes an event record for all resources.

resource_name

Writes an event record for the specific resource only.

See "LOCK / UNLOCK Trace" on page 61 for details of the resource-name formats.

SDAID then prompts you for the type (LOCK or UNLOCK):

4C08D SPECIFY TYPE OF REQUEST OR PRESS END/ENTER.+

Your response may be one of the following:

Press ENTER

Writes an event record for all types of request.

Lock Writes an event record for resource locking only.

Unlock

Writes an event record for resource unlocking only.

SDAID then prompts you for the scope (INTERNAL or EXTERNAL):

4C08D SPECIFY THE SCOPE OR PRESS END/ENTER.+

Your response may be one of the following:

Press ENTER

Writes an event record for all scopes.

INTernal

Writes an event record for Internal Locks or Unlocks.

EXTernal

Writes an event record for External Locks or Unlocks.

SDAID then prompts you for a volume ID:

4C08D SPECIFY VOLUME-ID OR PRESS END/ENTER.+

Your response may be one of the following:

Press ENTER

Writes an event record for all volume IDs.

volume_id

Writes an event record for the specified volume ID only.

SDAID then prompts you for a return code(s):

4C08D SPECIFY RETURN CODE(S) OR '*' .+

Your response may be one of the following:

* Writes an event record for all return codes.

return_code(s)

Writes an event record for the specified return code(s) only.

See "LOCK / UNLOCK Trace" on page 61 for details of how to specify return code(s).

SDAID then prompts you for the trace area or the job name:

4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+

See “AREA Definition” on page 156 and “JOBNAME Definition” on page 157.

MONitorcall Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
monitorcall ← Your response
```

SDAID prompts you for additional control information as follows:

4C08D SPECIFY MONITOR CALL CLASS(ES) OR *.+

Your response may be one of the following:

monitor classes

defines the MC instructions to be traced by one or up to eight monitor classes.

Monitor classes must be specified as one-digit hexadecimal values. If you specify two or more classes, separate them by one or more blanks, or by a comma with or without blanks.

You may specify any valid monitor class; however, SDAID ignores a specification of class 2.

asterisk (*)

defines all classes except class 2.

Sample responses:

```
3 5, 8,c d e f
a
*
```

SDAID then prompts you for the definition of the trace area or the job name:

4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+

See “AREA Definition” on page 156 and “JOBNAME Definition” on page 157.

PGMCheck (Program Check) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
pgmcheck ← Your response
```

SDAID prompts you for additional control information as follows:

4C08D SPECIFY PROGRAM INTERRUPT CODE(S) OR '*' .+

Your response may be one of the following:

Program interrupt codes

(one to 16) must be specified in hexadecimal notation, leading zeros may be omitted. An asterisk (*) indicates all program check interruption codes - except those page or segment translation exceptions which are caused by the temporary absence of a storage page. The specification 10 11 traces all page or segment translation exceptions.

If you specify more than one program interrupt code, separate them by one or more blanks, or by a comma with one or more blanks.

Sample specifications: 1 13,05, 10, 0A
9
*

SDAID then prompts you for the definition of the trace area or the job name:

4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+

See "AREA Definition" on page 156 and "JOBNAME Definition" on page 157.

PGMLoad (Program Load) Trace

trace 4C08D SPECIFY TRACE TYPE.+ pgmload ← Your response

SDAID prompts you for additional control information as follows:

4C08D SPECIFY PGMLoad TRACE OPTIONS.+

Your response may be one of the following:

Press ENTER

Writes an event record for all program load events (phase load, fetch request, or actual phase-load operation) within the specified trace range.

req Writes an event record each time loading/fetching a phase is requested (see Notes 1 and 3 below).

hdl Writes an event record each time a phase load/fetch request is handled; that is, when a requested phase is actually loaded into storage for execution (see Notes 1 and 3 below).

all Writes an event record each time a phase load/fetch request occurs, and also each time a phase is actually loaded into storage for execution. This is the default (see Notes 1 and 2 below).

Notes:

1. When you have entered a response to the above prompting message, SDAID repeats the prompting message until you respond by pressing ENTER.
2. If you want all program-load events to be traced, respond by pressing ENTER when SDAID displays the above prompting message *for the first time*.
3. If you want only one phase to be traced, submit the name of this phase after specifying HDL, REQ or ALL.

SDAID prompts you for additional control information as follows:

4C08D SPECIFY NAME OF PHASE TO BE MONITORED, OR PRESS ENTER

Your response may be one of the following:

Press ENTER

Defines all phases to be traced.

phase name

Defines the phase to be traced.

Figure 52 is an example of a prompting sequence for a program load trace request.

SDAID then prompts you for the definition of the trace area or the job name:

4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+

See "AREA Definition" on page 156 and "JOBNAME Definition" on page 157.

```
trace ⅃
4C08D SPECIFY TRACE TYPE +
pgmload ⅃ ←————— your response
4C08D SPECIFY PGMLOAD TRACE OPTIONS.+
hdl ⅃ ←————— your response
4C08D SPECIFY PGMLOAD TRACE OPTIONS.+
⅃ ←————— your response (A)
4C08D SPECIFY NAME OF PHASE TO BE MONITORED OR PRESS END/ENTER.+
myphase ⅃ ←————— your response (B)
```

(A) Indicates that no further control information of this type is to be entered.

(B) Restricts the trace to the loading of the named phase.

Figure 52. Prompting for a PGMLoad Request

Start Subchannel Instruction Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
ssch ←————— Your response
```

SDAID then prompts you for the definition of the trace area or the job name:

4C08D SPECIFY ONE OF THE KEYWORDS AREA, JOBNAME OR ALL.+

See "AREA Definition" on page 156 and "JOBNAME Definition" on page 157. ALL means all tasks of the system.

SDAID further prompts you for the definition of the specific I/O channel(s), control unit(s), or unit(s), as described under "I/O Definition" on page 158.

Storage Alteration Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
storage ← Your response
```

SDAID then prompts you for additional control information as follows:

```
4C08D SPECIFY STORAGE PATTERN.+
```

Your response to this prompting message may be either of the following:

Press ENTER

Requests an event record to be written whenever storage within the trace range *is altered*.

hexvalue

Requests an event record to be written whenever storage within the trace range is *set to the specified value* (any hexadecimal value of up to four bytes). If you specify an odd number of digits, a zero is inserted to the left of the first specified hexadecimal digit.

Note: This option traces only program-altered storage, not that altered by I/O operations.

SDAID then prompts you for the definition of the trace area or the job name:

```
4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+
```

See “AREA Definition” on page 156 and “JOBNAME Definition” on page 157.

Define ALL if you want to get all tasks of your system watched.

SVC (Supervisor Call) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
svc ← Your response
```

SDAID then prompts you for additional control information as follows:

```
4C08D SPECIFY SVC NUMBER(S) OR '*' .+
```

Your response to this prompting message may be either of the following:

one to 16 SVCs

Specify the SVC number in hexadecimal notation. If you specify two or more SVC numbers, they must be separated by one or more blanks, or by a comma with or without blanks.

asterisk (*):

Defines all SVC instructions to be traced.

Sample SVC specifications:

```
02 9,A 26
25
*
```

SDAID then prompts you for the definition of the trace area or the job name:

```
_____
4C08D SPECIFY ONE OF THE KEYWORDS AREA OR JOBNAME.+
_____
```

See “AREA Definition” on page 156 and “JOBNAME Definition” on page 157.

VTAMBU (VTAM Buffer) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
vtambu ← Your response
```

The SDAID prompts you for the OUTPUT definition. See “Additional Output Definition” on page 160.

VTAMIO (VTAM I/O) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
vtamio ← Your response
```

SDAID then prompts you for the definition of the partition-ID as shown below:

```
_____
4C08D SPECIFY PARTITION ID.+
_____
```

Specify the partition where VTAM is running, F3 for example.

SDAID then prompts you next for the definition of the specific I/O channel(s), control unit(s), or unit(s), as described under: “I/O Definition” on page 158.

XPCC (Partition Communication) Trace

```
trace
4C08D SPECIFY TRACE TYPE.+
xpcc ← Your response
```

The SDAID prompts you for a specific application name:

```
_____
4C08D SPECIFY APPLICATION NAME OR '*'.+
_____
```

Your response may be one of the following:

* Writes an event record for all applications.

application_name

Writes an event record for the specific application only.

SDAID then prompts you for a specific to_application name:

4C08D SPECIFY TO_APPLICATION NAME OR '*' .+

Your response may be one of the following:

- * Writes an event record for all to_applications.

to_application_name

Writes an event record for the specific to_application only.

SDAID then prompts you for an xpc function:

4C08D SPECIFY XPC FUNCTION.+

Your response may be one of the following:

- * Writes an event record for all xpc applications.

xpc_function

Writes an event record for the specified xpc_function only.

SDAID then prompts you for the direction of tracing:

4C08D SPECIFY DIRECTION OF TRACEING.+

Your response may be one of the following:

Press ENTER

Writes an event record for incoming and outgoing xpc requests.

in Writes an event record for incoming xpc requests.

out Writes an event record for outgoing xpc requests.

both Writes an event record for incoming and outgoing xpc requests.

SDAID then prompts you for a specific return code:

4C08D SPECIFY DEDICATED RETURN CODE.+

Your response may be one of the following:

- * Writes an event record for all return codes.

return_code

Writes an event record for the specified return code only.

SDAID then prompts you for a comparator (SUSR / IJBXSUSR):

4C08D SPECIFY COMPARATOR TO COMPARE SUSR AGAINST IJBXSUSR.+

Your response may be one of the following:

nocomp

No comparison is required.

EQ | NE | GT | GE | LT | LE

If your response is *not* **nocomp**, SDAID prompts you for a hex value which will be compared against the content of IJBXSUSR.

```
4C08D SPECIFY SUSR- HEXVALUE.+
```

See “XPCC Trace” on page 69 for details of how to prepare a compare field.

SDAID then prompts you for a comparator (RUSR / IJBXRUSR):

```
4C08D SPECIFY COMPARATOR TO COMPARE RUSR AGAINST IJBXRUSR.+
```

Your response may be one of the following:

nocomp

No comparison is required.

EQ | NE | GT | GE | LT | LE

If your response is *not* **nocomp**, SDAID prompts you for a hex value which will be compared against the content of IJBXRUSR.

```
4C08D SPECIFY RUSR-HEXVALUE.+
```

See “XPCC Trace” on page 69 for details of how to prepare a compare field.

SDAID then prompts you for an XPCC from_area:

```
4C08D SPECIFY XPAREA(FROM-PARTITION).+
```

Your response may be one of the following:

Press ENTER

Writes an event record for all from_areas.

syslog_id

Writes an event record for the specified from_area only.

SDAID then prompts you for an XPCC to_area:

```
4C08D SPECIFY XPTOAREA(TO-PARTITION).+
```

Your response may be one of the following:

Press ENTER

Writes an event record for all to_areas.

syslog_id

Writes an event record for the specified to_area only.

AREA Definition

This section describes the SDAID promptings and gives some examples of the possible replies to the AREa Definition. More detailed information about the AREa definition and the corresponding storage region definitions is given in the following sections:

- “Defining the Area to be Traced: AREA Definition” on page 70.
- “Defining the Storage to be Traced: OFFset, ADDRess, PHase, LTA” on page 71.

SDAID prompts you for the required area definition by displaying the message:

```
4C08D SPECIFY TRACE AREA.+
```

Enter one of the following responses:

```
partition-ID
SUP
ALL
```

JOBNAME Definition

The JOBNAME (and JOBNUMBER) definition allows you to trace a VSE/POWER job in a dynamic or static partition. You can use either the AREA or the JOBNAME definition, but not both. For details about the JOBNAME definition, refer to “Defining the Job to be Traced: JOBNAME Definition” on page 70.

The possible prompts are the same as for the AREA definition.

SDAID prompts you for the optional JOBNUMBER definition by displaying the message:

```
4C08D SPECIFY JOBNUMBER.+
```

Enter the job number of the VSE/POWER job to be traced.

Prompts after AREA and JOBNAME Definitions

SDAID prompts you for the definition of the storage area to be traced as follows:

```
4C08D SPECIFY TYPE OF LIMITS.+
```

Your response to this prompting message is one of the following:

```
END/ENTER
OFFset
PHase
ADDRess
LTA
```

Press ENTER

To trace the requested events within the storage occupied by the defined partition.

OFFset

SDAID prompts you for the actual offset values:

```
4C08D SPECIFY OFFSET RANGE.+
```

Your response to this message, a pair of offsets, is discussed below, under “PHase”.

Note that OFFset does not apply to ARea=All.

PHase SDAID prompts for the phase name as follows:

4C08D SPECIFY NAME OF PHASE

Then SDAID prompts you for further limitation of the trace range:

4C08D SPECIFY OFFSET RANGE.+

Your response to this message (for either partition or phase offset) is one of the following:

ENTER
reladdr1:reladdr2
reladdr1:*

Press ENTER

to trace the events in the entire area allocated to the specified partition, supervisor or phase.

reladdr1:reladdr2

to define an address range in hexadecimal notation.

reladdr1:*

to define an address range starting with 'reladdr1' up to the end of the specified partition, supervisor or phase.

ADDRESS

SDAID prompts you for the actual address range:

4C08D SPECIFY ADDRESS RANGE.+

Your response to this message is either of the following:

END/ENTER
addr1:addr2
addr1:*

Press ENTER

to trace the events of the tasks with the defined partition-id without address limitation.

addr1:addr2

to define a certain address area within the partition or supervisor.

addr1:*

to define an address area starting with 'addr1' up to the end of the partition or supervisor.

LTA SDAID traces the events of the specified partition or supervisor which occur in the Logical Transient Area.

I/O Definition

The I/O definition limits the range of an **IO**, **SSCH**, or **VTAMIO** trace to one or more devices, to one or more control units, or to one or more channels. This section describes the SDAID promptings. Detailed information about the I/O definitions is provided under "Defining the Traced I/O Devices" on page 82.

SDAID prompts you for the definition as follows:

4C08D SPECIFY KEYWORD UNIT OR CU OR CHANNEL.+

Your response to this prompting message may be one of the following: (detailed descriptions follow)

ENTER
UNit
CU
CHannel

Press ENTER

To define all I/O devices.

UNit SDAID prompts you for the hexadecimal specification of up to 8 unit addresses as follows:

4C08D SPECIFY UNIT ADDRESS(ES).+

If you specify more than one address, separate them by one or more blanks, or by a comma with one or more blanks or without a blank.

If you specify a 1-digit device address, SDAID assumes channel 0 and control unit 0; for a 2-digit device address, SDAID assumes channel 0.

Sample device-address list specifications:

003, e 181 281
282
e (same as 00e)
0e (same as 00e)

CU SDAID prompts you for the hexadecimal definition of up to 16 control unit addresses as follows:

4C08D SPECIFY CONTROL UNIT ADDRESS(ES).+

If you specify more than one address, separate them by one or more blanks, or by a comma with one or more blanks or without a blank.

Sample control-unit address list specifications:

1, 2a 3f
1c
2 (same as 02)

CHannel

The program prompts you for one or up to 16 channel addresses as follows:

4C08D SPECIFY CHANNEL ADDRESS(ES).+

If you specify more than one address, separate them by one or more blanks, or by a comma with one or more blanks or without a blank.

Sample channel address specifications:

1
0 2, 3

Additional Output Definition

This section gives information on the various responses to the SDAID promptings. If you want more detailed information on the output definitions, refer to “Defining Additional Trace Output: OUTPut Definition” on page 73 and Table 4 on page 73.

SDAID prompts you to specify additional output in the following way:

```
4C08D SPECIFY OUTPUT.+
```

You may respond with the following output definitions:

BUffer	FReg	PTAB
CCB	GReg	SUPvr
CCW	IOTab	SYSCom
CCWD	LOCKTE	TOD
COMReg	LOWcore	TTAB
CReg	LTA	XPCCB
DUMP	PTA	XPDATABU

For each prompt, you may specify one definition. Prompting continues until you press ENTER without a definition. This ends the output definition.

Option Definition

This section gives information on the various responses to the SDAID promptings. If you want more detailed information on the option definitions, refer to “Defining the Trace Options: OPTion Definition” on page 81.

SDAID prompts you for an option definition as follows:

```
4C08D SPECIFY OPTIONS.+
```

You respond with one of the following:

- Halt**
- NOJCL** Specification
- NOSource**
- NOTarget**
- OCcurrence** Definition
- SUPervisor**
- Termination** Specification

If you define OCcurrence, SDAID prompts you as follows:

```
4C08D SPECIFY OCCURRENCE RANGE.+
```

Respond with:

```
ENTER  
value1:value2
```

Press ENTER

To indicate that you want all occurrences of the specified event to be traced (same as if you defined 1:*)

value1:value2

To limit tracing (value2 must be higher than or equal to value1). See the examples below.

Sample occurrence definitions:

1:1 trace only the first occurrence
of the event
1:* trace all occurrences of the
event (this is the default value)
5:12 trace selected occurrences
(5 to 12) of the specified event

Chapter 12. Start/Stop and End the Trace

This chapter describes how you can start, stop, or terminate an initialized SDAID trace, and how to control the trace under exceptional conditions.

The Required Commands

STARTSD/STOPSD Commands: Starting and Stopping

Once you have entered the READY command which ends the initialization process, you can activate the trace at once or later. To start or restart the trace operation, enter the command

▶▶—STARTSD—◀◀

without any operand.

Note: If the trace was stopped by an event itself (TERMinate specified with the TRACE command), the trace operation can be restarted by issuing the STOPSD command followed by the STARTSD command.

The STARTSD command is rejected if the interactive trace program is active for any partition.

To interrupt the trace operation with the restart capability retained, enter the command

▶▶—STOPSD—◀◀

without any operand.

Note: When a tape is defined as output device, every STOPSD or ENDS command writes a tapemark on the tape if there was any trace event. If, for example, you specify three times STARTSD/STOPSD within an SDAID session, you get three trace files on your trace output tape. However, if there was no trace event since the last STARTSD command, the tape remains unchanged.

ENDSD Command: Ending Execution

You can end the SDAID session by issuing the command:

▶▶—ENDSD—◀◀

without any operands. The ENDS command releases all resources that were used by the program during the session, including the storage space that was occupied by SDAID and closes the trace output device. You may enter this command at any time during a session.

Attention Routine Command Example

The example in Figure 53 shows how an initialized SDAID trace is started, interrupted and ended. After the ENDSID command has been processed all of the initialized trace information is released.

```
•
startsd □
4C05I PROCESSING OF 'STARTSD' COMMAND SUCCESSFUL
•
•
•
stopspd □
4C05I PROCESSING OF 'STOPSD' COMMAND SUCCESSFUL
•
•
•
startsd □
4C05I PROCESSING OF 'STARTSD' COMMAND SUCCESSFUL
•
•
•
endsid □
4C05I PROCESSING OF 'ENDSID' COMMAND SUCCESSFUL
```

Figure 53. Attention Routine Commands to Start, Stop and End the Trace

How to Control the Trace under Exceptional Conditions

The start and stop procedures described above can be used only when the attention routine is available. When you try to start or stop a trace, the attention routine may be unavailable because of the problem you are trying to identify, or because SDAID is in a wait state.

This section tells you how to control traces:

- When the system is in an unintended loop;
- When a trace is running and the attention routine is not available;
- When the system is in a wait state.

Tracing an Unintended Loop

Perform the following steps to use the SDAID branch, instruction or storage-alteration traces to gather information about an unintended loop:

1. Initialize one of the trace types mentioned above in the normal way.
2. Start the trace with the STARTSD command.
3. Display the contents of control register 9 with the control processors alter/display feature.
4. Notice the contents of this control register for later use.
5. Set bits 0 through 3 to zeros with the alter display feature. This stops the trace.
6. Recreate the loop condition by submitting the same job mix that existed when the particular loop occurred the first time.
7. When the loop appears again, restart SDAID operation by setting those bits of control register 9 to a value of 1 which you have set to zeros before.

Bit	Effect if set to 1
0	Successful branches are traced.
1	Instruction executions are traced.
2	Storage alterations are traced.

(See “Hardware Alter/Display” on page 236 for information on how to use the Alter/Display feature.)

Control register A contains the start and control register B the end address of the trace. You may change this address range by varying the addresses stored in control registers A and B.

To resume operation after SDAID has collected sufficient information about the loop, and if you cannot exit from the loop, re-IPL VSE.

Terminating SDAID Program Without the Attention Routine

It may happen that you can no longer request the attention routine to gain control of your processor. At that point, SDAID operation cannot be stopped as usual by entering the command STOPSD. Instead you can perform the following steps:

If the trace type is INST, BR, or STORAGE, you can use the following method:

1. Change your processor’s mode of operation to manual.
2. Alter bits 0 through 2 of control register 9 to zero (using the alter/display feature).
3. Let SDAID finish execution by changing your processor’s mode of operation back to normal.

For more information on the values to be set into the control registers used by SDAID, consult the *Principles of Operation* manual pertaining to your processor.

If OUTDEV is a printer, the following method to stop the trace is possible:

1. Stop the printer device.
2. Wait until the system goes into the wait state.
3. Press the external interrupt key to stop the trace output.
4. Stop the trace with the STOPSD command.

Starting/Terminating Tracing in a System Wait Condition

In some cases SDAID forces a system wait condition. How you can restart the system by starting or terminating the trace options is the subject of this section.

Wait Due to OPTION=HALT

You may define that the system enters the wait state at occurrence of a specific event. This is accomplished by the option ‘OPTION=HALT’ defined together with the desired event.

When the system has entered the wait, the address part of the wait PSW contains the value X’00EEEE’. The following actions may be taken to get out of the wait state:

1. If you want to continue tracing:
Press the external interrupt key once. The system will enter the wait state again on the next occurrence of the traced event.
2. If you want to continue tracing but without OPTION=HALT:
Enter X’FF’ in storage location zero,
Press the external interrupt key.

This removes the OPTION=HALT specification. The system continues tracing but does not enter the wait state on the next occurrence of the same event again.

System Wait Due to Intervention Required at the Output Device

SDAID loads a wait PSW with the value of X'EEEEEE' in the address part. The required operator action is described under "Exceptional Conditions on the Output Device" on page 55.

Part 4. Info/Analysis

Info/Analysis is a tool for:

- Dump file management
- Problem source identification
- Problem analysis

Chapter 13. Info/Analysis: Introduction

With Info/Analysis, you can simplify the task of using dump data to solve software problems. Info/Analysis assists you in this task through the following functions:

- Dump Management - to list the dumps being managed by Info/Analysis, to add or delete dumps from that list, and to delete dumps from the system.
- Dump Symptoms - to display problem failure information collected by the dumping component and by subsequent analysis routines.
- Dump Viewing - to display dump data in hexadecimal and character format, format, and display control blocks and other dump data that may be pertinent to the problem, to invoke dump analysis routines, and to display the results of those routines.
- Dump Offload - to copy a dump to tape for later retrieval.
- Dump Onload - to copy a dump to a dump sublibrary.

Operating Environment

Info/Analysis runs in a z/VSE partition with a size of at least 1M of storage.

Info/Analysis processes storage dumps that result from errors within the system, subsystems or user programs running on the system. The dumps are created by system dump and stand-alone dump programs. Info/Analysis does not directly access the dump data. Rather, it uses system facilities to retrieve and update dump data and the symptom record. The symptom record is a collection of problem-related information stored in the dump and its extensions.

Info/Analysis uses a dump management file to maintain information about dumps. A dump must be identified in this file before it can be processed by Info/Analysis. This file is maintained using the Dump Management function.

Info/Analysis also uses an external routines file. This file contains a list of analysis routines that you may invoke to process dump data. The file also identifies user exit routines and dump access routines called by Info/Analysis.

You may enter control statements in two modes:

- Line mode - from the operator console
- Reader mode - from the system input device reader defined as the input area

From a z/VSE partition, all output of batch operations is routed to the SYSLST device assigned to the partition. In line mode, messages are sent to the console as well as to SYSLST. SYSLST must always be assigned to a unit record device. When running in reader mode, SYSIPT must be assigned to a unit record device.

The Dump Management File

The dump management file BLNDMF contains information about dumps managed by Info/Analysis.

&ia adds this information either during dump management invocation (&ia searches the dump sublibraries for new dumps automatically), or when you

specify the name of a new dump. For a dump produced as a result of a DUMP attention routine command or for a stand-alone dump you want to onload, supply a name via the Info/Analysis statement

```
DUMP NAME (specify the current dump)
```

Once information about a dump has been added to the dump management file, the Info/Analysis functions can be used to process the dump.

A dump entry remains in the dump management file until the dump is deleted using the &ia function.

Initializing the Dump Management File

Before you can use the functions of Info/Analysis the dump management file has to be initialized.

This initialization is accomplished by the UTILITY statement of Info/Analysis. The statement is used at system installation time and whenever you want to initialize or recreate the dump management file, for example after you have increased the size of the file, or after the file has been damaged.

For an explanation of how to change the size of the dump management file see "UTILITY - Initialize Dump Management File" on page 180. Figure 54 shows a job example to initialize the dump management file.

Sample Initialization Job

```
// JOB      INIT
// ASSGN   SYSLST,00E
// ASSGN   SYS020,252      Dump library
// ASSGN   SYS016,252      Dump management file
// ASSGN   SYS017,252      External routines file

// DLBL   SYSDUMP,'VSE.DUMP.LIBRARY',1999/365,SD
// EXTENT SYS020,SYSWK1,1,0,3150,600
// DLBL   BLNDMF,'INFO.ANALYSIS.DUMP.MGNT.FILE',0
// EXTENT SYS016,SYSWK1,1,0,9030,15
// DLBL   BLNXTRN,'INFO.ANALYSIS.EXT.RTNS.FILE',1999/365,SD
// EXTENT SYS017,SYSWK1,1,0,9045,15

// EXEC   INFOANA,SIZE=300K

SELECT DUMP MANAGEMENT      | use the
UTILITY                     | utility
RETURN                      | function

SELECT END

/*
/ &
```

Figure 54. Sample Job: Dump Management File Initialization

The External Routines File

The external routines file contains the names of dump analysis exit routines. These routines are used to analyze dumps stored in one of the dump sublibraries. The external routines file contains the name and, optionally a description of each routine available for use with Info/Analysis.

Presently the external routines file contains four name lines. DFHPD410 analyses a dump of the CICS Transaction Server partition. IJBXDEBUG is the common analysis routine for stand alone dumps. IJBXSDA formats the SDAID buffer in a stand-alone dump. IJBXCSMG formats the console buffer in a stand-alone dump.

The name of the external routines file is BLNXTRN. The job INITDUMP.Z creates the external routines during the system build process. If the external routines file is damaged you may recreate it via a DITTO job (see Figure 55), or if the DITTO program is not available in your system, via an OBJMAINT job (see Figure 56).

Loading the Info/Analysis External Routines File

The sample jobs shown in Figure 55 and Figure 56 record the names of the analysis routines DFHPD410, IJBXCSMG, IJBXDEBUG and IJBXSDA in the Info/Analysis external routines file

Sample Jobs of External Routines File

```
// JOB LOAD1
// DLBL BLNXTRN,'INFO.ANALYSIS.EXT.RTNS.FILE',1999/365,SD
// EXTENT SYS017,SYSWK1,1,0,9045,15
// UPSI 1
// EXEC DITTO
$$DITTO CSQ BLKFACTOR=1,FILEOUT=BLNXTRN
ANEXIT DFHPD410 CICS DUMP ANALYZER
ANEXIT IJBXCSMG ANALYSE CONSOLE BUFFER
ANEXIT IJBXDEBUG ANALYSE STANDALONE DUMP ROUTINE
ANEXIT IJBXSDA SDAID BUFFER FORMATTING ROUTINE
/*
$$DITTO EOJ
/*
/&
```

Note: The example shows the DITTO statements for an external routines file on CKD disk. If the external routines file is on an FBA disk the DITTO command line reads like
\$\$DITTO CSQ BLKFACTOR=1,FILEOUT=BLNXTRN,CISIZE=512

Figure 55. Sample Job: Loading the External Routines File via DITTO

```
// JOB LOAD2
// ASSGN SYS004,00C
// ASSGN SYS005,SYSWK1
// DLBL UOUT,'INFO.ANALYSIS.EXT.RTNS.FILE',1999/365,SD
// EXTENT SYS005,SYSWK1,1,0,9045,15
// EXEC OBJMAINT
./ CARD DLM=$$
./ Copy
ANEXIT DFHPD410 CICS DUMP ANALYZER
ANEXIT IJBXCSMG ANALYSE CONSOLE BUFFER
ANEXIT IJBXDEBUG ANALYSE STANDALONE DUMP ROUTINE
ANEXIT IJBXSDA SDAID BUFFER FORMATTING ROUTINE
$$
/*
/&
```

Figure 56. Sample Job: Loading the External Routines File via OBJMAINT

Label Information for Info/Analysis

Figure 57 shows an example of the DLBL and EXTENT information to submit if you want to use the functions of the Info/Analysis program. These labels should be stored in the system standard label area.

```
* LABELS FOR THE SYSDUMP LIBRARY,  
* THE DUMP MANAGEMENT FILE, AND  
* THE EXTERNAL ROUTINES FILE  
  
... ..  
// DLBL SYSDUMP, 'VSE.DUMP.LIBRARY',1999/365,SD  
// EXTENT ,SYSWK1,1,0,3150,600  
// DLBL BLNDMF, 'INFO.ANALYSIS.DUMP.MGNT.FILE',0  
// EXTENT SYS016,SYSWK1,1,0,9030,15  
// DLBL BLNXTRN, 'INFO.ANALYSIS.EXT.RTNS.FILE',1999/365,SD  
// EXTENT SYS017,SYSWK1,1,0,9045,15  
... ..
```

Figure 57. Example: File Labels for Dump Processing

Functional Overview

When a dump is created, you can use it to solve a problem by taking actions that range from printing the dump symptoms to analyzing the dump in detail. The actions that you take depend on local procedures for dealing with dumps and your own techniques of dump analysis. &pdm is a tool that can be used to enhance these procedures and techniques.

This section presents the stages of a dump's life cycle from problem occurrence to resolution. The ways in which you can use Info/Analysis at each stage are briefly presented.

When a problem occurs during system operation, the detecting component captures the condition of the system in a dump. The component stores the dump in the dump sublibrary designated for the partition that failed. Sometimes, a system operator may detect a problem and use stand-alone dump or other dumping facilities to create a storage dump. Stand-alone dumps are stored on tape or disk.

In either case, a dump contains a copy of system storage, and a symptom record. The symptom record is a collection of failure-related information gathered by the dumping component when the dump is taken or added later by dump analysis routines.

The symptom record may contain:

- A description of the operating environment at the time the problem occurred.
- Symptoms that provide clues to the problem's origins.
- Free-form text and hexadecimal information that may describe the problem.
- Entries that define the format and location of dump data that may be pertinent to the problem. These entries are used when data is displayed in formatted mode.

For further information about the symptom record, see Appendix A, "Symptom Records Overview."

Chapter 14. Dump Symptoms

The first step in dump analysis is to examine any symptoms that are recorded for the problem. This chapter discusses the Dump Symptoms function with which you may display or print the problem symptoms collected by the dumping component at the time of a failure or by subsequent analysis routines. The list of symptoms may indicate a new problem or a duplicate of a previously encountered problem. If sufficient symptoms are provided, they may pinpoint the cause of the failure.

The successful use of the Dump Symptoms function is dependent on the presence of a symptom record in the dump you are processing. The symptom record is created by the dumping component when the dump is taken (see Appendix A, "Symptom Records Overview," on page 217). Figure 58 shows an example of the symptom part of a dump.

```
                                DUMP SYMPTOMS

DUMP NAME .... SYSDUMP.F6.DF600010

ENVIRONMENT:
CPU MODEL ..... 3090
CPU SERIAL ..... 060707
TIME ..... 08:44:51:42
DATE ..... 92/04/12
SYSTEM ID ..... 568606606
RELEASE ..... 1
FEATURE ..... 5C
DUMPTYPE ..... SCPREQ
PROBLEM NUMBER ..

REQUIRED SYMPTOMS:
AB/S0234===>
PIDS/5745SCBTM
RIDS/DFHTUB
REGS/0C14E
REGS/0A050
```

Figure 58. Dump Symptoms Part

Types of Dump Symptoms

The symptoms are organized into the following sections:

- Environment data
- Required symptoms
- Optional symptoms in structured data base (SDB) format
- Optional symptoms in non-SDB format

The symptoms are initially provided by the dumping component. If you subsequently execute analysis routines for the dump, these routines may add symptoms. A plus sign (+) appears before a symptom if it was added by an analysis routine.

The displayed symptoms may pinpoint the cause of the failure. If not, you may compare these symptoms to the symptoms of other locally reported problems. Your installation should set up a procedure whereby a file of problem symptoms is kept and the symptoms can be compared to one another.

If a satisfactory match is found, the problem is considered a local duplicate. If no match is found and the problem is related to an IBM product, a search of known IBM problems would be a logical next step. You may contact IBM service personnel to request this search. If a duplicate set of symptoms is found in either search, a solution may be immediately available or already under investigation. If no match or too many matches occur after a search, additional analysis is necessary. You may perform this analysis using the Dump Viewing function.

Environment

The environment section of the symptom record describes the environment at the time the dump was created. This section is provided by the dumping component. The CPU, operating system, type of dump, and date and time that the dump was taken are identified. Additional items such as release level may be included.

Required Symptoms

Required symptoms are those considered essential for problem analysis. They are provided by the dumping component or by the dump analysis routines.

All of the required symptoms are likely to occur each time the same error occurs. The format of the required symptoms is standardized so that more effective keyword searches for duplicate problems may be conducted.

Each symptom is formatted with a prefix and the specific data connected by a slash. The required symptoms are described under "Symptom Part Description" on page 185.

Optional Symptoms

Optional symptoms may be provided by the dumping component or by the dump analysis routines. These are additional symptoms that apply to the problem and may be present if the problem recurs. Some of these symptoms, for example the component level, are formatted like the required symptoms. There are also free-form symptoms that may be used in problem analysis, but which are not in standardized format.

Chapter 15. Invoking Info/Analysis

You may use Info/Analysis in two ways:

- Line mode - from the operator console
- Reader mode - from the system input device

In either case, you invoke Info/Analysis by submitting a series of job control statements (JCL) followed by control statements that request Info/Analysis functions. All output is routed to the SYSLST device. The output includes the input control statements, the results of processing, and any messages issued by Info/Analysis. If you are working at a console, Info/Analysis also routes messages there. For information on how to print dumps, see “Printing Dump Information” on page 183.

This chapter describes:

- How to invoke Info/Analysis.
- Syntax rules for control statements.
- Each Info/Analysis function and the control statements needed to request it.
- How to end Info/Analysis.

To illustrate this information, the chapter includes example sequences of control statements.

Submitting a Job to Invoke Info/Analysis

You invoke Info/Analysis by submitting the necessary JCL followed by control statements that request functions. You may submit the job either in line mode by entering statements on the console, or in reader mode by submitting a job to the system input device.

Info/Analysis requires a program area of at least 300K bytes and a 24-bit partition GETVIS area of at least 600K bytes. Thus, the partition used for the execution of Info/Analysis should have a minimum of 900K bytes.

If user-written exits are to be used when running Info/Analysis, this size has to be adjusted accordingly. Especially the CICS Transaction Server uses such exits; therefore, a size of at least 4M bytes is necessary to analyze a CICS dump.

With JCL, you must specify any nonstandard system device assignments and pre-allocate and assign any files that you require other than the system libraries. A sample of the JCL for invocation is:

```

// JOB      JCL
// ASSGN   SYSLST,00E
// ASSGN   SYS020,252      Dump library
// ASSGN   SYS016,252      Dump management file
// ASSGN   SYS017,252      External routines file

// DLBL   SYSDUMP,'VSE.DUMP.LIBRARY',1999/365,SD
// EXTENT ,SYSWK1,1,0,3150,600
// DLBL   BLNDMF,'INFO.ANALYSIS.DUMP.MGNT.FILE',0
// EXTENT SYS016,SYSWK1,1,0,9030,15
// DLBL   BLNXTRN,'INFO.ANALYSIS.EXT.RTNS.FILE',1999/365,SD
// EXTENT SYS017,SYSWK1,1,0,9045,15

// EXEC   INFOANA,SIZE=300K

....
....      | Info/Analysis
....      | control statements
....

/*
/&

```

Figure 59. Sample Job: Invoke Info/Analysis

Standard Info/Analysis Job Stream

Figure 59 shows a sample job to invoke the Info/Analysis program (`// EXEC INFOANA`). Assume that the dump sublibraries in the library SYSDUMP have already been defined and the label information for the SYSDUMP library resides in the standard label area. Information on the SYSDUMP library can be found under “The SYSDUMP Sublibraries” on page 11.

Once the JCL has been processed, you are at the selection level in Info/Analysis. The program reads for your control statements. An end of input (`/*`) statement marks the end of these statements. To end your job, enter an end of job (`/&`) statement.

Control Statement Syntax

You operate Info/Analysis by entering control statements. These control statements specify the major functions you wish to perform (Dump Management, Dump Viewing, etc.) and the information necessary to perform each function. This section describes the syntax rules for entering the control statements.

For a description of the syntax diagrams please read “Understanding Syntax Diagrams” on page xv.

Entering Control Statements

The rules for entering the control statements that request Info/Analysis functions are:

- Each card or input line may contain only one control statement.
- A control statement may begin in any column.
- Control statements and their operands may be entered in uppercase or lowercase.

- Control statements must be entered in their complete form; no abbreviations are allowed.
- Each word in a control statement must be a contiguous string of characters.
- Some blanks, at least one, must appear between the words in a statement.
- A blank followed by an asterisk (*) signifies the start of a comment. If the first non-blank character in any control statement is an asterisk, the entire statement is treated as a comment.
- Sequence numbers or other characters should not be placed in columns 73 to 80 of Info/Analysis commands if they are not designated as a comment.

If you enter an invalid control statement in reader mode, the remaining control statements in the job are flushed and the session is canceled. The output indicates the erroneous statement. You should correct the statement and resubmit the job. In line mode, invalid input causes Info/Analysis to flush the control statement. You may then reenter the statement correctly.

Common Control Statements


Every function has its own set of control statements. In this chapter, the control statement descriptions are presented by function. The same control statement may have different uses within different functions. Therefore statements such as PRINT are described more than once. A summary of the control statements is given under “Control Statement Summary” on page 213.

The following control statements are common to all Info/Analysis functions. They may also be entered at the selection level; that is, when Info/Analysis is initialized and no function is currently selected.

- **SELECT** - invoke a function or end the Info/Analysis session
- **RETURN** - end the current function; return to the selection level
- **DUMP NAME** - specifies the name of the dump to be processed.

SELECT - Specify a Function or End Info/Analysis

To perform any function, you must first select it. Use the SELECT statement to specify the function you wish to perform or to end your Info/Analysis session.



You enter a SELECT statement as follows:

```
SELECT DUMP ONLOAD
```

Once you have specified a SELECT statement, you are operating at the function level. All subsequent statements apply to the selected function until you enter a RETURN statement.

Use SELECT END to conclude your Info/Analysis session. If you end your control statement sequence with an end of input (/*) and end of job (/&) statement

without first specifying a SELECT END or RETURN statement, Info/Analysis ends the job as though they had been included.

RETURN - End Current Function

Use the RETURN statement to end the current function. Each function you use must be requested by a SELECT statement and ended by a RETURN statement, before you can select another function.

▶▶—RETURN—▶▶

The RETURN statement has no operands. RETURN brings you back to the selection level where you may select another function or end your session. In the following sequence, RETURN ends the Dump Management function:

```
DUMP NAME SYSDUMP.F6.DF600013
SELECT DUMP MANAGEMENT
PRINT DATA
RETURN
SELECT END
```

You must enter RETURN to complete one function before selecting another. If you end your control statement sequence with an end of input (/*) and end of job (/&) statement without first specifying a RETURN and/or SELECT END statement, Info/Analysis ends the job as if they had been included.

DUMP NAME - Specify or Add Current Dump

Specify the name of the dump you wish to process by entering the DUMP NAME statement. You must specify a dump before you perform any Info/Analysis function except the Dump Management UTILITY or PRINT functions and the HELP function.

▶▶—DUMP NAME *dumpname*—▶▶

“dumpname” is a variable representing the dump name. The dump you specify on a dump name statement is considered to be the current dump; that is, all subsequently selected functions process that dump until you enter another dump name or until it has been deleted.

When you specify a dump, Info/Analysis searches the dump management file for the name. If it finds the name, the dump is made current.

If the dump you specify is not identified in the dump management file and does not reside in a dump sublibrary, Info/Analysis adds the dump name and the information “TO BE ONLOADED” to the file. A dump must reside in one of the dump sublibraries before you can act on it using any Info/Analysis function other than Dump Onload.

You may specify the DUMP NAME statement at the selection level. In other words, the DUMP NAME statement may be placed:

- Immediately preceding a SELECT statement, as follows:

```
DUMP NAME SYSDUMP.F3.DF300010
SELECT DUMP MANAGEMENT
PRINT DATA
RETURN
```

- After the SELECT DUMP MANAGEMENT statement and before the next RETURN statement, as follows:

```
SELECT DUMP MANAGEMENT
DUMP NAME SYSDUMP.F3.DF300010
PRINT DATA
RETURN
```

Recommendations (Restrictions) for the Generation of Dump Names

The VSE dump routines generate dump names for ABEND dumps, IDUMPs, and SDUMPS. These system generated dump names have the format 'Dppnnnnn', or 'Sppnnnnn', where the character 'D' identifies dumps of the executing space, and 'S' identifies dumps of associated data spaces. The character combination 'pp' denotes the partition identification, like BG or F7, and 'nnnnn' is a unique decimal number between 00000 and 99999

When Info/Analysis processes a dump via the dump viewing function, it enters additional members into the dump library to save intermediate analysis results. These dumps are named 'Maaaaaaa' or 'Xaaaaaaa', where 'M' or 'X' replace the heading character 'D' or 'S' and aaaaaaa is the unchanged trailing portion of the name of the analysed dump. If you have, for example, a dump named SX400002 in the dump library, then the Infoana program will create the library members MX400002 and XX400002 for its internal use.

You define the names for stand-alone dumps or attention dumps during the onload process. Here are some recommendations which help to build unique dump names.

- Do not use D, H, L, M, N, S, X, as the first character of a dump member name. These initial characters are reserved for use by the system.
- Instead use **one** heading letter, for example the letter 'A' as first letter of all dump names within a sublibrary.
- If you use a partition identification within a dump name, do not enter it in positions 2 to 3. This might generate conflicts with system generated names of ABEND dumps, IDUMPS, or SDUMPS.

Dump Management

The Dump Management function enables you to manage your dump data sets by manipulating the contents of the dump management file. This file contains descriptive information about the dumps being managed by Info/Analysis. You can use this information to keep track of the dumps on your system, those you have offloaded, and those you plan to onload.

To initiate Dump Management, use the following statement:

```
▶▶—SELECT DUMP MANAGEMENT—◀◀
```

After this selection, specify the desired combination of the following control statements:

UTILITY -

initialize a new or reallocated dump management file to contain the list of dumps managed by Info/Analysis.

DELETE -

erase the current dump from the dump sublibrary, if it resides there, and delete information about the dump from the dump management file.

PRINT DATA -

print the contents of the dump management file.

Info/Analysis responds by searching the dump sublibraries for dumps that are not yet identified in the dump management file. For each of these dumps, if any, the routine adds identifying information to the file.

UTILITY - Initialize Dump Management File

►►—UTILITY—◄◄

The UTILITY statement is intended for the system programmer at your installation who has responsibility for Info/Analysis. UTILITY initializes the dump management file at installation time or reinitializes the file when it is subsequently reallocated with more or less space.

The dump management file is allocated during installation of your VSE system. The size of the file is big enough to hold some hundred dump names. (For example, the DMF allocated on a 3380 disk is sufficient for about 1000 dump names.) For performance reasons you should not increase the size of the DMF. Instead it is recommended that you clean-up the dump library from time to time, and delete the dumps which are no more used or offload those dumps to tape which may be used at a later time.

You initialize the dump management file with the following statements:

```
SELECT DUMP MANAGEMENT
UTILITY
RETURN
SELECT END
```

UTILITY sets up the control information in the file for use by the dump management function. The control record indicates the number of dumps currently being managed and the maximum number that will fit.

DELETE - Delete Current Dump

You can erase a dump, and delete its corresponding information in the Info/Analysis dump management file, by using the DELETE statement.

►►—DELETE—◄◄

The example in Figure 60 on page 181. shows a job to delete two dumps in one dump management run. Note, that the DUMP NAME statement is valid prior as well as after the SELECT DUMP MANAGEMENT statement.

If you specify DELETE and the current dump does not reside in the library (that is, it has never been onloaded or it has been offloaded), information about the dump

is deleted from the dump management file. If a copy of the dump exists on tape, it is your responsibility to dispose that tape.

```
// JOB    DELETE
// ASSGN  SYSLST,00E
...
// EXEC  INFOANA,SIZE=300K
SELECT DUMP MANAGEMENT      | Calls dump management
DUMP NAME SYSDUMP.F5.DF500002 | Specifies the dump
DELETE
DUMP NAME SYSDUMP.F5.DF500003
DELETE

RETURN
SELECT END
/*
/ &
```

Figure 60. Sample Job: Delete Dumps

Info/Analysis indicates the successful execution of the DELETE function with a message (DUMP dumpname DELETED).

If you want to retain the information about the dump in the dump management file, you can use the SELECT DUMP OFFLOAD operation with the BYPASS operand. Please see Figure 73 on page 202, which shows a summary of the OFFLOAD and DELETE functions.

PRINT - Print List of Managed Dumps

Use the PRINT statement to print the contents of the dump management file.

▶▶—PRINT DATA—◀◀

Output from the PRINT statement is a list of the dumps being managed by Info/Analysis. For each dump, one line of information is printed. A sample job is shown in Figure 61 on page 182.

```

// JOB LIST
// ASSGN SYSLST,00E
...

// EXEC INFOANA,SIZE=300K

SELECT DUMP MANAGEMENT      | calls the &osq.list
PRINT DATA                  | managed dumps&csq.
RETURN                       | function
                               |
SELECT END

/*
/ &

```

Figure 61. Sample Job: List Managed Dumps

An output example of the PRINT DATA function is shown in Figure 62.

DUMP NAME	RELATED DUMP	ONLINE	DATE/TIME	TAKEN	VOLID	DATA SPACE NAME
SYSDUMP.BG.DBG00000	SBG00006	Y	91/02/21	11:10:00	VOLID0	
	SBG00007					
	SBG00008					
SYSDUMP.BG.DBG00001	--NONE--	Y	91/02/21	11:11:00	VOLID1	
SYSDUMP.BG.DBG00002	--NONE--	Y	91/02/21	11:12:51	VOLID2	
SYSDUMP.BG.DBG00003	SBG00009	Y	91/02/21	11:13:00	VOLID3	
	SBG00010					
SYSDUMP.BG.DBG00004	--NONE--	Y	91/02/21	11:15:53	VOLID4	
SYSDUMP.BG.DBG00005	--NONE--	Y	91/02/21	11:16:54	VOLID5	
SYSDUMP.BG.SBG00006	DBG00000	Y	91/02/21	11:10:10	VOLID6	DATSPAC1
SYSDUMP.BG.SBG00007	DBG00000	Y	91/02/21	11:10:20	VOLID7	DATSPAC2
SYSDUMP.BG.SBG00008	DBG00000	Y	91/02/21	11:10:30	VOLID8	DATSPAC3
SYSDUMP.BG.SBG00009	DBG00003	Y	91/02/21	11:13:20	VOLID9	DATSPAC4
SYSDUMP.BG.SBG00010	DBG00003	Y	91/02/21	11:13:40	VOLID10	DATSPAC5

Figure 62. Example: Dump Management PRINT DATA Output

The column headings represent:

- DUMP NAME - The identifier of the dump.
- RELATED DUMP - Displays the correlation between an ABEND or stand-alone dump and its data space dumps:

SYSDUMP.xx.Dxxnnnnn denotes an ABEND or stand-alone dump.

SYSDUMP.xx.Sxxnnnnn denotes a data space dump.

Note that in the column RELATED DUMP, the library and sublibrary of the dump is not displayed, since it is the same as in the corresponding DUMP NAME column.

For lines containing an ABEND or stand-alone dump in the column DUMP NAME, the column RELATED DUMP contains the names of the data space dumps that were created with the ABEND or stand-alone dump. Vice versa, for lines containing a data space dump in the column DUMP NAME, the second column contains the name of the ABEND or stand-alone dump which was taken together with the data space dump.

If the entry in the column RELATED DUMP contains --NONE--, the named dump did not access any data space or the data spaces are not indicated in the optional symptoms.

Information Analysis extracts this information from the optional symptoms of the symptom record.

When you delete a dump without deleting all related dumps, this overview listing becomes inconsistent, which means that there are no longer lines for all related dumps.

- ONLINE - An indication (Y, for yes) if the dump is currently stored in the dump sublibrary.
- DATE/TIME TAKEN - The date and time the dump was created or, if the actual date and time are not available, the date and time the dump was identified to Info/Analysis.
"TO BE UNLOADED" indicates that the dump is not in the dump sublibrary (for example a stand-alone dump named to the Info/Analysis management, but not yet unloaded).
- VOLID - The identifier of the tape volume to which the dump has been offloaded or from which the dump has been onloaded, if any. If a dump has been offloaded and then onloaded again, Info/Analysis retains the volume id in the dump management file. Consequently, a dump may be in a dump sublibrary ("Y" in ONLINE field) and still have a VOLID.
- DATA SPACE NAME - The name of the data space (as specified in the DSPSERV macro).

Printing Dump Information

The following sections describe how Info/Analysis is used to print information from the dumps stored on a tape or disk, or in a dump sublibrary.

The contents of a given dump depends on the function which created the dump. But the main form is the same for all dump requesting functions which are able to store dumps on tape, disk, or in a dump sublibrary. The functions described in the subsequent sections can be used for any type of dump stored in a dump sublibrary.

Figure 63 on page 184 gives an overview of the various parts of a dump. Info/Analysis can be used to print these parts selectively.

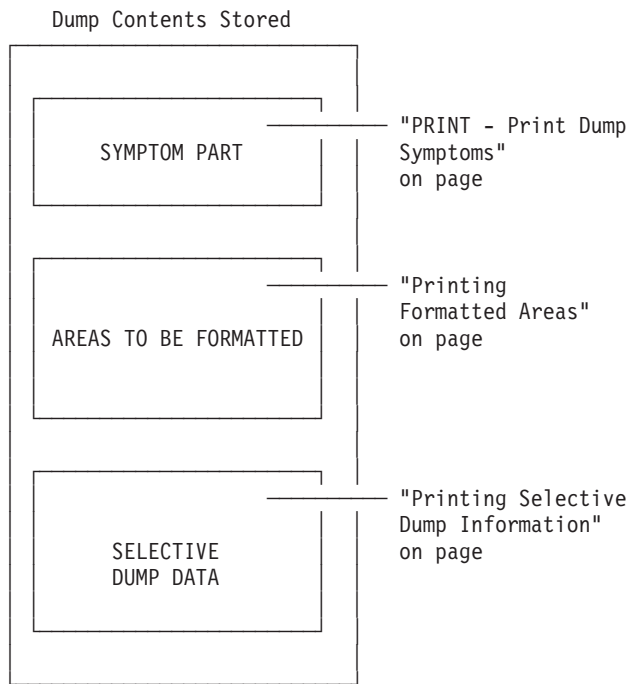


Figure 63. Overview: Dump Contents

Dump Symptoms

The Dump Symptoms function prints the failure information that is contained in the dump symptom record. You may use this information to identify duplicate problems locally and in an IBM maintenance data base. To initiate Dump Symptoms, enter the following statement:

▶▶—SELECT DUMP SYMPTOMS—▶▶

After making this selection, you may print the symptoms of the current dump using the PRINT statement.

PRINT - Print Dump Symptoms

▶▶—PRINT DATA—▶▶

Use the PRINT statement to print sections 1 through 5 of the symptom record of the current dump. For a description of a symptom record, see Appendix A, "Symptom Records Overview," on page 217.

Figure 64 shows a sample print dump symptoms job.

```

// JOB PRINT
.
.
.
// EXEC INFOANA,SIZE=300K

DUMP NAME SYSDUMP.F3.DF300010 | defines the
                                | dump name

SELECT DUMP SYMPTOMS          | calls the
PRINT DATA                   | print operation

RETURN

SELECT END
/*
/&

```

Figure 64. Sample Job: Print Dump Symptoms

Symptom Part Example

“Printed Output of the Main Dump File of a Stand-Alone Dump” on page 208 shows an output example of a print dump symptoms operation.

```

                                DUMP SYMPTOMS
SYSDUMP.BG.TEST

ENVIRONMENT:
CPU MODEL ..... 3090
CPU SERIAL ..... CF5A05
TIME ..... 07:56:09:00
DATE ..... 94/11/17
SYSTEM ID ..... 568606606
RELEASE ..... 1
FEATURE ..... 5C
DUMPTYPE ..... SADUMP
PROBLEM NUMBER ..

REQUIRED SYMPTOMS:

OPTIONAL SYMPTOMS (SDB):

OPTIONAL SYMPTOMS (NON-SDB):

```

Figure 65. Example: Output of Print Dump Symptoms

Symptom Part Description

ENVIRONMENT

The environment section of the symptom record describes the environment at the time the dump was taken. This data is provided by the dumping system component.

All of the information in the environment section is self-explanatory except the dump type entry, which may be:

SCPREQ

for ABEND dumps

IDUMP

for internal VSE/Advanced Functions dump requests

OPRREQ

for DUMP command dumps

SADUMP

for stand-alone dumps

SDUMP

for SDUMP and SDUMPX macro dumps

REQUIRED SYMPTOMS

Required symptoms are those considered essential to your problem analysis effort:

- The symptoms for ABEND dumps are described below.
- DUMP command dumps only provide the DUMP command itself (as typed in at the console).
- SDUMPs and IDUMPs contain the symptoms that were provided at the respective macro invocation.
- Stand-alone dumps do not contain required symptoms. Details on the dump contents can be obtained by executing the stand-alone dump analysis routine IJBXDEBUG, for example.

Each symptom is headed by a prefix; the specific data is connected by a slash. The possible REQUIRED SYMPTOMS (for an ABEND dump) are:

AB/Sxxyy

ABEND (cancel) code

xx ... first cancel code

yy ... second cancel code or 00 if none exists

Example: AB/S0900

The following symptoms are additional information to the various cancel codes.

ADRS/address

Reflects the absolute address of the instruction following the failing one if the failing address is outside the LTA, SVA, or partition areas (the address is extracted from the PSW).

MS/xxxxx

Message number

Example: MS/0V15I

OFFS/offset

Offset of the instruction following the failing one relative to the phase start address, if the phase resides in the LTA or in the SVA, or relative to the partition begin address (the address is calculated from the PSW).

OPCS/aaaaaa

aaaaaa Represents either

SVCnn (nn = SVC code) or

CODEmm (mm = program interruption code)

An entry for OPCS is generated only if an illegal SVC or if a program check occurs. If an illegal SVC occurs, the OPCS entry contains the failing SVC code in decimal. If a program check occurs, the entry contains the program interrupt code.

PIDS/comp.id

Component identifier

Example: PIDS/5686032V2

REGS/xyyy

- xx is the register number in hex which contains a value, less than and within 4K (K=1024 bytes) of the PSW address at the time of failure. This register contains an address that is close to the point of failure.
- yyy is the difference between the PSW and the register address.

Example: REGS/0C14E

0C number of the
 general purpose register X'C'
14E is the difference between the
 PSW address minus the contents
 of register (C).

RIDS/Caaaaaaaa

aaaaaaaa stands for:

- The name of a phase, if the failing instruction is in the SVA or LTA.
- The name of a program if the error occurred in a partition.

Example: RIDS/DFHTUB

VALU/Caaaaaaaa

aaaaaaaa represents either:

- phase name (for example, combined with AB/S2200 phase not found)
- SYSnnn (for example, together with AB/S2600 SYSnnn not assigned).

OPTIONAL SYMPTOMS (SDB)

SDB stands for structured data base. For information on the symptom record please refer to Appendix A, "Symptom Records Overview," on page 217.

These symptoms (SDB) may be provided by the function which produced the dump or by subsequently executed analysis routines. These additional symptoms apply to the problem and may be present if the problem recurs. They are in SDB format; for example, the component level may be included in this section.

OPTIONAL SYMPTOMS (NON-SDB)

These symptoms (non-SDB) are optionally provided by the dump originating component or by subsequently executed analysis routines. They are free-form symptoms that may be used in problem analysis but do not fit into the SDB format.

Note, that a symptom added by an analysis routine contain a preceding plus sign (+).

Optional (Non-SDB) Symptoms of a Stand-Alone Dump: In addition to the symptoms described above, a symptom record of a stand-alone dump shows the following entries (see also "Printed Output of the Main Dump File of a Stand-Alone Dump" on page 208):

ACTIVE_SPACE_ID

which identifies the address space that was active at the time the dump was taken.

DATA_DUMPED_FROM_SPACE_ID

which indicates the address space from which this dump data file was dumped.

JOB_NAME or DATA_SPACE_NAME

which indicates the job or data space name for this dump file.

DUMPED_DATA

which shows what data is in this dump file.

Optional (Non-SDB) Symptoms of an ABEND Dump: The optional symptoms part of a partition dump shows whether there are any appended data space dumps or not. The symptom gives a list of all data space names which belong to the dumped partition:

```
DATA_SPACES=(aaaa,bb,ccccc)
```

where

```
aaaa,bb,ccccc
```

is the name of the address space or data space.

The symptom record of the data space contains the following symptoms:

```
DATA_SPACE=aaaa
```

```
ALET=alet
```

```
RELATED_ABEND_DUMP=DF400001
```

Dump Viewing

Using Dump Viewing you may print dump data and analysis summary data. You may also call an analysis routine for execution. You cannot locate particular data or mask data for security purposes. To initiate Dump Viewing, enter the following statement:

```
▶▶—SELECT DUMP VIEWING—▶▶
```

After making this selection, you may perform Dump Viewing functions by specifying the following control statements:

- PRINT - print dump data
- CALL - initiate an analysis routine

PRINT - Print Dump Data

Use the PRINT statement to print dump data or analysis summary data. You must specify either the area of the dump you wish to print in hexadecimal mode or that you wish to print all formatted and all analysis summary data. The operands for the PRINT statement are mutually exclusive.

```
▶▶—PRINT—

|                                    |
|------------------------------------|
| <i>from_addr</i>                   |
| <i>from_addr to_addr</i>           |
| <i>from_addr</i> END               |
| <i>from_addr</i> FOR <i>length</i> |
| FORMAT                             |

—▶▶
```

The PRINT control statement for Dump Viewing requires either an address range or the FORMAT operand. The results of PRINT with the addr-range operand are printed in traditional hexadecimal format with EBCDIC translation (refer to "Printing Selective Dump Information" on page 189). The FORMAT operand prints the analysis summary if stored during a previous analysis routine run, and formatted dump display information (refer to "Printing Formatted Areas" on page 190). This data includes control blocks, text data, hexadecimal data, and control block linkage descriptors as defined in section 6 of the symptom records. For a description of the symptom records, see Appendix A, "Symptom Records Overview," on page 217.

PRINT

from_addr

'from_addr' marks the beginning of the 8192 byte area to be printed.

from_addr to_addr

'from_addr' marks the beginning of the area to be printed and 'to_addr' marks the end.

from_addr END

'from_addr' marks the beginning of the area to be printed and 'END' indicates that the data up to the high address end of the dumped storage is to be printed.

from_addr FOR length

'from_addr' marks the beginning of the area to be printed and 'length' represents the number of bytes in hexadecimal which are to be printed. For example, if you specify 10 as a length, 16 bytes are printed.

All addresses are 1- to 4-byte hexadecimal values representing valid addresses in the dump. Leading zeros are not required for an address specification.

FORMAT

causes the data to be printed with correlated field names and other identifiers. The data printed is determined by information in section 6 of the symptom records.

For print job examples, see Figure 66 on page 190 and Figure 67 on page 190.

Printing Selective Dump Information

The PRINT statement of the &ia SELECT DUMP VIEWING function can be used to print dump information selected by addresses.

All addresses and the length setting are 1- to 8-character hexadecimal values representing valid addresses in the dump. Leading zeros are not required for an address specification. The specification of

```
PRINT 0 END
```

for example, would cause the whole dump data to be printed.

The example in Figure 66 on page 190 shows the:

- Definition of the dump SYSDUMP.F3.DF300010 to be processed.
- Definition of the selective print operation for dump data, beginning at the address X'302000' and ending at dump end.

```

// JOB PRINT
.
.
.
// EXEC INFOANA,SIZE=300K

DUMP NAME SYSDUMP.F3.DF300010 | defines the
                               | dump name

SELECT DUMP VIEWING           | defines the
PRINT 302000 END              | area to be
                               | printed

RETURN
SELECT END
/*
/ &

```

Figure 66. Sample Job: Print Selected Dump Areas

Printing Formatted Areas

The printed output of a dump is called formatted if selected system information is extracted and printed separately.

The example in Figure 67 shows how to print the dump SYSDUMP.F3.DF300010 formatted on the device at SYSLST.

```

// JOB PRINT
.
.
.
// EXEC INFOANA,SIZE=300K

DUMP NAME SYSDUMP.F3.DF300010 | defines the
                               | dump name

SELECT DUMP VIEWING           | calls the
PRINT FORMAT                  | PRINT FORMAT
                               | function

RETURN
SELECT END
/*
/ &

```

Figure 67. Sample Job: Print a Dump in Formatted Form

CALL - Initiate Analysis Routine

Use the CALL statement to invoke an analysis routine.

►►—CALL *routine_name*—◄◄

The *routine_name* is required on a CALL statement and must be the name of an executable routine. These routines may be provided by system components or by your location. It is the responsibility of your location to maintain an external routines file containing the names of executable routines. Consult your system programmer for the names of these routines.

A call job example is given in Figure 68. It shows a job which

- Selects the dump SYSDUMP.BG.ADUMP10;
- Calls the analysis routine IJBXCSMG.
- Calls the analysis routine IJBXDEBUG.
- Calls the analysis routine IJBXSDA.

Analysis routines can be used to analyze the stored dumps. For example, the routine IJBXDEBUG (shipped as part of VSE/Advanced Functions) analyzes the output of stand-alone dumps and adds the analysis information to the dump sublibrary.

For a description of the analysis routines available with VSE, see “The Stand-Alone Dump Analysis Routine IJBXCSMG,” “The Stand-Alone Dump Analysis Routine IJBXDEBUG” on page 192 and “The Stand-Alone Dump Analysis Routine IJBXSDA” on page 198.

```
// JOB ANALYZE
// ASSGN  SYSLST,00E
...
// EXEC  INFOANA,SIZE=300K
DUMP NAME SYSDUMP.BG.ADUMP10
SELECT DUMP VIEWING
CALL IJBXCSMG
CALL IJBXDEBUG
CALL IJBXSDA
RETURN
SELECT END
/*
/ &
```

Defines the dump name
Defines the analysis routine call

Figure 68. Sample Job: Call the Analysis Routines IJBXCSMG, IJBXDEBUG and IJBXSDA

The Stand-Alone Dump Analysis Routine IJBXCSMG

IJBXCSMG is an exit routine of Info/Analysis which processes the console events by redisplaying about the last most recent 20 messages and inputs including the timestamp and the console name where the source is coming from. IJBXCSMG prints the last console entries on SYSLST.

Activating the routine

The routine must be contained in the Info/Analysis external routines file before you can call the routine.

Start execution of the analysis routine with the Info/Analysis statements:

```
SELECT DUMP VIEWING
CALL IJBXCSMG
```

Figure 68 illustrates the activation of an analysis routine.

The Stand-Alone Dump Analysis Routine IJBXDEBUG

The analysis routine IJBXDEBUG analyses the output of a stand-alone dump unloaded to a dump sublibrary.

When IJBXDEBUG receives control from Info/Analysis, it requests portions of dump data using Info/Analysis dump access routines. During analysis of the dump data, information from the dump is extracted and written back to the dump sublibrary for further Info/Analysis operations.

Activating the Routine

The routine name IJBXDEBUG must be contained in the Info/Analysis external routines file before you can call the routine.

The statements

```
SELECT DUMP VIEWING  
CALL IJBXDEBUG
```

of Info/Analysis start execution of the analysis routine.

Figure 68 on page 191 illustrates the activation of an analysis routine.

Output of the Routine

The output of the analysis routine may contain the following:

- General information.
- Specific information.

While general analysis information is provided in every dump, specific analysis information concerns only particular error situations. The output of the analysis routine is stored in the dump sublibrary. It can be printed together with formatted dump areas with the operation:

```
DUMP NAME .....  
SELECT DUMP VIEWING  
PRINT FORMAT
```

General Analysis Information

The general analysis information, which is provided for each dump, is not dependent on certain error conditions.

Header entry: This contains data as follows:

- Service level identifier
- Supervisor ID
- Supervisor name
- Date the dump was taken
- Dump type
- System status
- Current task
- Owner of LTA and transient name (if active)

Address Validation: When an address has been located or calculated during the analysis process of the dump data, it is validated. IJBXDEBUG checks whether the address is

- Within the range of high and low limits of the affected areas:
 - Supervisor
 - Partition

- SVA.

When an address is found to be invalid, the address and information about the expected contents of the address are added to the dump. An address validation entry might look like this:

```
INVALID ADDRESS FF0002 ENCOUNTERED DURING ANALYSIS.  
ADDRESS OF: PUB TABLE FROM BG COMREG
```

Specific Analysis Information

The data which is selected, analyzed, and finally stored in the dump sublibrary depends on the error situation. The following error conditions can be recognized by the IJBXDEBUG routine:

- Hard Waits
 - WAITFFA
 - WAITFFB
 - WAITFFF
 - WAITFF9
 - WAITFFE
 - WAITFF4
 - WAITFF8
 - WAITFD0
 - Other Hard Waits
- Soft Waits or System Running (Loop)

Hard Wait Dump Entries: For all dumps that indicate a hard wait, the routine supplies

- Wait state code
- Hard wait reason code
- General purpose registers at the time the failure occurred
- Access registers at the time the failure occurred

Besides this information, the following data, depending on the specific hard wait code, is provided:

For *WAITFFA*, *WAITFFB*, or *WAITFFF*

- Type of program check
- Program check address
- Instruction at the program check address
- Overwritten instruction information (if applicable)
- Name of transient which program-checked and the displacement within the transient (if applicable)
- Transient areas checked are LTA, PTA, DOC, and RTA
- Name of the SVA phase which program-checked and the displacement within the phase (if applicable)
- Registers at the time of failure
- 64 bytes of data pointed to by each register

For *WAITFF9* or *WAITFFE*

- Last device to which a sense was issued
- Sense data address
- Sense data
- Registers at the time of failure
- 64 bytes of data pointed to by each register.

For *WAITFF4*

- Error recovery phase which has not been found

- Registers at the time of the failure
- 64 bytes of data pointed to by each register

For *WAITFF8*

- CRT phase which has not been found
- Registers at the time of the failure
- 64 bytes of data pointed to by each register

For *WAITFD0*

- IPL cancel code
- IPL cancel reason code
- Registers at the time of the failure
- 64 bytes of data pointed to by each register

For *other Hard Waits*

- Registers at the time of the failure
- 64 bytes of data pointed to by each register

Soft Waits or System Running: On all dumps indicating a soft wait or a system running condition, the status of all active devices (non-telecommunication) and active tasks is supplied. Information is provided for active devices, excluding local telecommunication devices. A device is active if:

- It is flagged busy in the PUB table
- It has a channel queue entry queued to the PUB table.

A task is active if it is not unbatched, stopped, or flagged not active in the TIB (Task Information Block). Tasks of VSE/POWER are classified as not active if VSE/POWER has flagged the task as waiting for work. The following information is provided:

Device Status Information

- Device address
- Device type
- Task-id of first channel queue entry
- I/O request status (I/O started or not started)
- Reason I/O not started (CSW stored, intervention, etc.)
- CSW from channel queue entry if interrupt has been presented
- A list of additional tasks with I/O queued for this device
- The device, that last presented an interrupt to the system

Task Status Information

- Task name
- Serviced task name
- Main task name for subtask
- Status
- What the task is waiting for
- Task information (LTA active, ICCF pseudo partition, EOT active, etc.)
- Subsystems running in the partition

Output Examples

Examples of IJBXDEBUG output are given:

- For a hard wait X'FFF' - See Figure 69 on page 195.
- For a soft wait - See Figure 70 on page 196.
- For a system loop - See Figure 71 on page 197.

DATE DUMP WAS TAKEN: 11/17/94 DUMP TYPE: SADUMP
SUPERVISOR ID: Y-ESA10/27-16.20 SUPERVISOR NAME: \$\$A\$SUPX
SYSTEM STATUS: HARD WAIT HARD WAIT CODE: FFF
CURRENT TASK: CST TASK

HARD WAIT REASON CODE: 24 - PROGRAM CHECK IN SUPERVISOR

PROGRAM OLD PSW INDICATES BC MODE.

PROGRAM CHECK TYPE: 0000 (UNKNOWN TYPE)
ADDRESS OF PROGRAM CHECK: 0000B798
PROGRAM CHECK INSTRUCTION: 9098028058D0

INSTRUCTION LENGTH CODE ZERO.

SYSTEM STATUS: HARD WAIT
CURRENT TASK: CST TASK

DEVICE ANALYSIS FOR ACTIVE NON TP DEVICES ONLY:

DEV	TYPE	TSK	I/O REQUEST STATUS AND INFORMATION
009	1052	CST	I/O STARTED, AWAITING INTERRUPT DEVICE END POSTING REQUIRED
001	N/A	N/A	LAST I/O INTERRUPT WAS FROM THIS DEVICE DEVICE NOT FOUND IN PUB TABLE

TASK ANALYSIS FOR ACTIVE TASKS ONLY:

TASK NAME	STATUS	TASK INFORMATION
CST TASK	READY	READY TO RUN
T13 TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 03F1F0
AR TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 000000 SVC RETRY INDICATOR ON SVC: 1D (HEX)
BG MAIN TASK	READY	READY TO RUN LIBRARIAN SERVICE ACTIVE SYSTEM CODE ACTIVE (LIBR)

Figure 69. Example: WAITFFF Analysis Report

IBM VSE/ESA 5686-6606-15C ECLEV=0000

DATE DUMP WAS TAKEN: 11/17/94	DUMP TYPE: SADUMP
SUPERVISOR ID: Y-ESA10/27-16.20	SUPERVISOR NAME: \$\$\$SUPX
SYSTEM STATUS: SOFT WAIT	
CURRENT TASK: BG MAIN TASK	BASE PHASE: NO NAME

DEVICE ANALYSIS FOR ACTIVE NON TP DEVICES ONLY:

DEV TYPE TSK I/O REQUEST STATUS AND INFORMATION

(NO BUSY DEVICES AND NO DEVICES WITH I/O QUEUED)

TASK ANALYSIS FOR ACTIVE TASKS ONLY:

TASK NAME	STATUS	TASK INFORMATION
CST TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 000000 SVC RETRY INDICATOR ON SVC: 1D (HEX)
T13 TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 03F1F0
AR TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 000000 SVC RETRY INDICATOR ON SVC: 1D (HEX)
BG MAIN TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 403504 JOB CONTROL ACTIVE IN THIS PARTITION OCCF SERVICE REQUEST PENDING

Figure 70. Example: Soft Wait Analysis Report

IBM VSE/ESA 5686-6606-15C ECLEV=0000

DATE DUMP WAS TAKEN: 11/17/94 DUMP TYPE: SADUMP
SUPERVISOR ID: Y-ESA10/27-16.20 SUPERVISOR NAME: \$\$A\$SUPX
SYSTEM STATUS: RUNNING PSW: 01000F0F 0027D01C
CURRENT TASK: BG MAIN TASK BASE PHASE: GSXDEBUG

AREA POINTED TO BY PSW: SVA PHASE NAME: \$IJB LBR

SCANNING SVA/\$IJB LBR , APAR IDENTIFIERS WERE FOUND.
APAR NUMBERS: DY15017 DY20012

ASYNOC REPLIES OUTSTANDING: F3-003 F2-002

DEVICE ANALYSIS FOR ACTIVE NON TP DEVICES ONLY:

DEV	TYPE	TSK	I/O REQUEST STATUS AND INFORMATION
009	1052	F7	I/O NOT STARTED, REASON UNKNOWN OTHER TASKS QUEUED: SNS,FA,ASY CHANNEL QUEUE CHAIN ERROR
00E	1403	F4	I/O NOT STARTED, INTERVENTION REQUIRED
02A	2540	FB	I/O STARTED, AWAITING DEVICE END
100	3375	T37	I/O NOT STARTED, CSW STORED CSW: 00005980 001400007
104	3380	F2	I/O STARTED, AWAITING INTERRUPT OTHER TASKS QUEUED: T38,F5 I/O STARTED ON ALTERNATE CHANNEL
204	3375	F8	I/O STARTED, INTERRUPT RECEIVED CSW: 00105A98 0C40009E RAS RETRY INDICATOR ON FOR THIS DEVICE
260	3380	N/A	LAST I/O INTERRUPT WAS FROM THIS DEVICE

Figure 71. Example: System Loop Analysis Report (Part 1 of 2)

TASK ANALYSIS FOR ACTIVE TASKS ONLY:

TASK NAME	STATUS	TASK INFORMATION
SNS FOR F6	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 001212
ASY TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 002212
BG MAIN TASK	READY	READY TO RUN
F1 MAIN TASK	RURBND	WAIT FOR LOCKED RESOURCE RESOURCE: PAYROLL.FILE OWNER: BG SVC RETRY INDICATOR ON SVC: 6E (HEX)
F2 MAIN TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 1BF628 SUB SYSTEMS IN THIS PARTITION: ICCF CICS
F3 MAIN TASK	READY	READY TO RUN VSAM AUTO CLOSE ACTIVE FOR THIS TASK
F4 MAIN TASK	CHQBND	WAIT FOR CHANNEL QUEUE
F7 MAIN TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 5BF620 TASK IS LTA OWNER
F8 MAIN TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 474628 TERMINATOR ACTIVE FOR TASK
F9 MAIN TASK	LTABND	WAITING FOR LTA
FA MAIN TASK	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 333628 JOB CONTROL ACTIVE IN THIS PARTITION
T37 F2 SUB	WAITBND	WAIT FOR I/O OR ECB POST CCB/ECB ADDRESS: 222628 ICCF PSEUDO PARTITION
T38 F2 SUB	READY	READY TO RUN ICCF PSEUDO PARTITION

Figure 71. Example: System Loop Analysis Report (Part 2 of 2)

The Stand-Alone Dump Analysis Routine IJBXSDA

The analysis routine IJBXSDA formats the contents of the SDAID buffer in a stand-alone dump, if SDAID was active when the dump was taken. IJBXSDA prints the trace entries on SYSLST.

Activating the Routine

The routine name IJBXSDA must be contained in the Info/Analysis external routines file before you can call the routine.

Start execution of the analysis routine with the Info/Analysis statements:

```
SELECT DUMP VIEWING  
CALL IJBXSDA
```

Figure 68 on page 191 illustrates the activation of an analysis routine.

Dump Offload

Dump Offload places a dump file that resides on the dump library onto tape for later retrieval. You may choose whether or not to maintain the copy that is on the dump library; the default is to erase the dump. To initiate Dump Offload, use the following statement:

```
▶▶—SELECT DUMP OFFLOAD—◀◀
```

After making this selection, you may offload the current dump by specifying any of the following control statements that are necessary:

- VOLID - specify the output volume and the logical unit number
- BYPASS - skip the write-to-tape operation
- ERASE NO - does not delete the library copy of the dump

Dump Offload is valuable if you need to increase the available online space. Dump Offload does not remove the information about the dump from the dump management file.

An example of a dump Offload job is given in Figure 72 on page 201.

VOLID - Specify Output Volume

Use the VOLID statement to specify the identifier of the output tape.

```
▶▶—VOLID volume_id—┐
                       └SYSnnn┘
```

The *volume_id* is a 6-character alphanumeric value that is added to the entry for the current dump contained in the dump management file. The VOLID statement is required if the dump you are offloading has not been previously unloaded or offloaded.

For subsequent offloads of the same dump, Info/Analysis can retrieve the *volume_id* from the dump management file. To override the saved volume, use the VOLID statement. The most recent VOLID is the one saved in the dump management file.

The logical unit number (SYSnnn) can be assigned to a physical tape address via the job control 'ASSGN' statement. If you do not define a logical unit number allocation is done automatically by the system. The first available unit will be allocated and a message issued for the tape to be mounted on this unit.

BYPASS - Skip Offload

Use the BYPASS statement to free the library space used by the dump without writing the dump to tape. BYPASS is allowed only if a valid offloaded copy of the dump exists; that is, if both of the following conditions are met:

- The current dump has been previously offloaded, but is still in the library, and
- The current dump has not been modified by an analysis routine since it was last offloaded.

```
▶▶—BYPASS—▶▶
```

When BYPASS is processed, &pdm checks for the above conditions. If they are not met, the offload function is not performed.

The DUMP Offload BYPASS statement and the Dump Management DELETE statement differ in the following ways:

- BYPASS checks for a copy of the dump on tape. DELETE does not.
- DELETE removes references to the dump from the dump management file. BYPASS does not.

Thus, use Dump Management with DELETE only if you no longer need a dump. Use Dump Offload with BYPASS if you wish to remove a dump from the dump library but want to maintain a copy on tape and keep information about the dump in the dump management file.

The BYPASS and ERASE NO statements are contradictory and thus mutually exclusive.

ERASE - Delete or Retain Library Copy of Dump

The ERASE statement specifies whether or not Info/Analysis should delete the dump library copy of the dump when doing an offload.



If you want to maintain a copy of the dump on the dump library as well as on tape, specify ERASE NO. **ERASE YES** is the default. Therefore, if you specify ERASE, ERASE YES, or do not specify the ERASE control statement during Dump Offload, Info/Analysis erases the dump from the dump library after a copy is offloaded to tape.

By specifying ERASE NO, you can offload more than one copy of the dump. For example, you may offload a copy of the dump, then run analysis routines, then offload the modified copy. The ERASE statement updates the dump management file with offload information.

BYPASS and ERASE NO are contradictory and thus mutually exclusive statements.

Offloading a Dump to Tape

Figure 72 on page 201 shows how to offload a dump to a tape, erase the copy in the dump sublibrary and update the Info/Analysis dump management file.

<code>// JOB OFFLOAD</code>	
<code>// ASSGN SYS009,280</code>	Assigns tape to receive offloaded dump
<code>// MTC REW,280</code>	Rewind tape to loadpoint
<code>// ASSGN SYSLST,00E</code>	
<code>...</code>	
<code>// EXEC INFOANA,SIZE=300K</code>	
<code>DUMP NAME SYSDUMP.F4.DF400003</code>	Defines the dump name
<code>SELECT DUMP OFFLOAD</code>	Calls OFFLOAD
<code>VOLID T07111 SYS009</code>	Defines the tape volume and optionally the logical unit
<code>RETURN</code>	
<code>SELECT END</code>	
<code>/*</code>	
<code>/&</code>	

Figure 72. Offloading a Dump to Tape

If the tape unit is not defined via the 'ASSGN' and 'VOLID' statements during the OFFLOAD process, &ia searches for a free tape drive and issues a volume mount request message (MOUNT VOLUME volumename ON UNIT xxx ...).

Info/Analysis indicates the successful execution of the OFFLOAD function with a message (DUMP dumpname OFFLOADED...).

SELECT DUMP OFFLOAD versus SELECT DUMP MANAGEMENT DELETE

In comparison to the SELECT DUMP OFFLOAD the SELECT DUMP MANAGEMENT DELETE operation described in "Dump Management" on page 179 has the following functions:

- Erases the dump from the dump sublibrary;
- Erases the information about the dump from the dump management file.

Figure 73 summarizes the functions of the SELECT DUMP OFFLOAD and the SELECT DUMP MANAGEMENT DELETE operation and shows the differences between these two operations.

Dump written to Tape	Information kept in Dump Management File	Dump erased from Sublibrary	Info/Analysis Function
YES	YES	YES	OFFLOAD without additional operands
YES	YES	NO	OFFLOAD with ERASE NO specified
NO	YES	YES	OFFLOAD with BYPASS specified
NO	NO	YES	DELETE

Figure 73. Summary: SELECT DUMP OFFLOAD and DELETE Operation

Dump Onload

Dump Onload copies dumps which reside on tape or disk into the dump library so that they can be further processed by Info/Analysis. To initiate Dump Onload, use the following statement:

```
▶▶—SELECT DUMP ONLOAD—▶▶
```

After making this selection, you may onload the current dump by entering the VOLID and FILE control statements if necessary.

An example of a dump Onload job is given in Figure 74 on page 204.

VOLID - Specify Input Volume

Onloading a Dump from Tape

Use the VOLID statement to specify the volume identifier and (optional) the logical unit number of the tape on which the current dump resides.

```
▶▶—VOLID volume_id—▶▶
      └─SYSnnn─┘
```

The volume_id is a 6-character alphanumeric value that is added to the entry for the current dump contained in the dump management file. The VOLID statement is required if you are onloading a dump for the first time.

For subsequent offloads and onloads of the dump, Info/Analysis retrieves the volume_id from the dump management file. To override the saved value, use the VOLID statement. The most recent VOLID is the one saved in the dump management file.

The logical unit number (SYSnnn) can be assigned to a physical tape address via the job control 'ASSGN' statement. If you do not define a logical unit number allocation is done automatically by the system. The first available unit will be allocated and a message issued for the tape to be mounted on this unit.

Specification of SYSnnn is highly recommended to prevent difficulties in mixed tape environments.

Onloading a Stand-Alone Dump from Disk

Use the VOLID statement with the DISK operand to specify the disk device on which the current dump resides.

▶▶—VOLID DISK SYSnnn—▶▶

DISK indicates that the stand-alone dump is to be copied from the disk device with the logical unit number SYSnnn. Note that SYSnnn is mandatory when onloading a dump from disk.

FILE - Specify Dump File on Multiple-Dump Device

If the tape or disk you are using contains more than one dump file, use the FILE statement to specify the specific dump file you want to onload. In this way, you may onload more than one dump from a tape or disk during a session. Keep in mind that you must leave Dump Onload and specify another dump name before onloading the next file.

▶▶—FILE *file_number*—▶▶

└LAST┐

The default for the FILE statement is "1" if the file statement is omitted. Therefore, if you are onloading a dump from a single file tape/disk or if you are onloading the first file from a multiple-file tape/disk, you need not specify the FILE statement. Also, for the main dump file of a stand-alone dump, the FILE statement need not be specified.

The file number must designate an existing file on the input device. This sequence number is used for searching by Dump Onload when the input file is opened.

When multiple dumps are onloaded during an Info/Analysis session, their file numbers do not have to be in ascending order.

The LAST parameter indicates that this is the last file to be onloaded from the current volume. Specifying LAST de-allocates the device from Info/Analysis.

Loading a Dump into a Dump Sublibrary

Dumps can be stored on tape or disk as output of the following functions:

- DUMP command (tape)
- Stand-alone dump program (tape or disk)
- Info/Analysis Offload operation (tape)

Before you can process these dumps, they have to be onloaded into a dump sublibrary.

The example in Figure 74 shows such an onload job; it stores the dump with the name SYSDUMP.BG.DMPLO3 in the dump sublibrary assigned to the BG partition.

// JOB ONLOAD	
// ASSGN SYS009,281	1. Assigns dump tape to SYS009
// ASSGN SYSLST,00E	
...	
// EXEC INFOANA,SIZE=300K	
DUMP NAME SYSDUMP.BG.DMPLO3	2. Defines dump name
SELECT DUMP ONLOAD	3. Calls ONLOAD
VOLID T03111 SYS009	4. Specifies volume and optional logical unit
FILE 2 LAST	5. Specifies second file
RETURN	
SELECT END	
/*	
/&	

Figure 74. Sample Job: Onload a Dump from Tape into a Dump Sublibrary

The sample job in Figure 74:

1. Assumes that the tape containing the dump is mounted on the tape drive at address 281. Assigns programmer logical unit SYS009 to this drive.
2. Specifies the dump for processing, by name. The dump sublibrary is determined by the dump name.
3. Calls the Info/Analysis DUMP ONLOAD function.
4. Specifies the tape volume on which the dump resides.
The volume name is provided in the list of managed dumps if the dump in question has been offloaded before. See "PRINT - Print List of Managed Dumps" on page 181, for the list function.
If the dump on tape you want to onload has not been offloaded before, the volume id is an identifier of your own choosing. It is used to identify the volume in subsequent dump operations.
5. Defines the file sequence number 2 with the LAST operand in the FILE statement (the dump resides on a multifile tape volume, sequence number 1 is the default value).

When the dump to be onloaded is on a disk extent, make sure that the DLBL/EXTENT statements for the IJSYSDU file are available (see "Dump Program File and Dump Data Set" on page 27). During the ONLOAD process &ia searches for a free tape drive if the tape unit is not defined via the 'ASSGN' and 'VOLID' statements and issues a volume mount request message (MOUNT VOLUME volumename ON UNIT xxx ...).

Info/Analysis indicates the successful execution of the ONLOAD function with a message (DUMP dumpname ONLOADED).

Printing a Dump Stored on Tape or Disk

The subsequent sections describe methods to print dumps that were written to tape or disk, like the stand-alone dump, or to tape, like the DUMP command dump.

The **Info/Analysis program** can be used to format and print these dumps after they have been onloaded to a dump sublibrary. The steps which have to be performed to print the dump are described under “Processing and Printing a Dump with Info/Analysis.”

The **DOSVSDMP utility program** can be used to write dumps in unformatted form directly from the dump device to the printer. This utility program **must** be used to print dumps produced in response to the attention routine command:


```
DUMP BUFFER,cuu
```

These dumps must not be onloaded to the dump sublibrary.

DOSVSDMP can also be used to print other dumps which are too large for the dump sublibrary and therefore cannot be handled by the Info/Analysis program.

Figure 75 shows the different steps which you have to perform depending on the print method you choose.

Dump Printed in Formatted Form with Info/Analysis



Dump Printed in Unformatted Form with DOSVSDMP




Figure 75. Overview: Print a Dump

The following sections describe how to print a dump with Info/Analysis and also how to use the DOSVSDMP utility for the same purpose. Sample jobs are included and the output is explained.

Processing and Printing a Dump with Info/Analysis

This section describes the steps which have to be performed in order to get a printed output of the stand-alone dump or the DUMP command dump. The dump has to be in one of the dump sublibraries before Info/Analysis can be used to print information from it. Therefore the first operation to be done is to onload the dump into one of the dump sublibraries.

Onloading the Dump into the Dump Sublibrary

The following describes in detail the steps which have to be performed.

1. Define the name of the dump.

The name of the dump determines the dump sublibrary. For example:

```
SYSDUMP.BG.DUMPSA2
```

- defines the sublibrary BG as the target library for the subsequent onload process.
- 2. Select the Info/Analysis Onload function.
- 3. Use the VOLID statement to define the tape or disk volume on which the dump resides.
- 4. If the dump resides on tape, mount the dump tape on the device which Info/Analysis requests during processing, or which you defined with the VOLID statement.
- 5. When the dump has been stored in the desired dump sublibrary, Info/Analysis can print the dump.

Printing a Stand-Alone Dump with Info/Analysis

Following are two examples of typical stand-alone dump print jobs and their output.

Sample Job to Print the Main Dump File of a Stand-Alone Dump

The example shown in Figure 76 defines a job which

- Onloads a stand-alone dump with the user defined name SYSDUMP.BG.ADUMPSA2.
- Prints the following on the device assigned to SYSLST:
 - The symptoms of the dump.
 - The formatted areas of the dump.

// JOB PRINT	
// ASSGN SYS009,281	
// ASSGN SYSLST,00E	
...	
// EXEC INFOANA,SIZE=300K	
DUMP NAME SYSDUMP.BG.ADUMPSA2	1. Defines the dump name
SELECT DUMP ONLOAD	2. Selects ONLOAD
VOLID T03111 SYS009	3. Defines the input device
RETURN	
SELECT DUMP SYMPTOMS PRINT DATA	4. Calls the DUMP SYMPTOMS function
RETURN	
SELECT DUMP VIEWING PRINT FORMAT	5. Calls the DUMP VIEWING function
RETURN	
SELECT END /* /&	

Figure 76. Sample Job: Print a Stand-Alone Dump (Main Dump File)

You can also print the dump data using

```
SELECT DUMP VIEWING
PRINT 0 END
```

which does no formatting of the dump. See also “Printing Selective Dump Information” on page 189.

Sample Job to Print an Additional File of a Stand-Alone Dump

The example shown in Figure 77 defines a job which

- Onloads a stand-alone dump with the user defined name SYSDUMP.BG.ADUMPSA4.
- Prints the following on the device assigned to SYSLST:
 - The symptoms of the dump.
 - The formatted areas of the dump.

```
// JOB    PRINT
// ASSGN  SYS009,281
// ASSGN  SYSLST,00E
...
// EXEC  INFOANA,SIZE=300K
DUMP NAME SYSDUMP.BG.ADUMPSA4 1. Defines the dump name
SELECT DUMP ONLOAD              2. Selects ONLOAD
VOLID T03111 SYS009            3. Defines the input device
FILE 4                          4. Selects file 4
RETURN
SELECT DUMP SYMPTOMS           5. Calls the DUMP SYMPTOMS
PRINT DATA                    function
RETURN
SELECT DUMP VIEWING           6. Calls the DUMP VIEWING
PRINT FORMAT                  function
RETURN
SELECT END
/*
/ &
```

Figure 77. Sample Job: Print a Stand-Alone Dump (Additional Dump File)

You can also print the dump data using

```
SELECT DUMP VIEWING
PRINT 0 END
```

which does no formatting of the dump. See also “Printing Selective Dump Information” on page 189.

Printed Output of the Main Dump File of a Stand-Alone Dump

Figure 78 shows the output (symptom part) of the sample job given in Figure 76 on page 206. A list of the formatted dump areas which are printed is given in Figure 79 on page 209.

```
ENVIRONMENT:
CPU MODEL ..... 3090
CPU SERIAL ..... CF5A05
TIME ..... 07:56:09:00
DATE ..... 94/11/17
SYSTEM ID ..... 568606606
RELEASE ..... 1
FEATURE ..... 5C
DUMPTYPE ..... SADUMP
PROBLEM NUMBER ..

REQUIRED SYMPTOMS:

OPTIONAL SYMPTOMS (SDB):

OPTIONAL SYMPTOMS (NON-SDB):
DATE_NOT_AVAILABLE
MACHINE=ESA                               Shows the hardware type.
MODE=PAGING                               Indicates paging.
ACTIVE_SPACE_ID=0                         Shows which address space was active
                                           at the time the dump was taken.
DUMPED_DATA_FROM_SPACE_ID=0              Shows which address space this dump
PMR_ADDRESS_SPACE_ID=00                  data file was dumped from.
DUMPED_DATA=SUPERVISOR+SVA              Shows what data is in this dump file.
```

Figure 78. Sample: Symptom Part of the Stand-Alone Dump Output (Main Dump File)

The following stand-alone dump information is selected and printed.

————— **Diagnosis, Modification or Tuning Information** —————

Formatted Areas of the Stand-Alone Dump:

PSW	Program Status Word (at time of failure)
AREGS	Access Registers
FREGS	Floating Point Registers
GREGS	General Purpose Registers
CREGS	Control Registers
MESSAGE	Error messages and the last 200 messages from the Hard-Copy File

Hexadecimally Displayed Areas of the Stand-Alone Dump Output:

LOWCORE	Low address storage
SYSCOM	System Communication Region
UNATTCB	Re-IPL control block (previous Re-IPL invocation)
UNATTCBN	Hard Wait information (last Re-IPL invocation)
SMCOM	Storage Management Communication Area
CLIM	Class/System Limits Control Block
PCB	Partition Control Block
SCB	Space Control Block
PMRAS	Page Manager Address Space
ASTE	Address Space Number Second Table Entry
PASNAL	Primary Address Space Number Access List
COMREG	Partition Communication Region
PIBTAB	Partition Information Block
PIB2TAB	Partition Information Block Extension
LUBTAB	Logical Unit Block
PUBTAB	Physical Unit Block
PUB2TAB	Physical Unit Block Extension
ERBLOC	Error Recovery Block
CHQTAB	Channel Queue Table
CHNTAB	Channel Control Table
TIBATAB	Task Information Block Address Table
TIB	Task Information Block
TCB	Task Control Block
SAVAREA	Partition Save Areas
ACCREGS	Access Registers
TDSE	Task's Data Space Extension
DUCT	Dispatchable Unit Control Table
DUAL	Dispatchable Unit Access List
DSCB-SCB	Data Space Space Control Block
LPT	Library Pointer Table
LDT	Library Definition Table
SDT	Sublibrary Definition Table
EDT	Extent Definition Table
DDT	Device Definition Table
SDBUFFER	SDAID Buffer

————— **End of Diagnosis, Modification or Tuning Information** —————

Figure 79. Summary of the DUMP VIEWING, PRINT FORMAT Operation Output

DUMP Command Dump Printed with Info/Analysis

This section gives an example of how to print a DUMP command dump that is on tape; it also shows the output of the sample job.

Sample Job to Print a DUMP Command Dump

The example shown in Figure 80 defines a job to print the dump symptoms and the formatted dump areas of a DUMP command dump named SYSDUMP.BG.ADUMPC02 on the device assigned to SYSLST.

An example of the output of this job is shown under Figure 81 on page 211 and Figure 82 on page 211.

```
// JOB PRINT
// ASSGN SYS009,280
// ASSGN SYSLST,00E
...
// EXEC INFOANA,SIZE=300K
DUMP NAME SYSDUMP.BG.ADUMPC02 1. Defines the dump name
SELECT DUMP ONLOAD 2. Selects ONLOAD
VOLID T03111 SYS009 3. Defines the input
RETURN tape volume
SELECT DUMP SYMPTOMS 4. Defines the
PRINT DATA print operation
RETURN
SELECT DUMP VIEWING 5. Defines the
PRINT FORMAT print operation
RETURN
SELECT END
/*
/ &
```

Figure 80. Sample Job: Print the Output of a DUMP Command

Output of the DUMP Command Dump Printed by Info/Analysis

The output of the DUMP command dump printed by Info/Analysis consists of two parts:

- The dump symptom part (example shown in Figure 81 on page 211)
- The formatted dump areas (summary of the areas shown in Figure 82 on page 211).

Figure 80 shows the job with which these two dump parts can be produced.

Output of the DUMP SYMPTOMS, PRINT DATA Operation:

```
ENVIRONMENT:
CPU MODEL .....9121
CPU SERIAL .....110294
TIME .....08:44:51:00
DATE .....94/04/12
SYSTEM ID .....568603206
RELEASE .....1
FEATURE .....5C
DUMP TYPE .....OPRREQ
PROBLEM NUMBER .....
```

REQUIRED SYMPTOMS:

DUMP		which is the command
50000-65000		that requested the
		dump

Figure 81. Sample: Symptom Part of the DUMP Command Dump

Output of the DUMP VIEWING, PRINT FORMAT Operation: Figure 82 gives a list of the dump areas printed.

PSW Program status word at time of failure

AREGS
Access registers

GREGS
General purpose registers at time of failure

FREGS Floating point registers

LOADLS
Phase load list of the partition

Figure 82. Summary: Areas to be Printed with DUMP VIEWING, PRINT FORMAT

Ending the Info/Analysis Job

You end an Info/Analysis job by submitting the `SELECT END` statement while you are at the selection level. The selection level is the point in a sequence after a `RETURN` statement and before a function is selected. If you wish to end your session and are at the function level, enter `RETURN` followed by `SELECT END`. The function level is the point in a sequence after a function is selected and before a `RETURN` statement is entered.

`SELECT END` should be followed by an end-of-input statement (`/*`) and an end-of-job statement (`/&`). If you enter an end-of-input or end-of-job statement at any point in the sequence, the job is canceled at that point. Any valid control statement sequences preceding the end-of-input or end-of-job statement are performed as specified.

Control Statement Sequence Examples

The following are examples of batch execution sequences. Each example describes a possible sequence of functions and presents the control statements to perform those functions.

Each function and the statement that selects that function are labeled with the same letter so that you may make comparisons easily. The example in Figure 83 contains the following operations:

1. Select the Dump Management function and request the printing of the list of managed dumps.
2. On the selection level, specify SYSDUMP.F6.DF600007 as the current dump.
3. Use the Dump Symptoms function to print the dump symptoms that are contained in the symptom record.
4. Use the Dump Viewing function to print selective areas of the dump. The assumed areas are written in the comments on each statement.
5. Use Dump Offload to offload SYSDUMP.F6.DF600007 to the tape with VOLID T02512.
6. End your Info/Analysis session.

```

1. SELECT DUMP MANAGEMENT
   PRINT DATA
   RETURN
2. DUMP NAME SYSDUMP.F6.DF600007
3. SELECT DUMP SYMPTOMS
   PRINT DATA
   RETURN
4. SELECT DUMP VIEWING
   PRINT 0 20880          * PRINT SUPERVISOR DATA
   PRINT C80000 END      * PRINT TO END OF STORAGE
   PRINT FORMAT          * PRINT ALL FORMATTED DATA
   RETURN
5. SELECT DUMP OFFLOAD
   VOLID T02512
   RETURN
6. SELECT END

```

Figure 83. Control Statement Sequence Example

The example in Figure 84 on page 213 contains the following operations:

1. On the selection level, specify SYSDUMP.BG.ADUMPSA6 as the current dump.
2. Use Dump Onload to load the current dump (file 3 on tape T300U1) into the dump sublibrary so that you can work with it.
3. Use Dump Viewing to call routine IJBXDEBUG to analyze the stand-alone dump. Results of the routine are printed together with all formatted data.
4. On the selection level, specify SYSDUMP.BG.ADUMPSA2 as the current dump.
5. Use Dump Offload to offload SYSDUMP.BG.ADUMPSA2, specifying the output volume and choosing to bypass the write operation because a valid copy of the dump already exists on tape. (The information concerning this dump in the dump management file will be kept.)
6. On the selection level, specify SYSDUMP.BG.ADUMPSA7 as the current dump.
7. Use Dump Onload to load the current dump (file 5 on tape T300U1) into a dump sublibrary, specifying LAST because it is the last dump to be onloaded from the tape.
8. End your Info/Analysis session.


```
1. DUMP NAME SYSDUMP.BG.ADUMPSA6
2. SELECT DUMP ONLOAD
   VOLID T300U1
   FILE 003
   RETURN
3. SELECT DUMP VIEWING
   CALL IJBXDEBUG
   PRINT FORMAT
   RETURN
4. DUMP NAME SYSDUMP.BG.ADUMPSA2
5. SELECT DUMP OFFLOAD
   VOLID T03417
   BYPASS
   RETURN
6. DUMP NAME SYSDUMP.BG.ADUMPSA7
7. SELECT DUMP ONLOAD
   VOLID T300U1
   FILE 5 LAST
   RETURN
8. SELECT END
```

Figure 84. Control Statement Sequence Example

Control Statement Summary

This section contains a summary of the control statements for Info/Analysis. The statements are presented in alphabetical order. The “Valid Functions” column represents the functions during which the control statement may be entered as follows:

- M - Dump Management
- S - Dump Symptoms
- V - Dump Viewing
- OF - Dump Offload
- ON - Dump Onload
- SEL - Selection level
- T - Tutorial

Control Statement	Description	Valid Functions					
		M	S	V	OF	ON	SEL
BYPASS	skip offload				X		
CALL routine	call analysis routine			X			
DELETE	delete dump	X					
DUMP NAME dumpname	specify dump	X					X
ERASE YES NO	delete/retain system copy of dump				X		
FILE number LAST	specify dump file					X	
PRINT addr-range FORMAT DATA	print dump data or formatted dump print dump management file	X	X	X			
RETURN	end function	X	X	X	X	X	X
SELECT function END	select function						X
UTILITY	initialize dump management file	X					
VOLID volume-id SYSnnn	specify input or output tape				X	X	
VOLID DISK SYSnnn	specify input disk					X	
<p>In the PRINT command, addr-range can be:</p> <p>from-addr from-addr to-addr from-addr END from-addr FOR length</p> <p>In the SELECT command, function can be:</p> <p>DUMP MANAGEMENT DUMP SYMPTOMS DUMP VIEWING DUMP OFFLOAD DUMP ONLOAD</p>							

Part 5. Appendixes

Appendix A. Symptom Records Overview

The symptom records contain a collection of failure-related symptoms in a standard format. At the time of problem detection, the failing component creates the symptom records. Subsequently, analysis routines may be run to collect additional symptoms and may add them to the symptom records. The ultimate goal of the symptom records is to reduce the amount of time necessary to analyze a dump.

Symptom Records Structure

The symptom records have six sections. Figure 85 shows an overview of the symptom records contents and the information used in a dump.

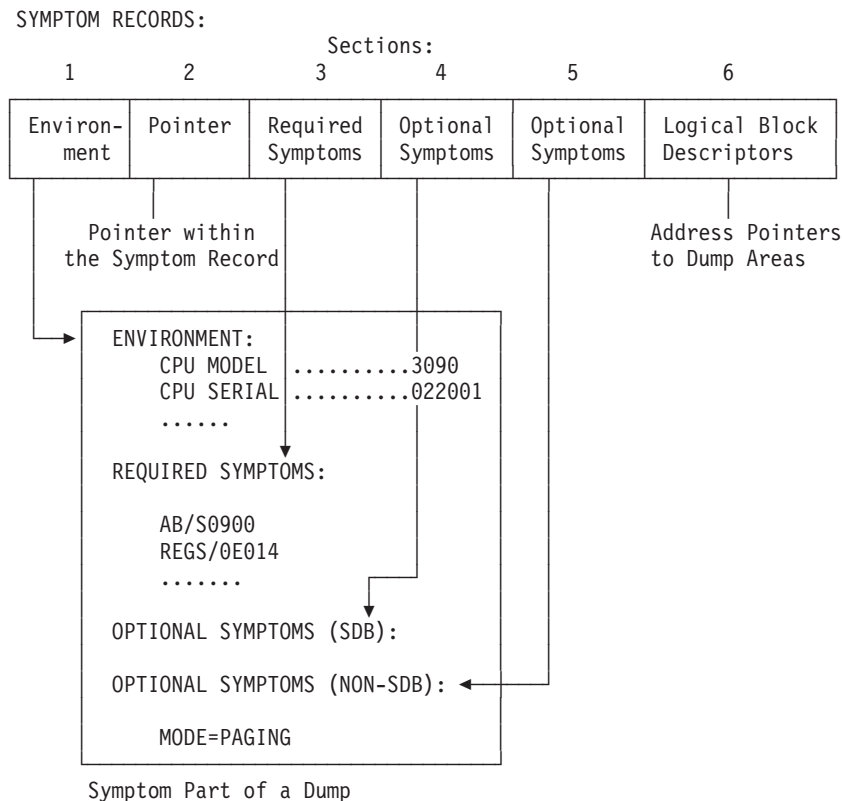


Figure 85. Symptom Records

The sections of the symptom records are:

1. The environment section, which describes the operating system at the time the problem is detected.
2. The pointer section, which describes the offsets of each of the other sections from the beginning of the record and the length of each section. This data is used by the analysis tool and is not accessible to you.
3. The required symptoms section, which contains symptoms considered essential to your dump analysis tasks. Required symptoms are in the structured data base (SDB) format.

Symptom Records

4. The optional symptoms (SDB) section, which contains additional symptoms in the SDB format.
5. The optional symptoms (non-SDB) section, which contains symptoms that do not conform to the SDB format but provide failure-related information.
6. The control block section, which contains descriptions and locations of control blocks that are necessary for problem analysis. This section also contains text and hexadecimal entries that may be related to the problem and descriptions of control block chains and arrays.

The contents of Sections 1, 3, 4, and 5 is displayed when you select Dump Symptoms. This information can help you determine the nature of the problem and where it occurred.

The contents of Section 6 is used by &ia to display dump information when you select Dump Viewing. The control block data presented via the Section 6 entries can help you determine why the problem occurred. See the discussion of Section 6 below.

Note: The structured data base format is used to standardize problem data so that searches for duplicate problems in the data base of existing problems used by customer engineering are accurate.

Symptom Record Creation

The symptom records are built when:

- A system component detects an error that may or may not result in a dump. The dumping component builds the symptom records, completing the required symptoms section and as much optional data as possible. The dumping component then calls a system dumping routine that fills out the environment section, merges it with the rest of the symptom records, and possibly takes a dump.
- A stand-alone dump is taken. The environment section is completed by the dump program.

Section 6

Section 6 acts as a table of contents for dump data. You can use it to locate certain control blocks without having to manually follow pointers through the dump. The component which originates the dump is responsible for providing information about the control blocks that are pertinent to the error.

The section provides for example:

- Names and locations of dumped control block storage.
- Descriptions of conditions present at the time of the dump.
- Hexadecimal data that may not be contained in the dumped storage (such as registers).
- Algorithms for control block relationships.
- The names of control block fields.

The location of Section 6 entries is in separate records of the dump that are classified as symptom record extensions.

At dump time, the failing component may designate the control blocks that are suspected of being in error or are necessary for problem determination. Ideally, the

dump includes only the storage in use when the error occurred. Host system storage that provides pointers to the component address space or partition may not be needed. This practice reduces the volume of dump output.

The failing component may include descriptions of the related control blocks in Section 6. Keep in mind that the data for the control blocks is within the main body of the dump. The information in Section 6 describes the addressing method, the content, the format, and the chaining structure of these related dump areas.

Each of these descriptive entries in Section 6 is called a locating block descriptor (LBD). LBDs come in a variety of forms to describe the structure and relationships of control blocks. Each LBD consists of a header segment and, optionally, a variable segment. The header identifies the data being described by the LBD by providing a name, its length, and usually, its location in storage. The optional portion may be:

- One or more extensions
- A formatting descriptor
- A linkage descriptor

By including a variety of LBDs, Section 6 becomes a table of contents for dump data. Through the Dump Viewing function of Info/Analysis, you can use Section 6 to analyze the dump.

Locators

If a header portion of an LBD provides a control block address, a specific instance of a control block has been identified. An LBD may still be a locator if the address is not provided. This would be the case for a linkage descriptor, for example.

A locator may be simple or complex. A simple locator names a control block and defines its length, address qualifications, and other pertinent information. A simple locator is appropriate if there is one occurrence of a particular control block in a dump or in a linkage. A complex locator consists of a header portion resembling a simple locator and one or more extensions that:

- Provide a way to find all occurrences of the control block that is defined in the header.
- Associate additional data with the control block defined in the header.

There are five kinds of locator extensions:

- Chain extension - describes a string of all occurrences of one type of control block in a dump. If you select a linkage descriptor while viewing the analysis summary main display, and the locator for a particular type of control block has a chain extension, Info/Analysis uses the locator and its chain extension to display the chain of control blocks of that type.
- Array Extension - describes contiguous occurrences of one type of control block in a dump. If you select a linkage descriptor while viewing the analysis summary main display, and the locator for a particular type of control block has an array extension, Info/Analysis uses the locator and its array extension to display the array of control blocks of that type.
- Text Extension - contains character data added to the dump by the dumping component. The text is associated with the name defined in the locator's header. You may view text for which the header merely provides a name by selecting this text entry from the analysis summary main display. Alternatively, the header may describe a control block with which the text is associated. If you select a

Symptom Records


locator while viewing the analysis summary LOCATORS display, and the locator by that name has a text extension, the lines of text are displayed along with the control block data.

- Hexadecimal Extension - contains hexadecimal information such as register contents or storage data areas that may pertain to the software problem that caused the dump. This data is collected and associated with the name defined in the header portion of the locator. You may view hexadecimal data for which the header merely provides a name by selecting the hexadecimal entry from the analysis summary main display.
- Keyfield Extension - contains the location of a particular field (such as a completion code) that is significant or pertinent to the associated control block identified in the header portion of the locator. This field may be in (or apply to) the control block. According to the status of this field, you can decide whether or not to continue to examine the control block. The keyfield is included in an analysis summary linkage or LOCATORS display.

Linkage Descriptors

A linkage descriptor is a special type of extension to a header that defines the relationships for a set of control blocks. That is, the control block named in the linkage is used to locate one or more control blocks that are logically related.

Linkage descriptors are used in conjunction with locators. Together, they enable Info/Analysis to display an ordered list of control blocks. For example, a component might create a linkage descriptor that defines control block A pointing to B and a chain of Cs, and each C pointing to a chain of Ds. When you select the linkage entry for that component while viewing the analysis summary main display of Dump Viewing, the displayed data depicts the following situation:



Info/Analysis shows the linkages by indenting control block names. To create the display, Info/Analysis uses both locators with their extensions, if any, and a linkage descriptor.

Formatting Descriptors

A formatting descriptor is a special type of extension to a header that defines a set of simple formatting instructions for a control block. The header portion of the formatting descriptor cannot be used to locate a control block in the dump.

A formatting descriptor for a control block maps out some or all of the fields of the control block named in the associated locator, their offsets from the beginning of the control block, and their lengths. When you select a control block for display while in the formatted mode of dump display, Info/Analysis displays the contents of each field, one or more per line. Depending on the options you have set, the field labels and offsets may appear with the data. To create such a display, Info/Analysis uses both a locator and a formatting descriptor.

Appendix B. Other Diagnosis Tools

This appendix describes various commands and facilities to process information relevant for problem analysis.

ACTION: Print Linkage Editor Map

»—ACTION MAP—«

To obtain a linkage editor map for a program, specify, for the program's linkage editor run, the linkage editor control statement ACTION MAP.

Figure 86 on page 224 shows a sample output from this routine.

Linkage Editor Map Warning Messages

The following messages may be included in the map output, except when NOMAP was specified in the ACTION statement for the linkage editor run.

Root structure overlaid by succeeding phase

When this message appears, "OVERROOT" is printed to the left of the phase name that overlays the root phase.

Possible invalid entry point duplication in input

An entry label appeared at least twice in the input. At the second (or later) appearance it was not possible to validate it as being a true duplication. The most common reason for this message is sub-modular structure with (source) entry labels defined before the CSECT in which the entry point appears.

Invalid transfer label on end or entry statement ignored

An overriding transfer label in the entry statement was not defined within the first phase, or a transfer label was not defined in an end statement in its own module.

Control sections of zero length in input

The COBOL, FORTRAN, RPG*, AND PL/1 (D) compilers do not supply all of the information required by the linkage editor in the ESD records. Specifically, the control section length is provided in the end record. If a control section defined in the ESD information has a length of zero, it normally indicates that the length is to appear in the end record. It is possible to generate zero-length control sections through Assembler. Such a condition produces this message. This is not an invalid condition if it is not the last control section that is of zero length. If the last control section is of zero length, the length is implied to be in the end record and, if not present, causes an error condition.

Unresolved external references

These labels indicate external references that cannot be matched with a corresponding entry point. ESD items from unused control sections may also cause this message.

Other Diagnosis Tools

```

// JOB LINK MODABCD                                DATE 03/30/94,CLOCK 14/12/11
// LIBDEF PHASE,CATALOG=IJSYSRS.SYSLIB
// LIBDEF OBJ,SEARCH=(IJSYSRS.GL)
// OPTION CATAL
  INCLUDE MODLNK
  PHASE PHASE2,*,SVA
  MODE AMODE(31),RMODE(24)
  ACTION ERRLMT(1000),XXXX
  INCLUDE MODC
  ENTRY CS1B
// EXEC LNKEDT,PARM='AMODE=ANY'
JOB LINK 03/30/94 5686-066-06-15C-0 LINKAGE EDITOR DIAGNOSTIC OF INPUT
INVOCATION PARAMETERS: AMODE=ANY RMODE=24 (A)
ACTION TAKEN MAP (B)
  INCLUDE MODLNK (C)
    ** MODULE MODLNK 94-03-30 11.30 94-03-30 12.36 INCLUDED FROM IJSYSRS .GL VOLID=ESA132
-----
  PHASE PHASE1,* (C, G)
-----
  INCLUDE MODA (C) 0000001
    ** MODULE MODA 94-03-30 07.51 94-03-30 12.14 INCLUDED FROM IJSYSRS .GL VOLID=ESA132
  INCLUDE MODB (C) 0000001
    ** MODULE MODB 94-03-30 07.59 94-03-30 13.56 INCLUDED FROM IJSYSRS .GL VOLID=ESA132
-----
  PHASE PHASE2,*,SVA (C, G)
-----
  MODE AMODE(31),RMODE(24)
  ACTION TAKEN ERRLMT=1000
2135I INVALID OPERAND IN ACTION STATEMENT (K)
  ACTION ERRLMT(1000),XXXX (D) (E) (F)
  INCLUDE MODC (C)
    ** MODULE MODC 94-03-30 08.05 94-03-30 09.20 INCLUDED FROM IJSYSRS .GL VOLID=ESA132
    ** MODULE MODD 94-03-30 08.06 94-03-30 12.13 AUTOLNKD FROM IJSYSRS .GL VOLID=ESA132
-----
  ENTRY CS1B
03/30/94 COMMON AREAS:
      NAME      LOADED AT      LENGTH
      COM1A     2AB078           8
      COM1B     2AB080          280
      COM1B     2AB300           20
      COM2B     2AB320           68
03/30/94 PSEUDOREGISTERS:
      NAME      ORIGIN      LENGTH
      PSEU1     0            19
      PSEU2     20           78
TOTAL LENGTH OF PSEUDOREGISTERS:
      PSEU2     20           98

```

Figure 86. Sample: Linkage Editor Output (ACTION MAP) (Part 1 of 2)

- (A) Possible invocation parameters on the PARM field of the EXEC LNKEDT statement are: MSHP, AMODE, RMODE.
- (B) Option MAP is default if SYSLST is assigned.
- (C) Listing of control statements as submitted to linkage editor. (From Job Control or an included module.)
- (D) Date and time the module has been cataloged the first time.
- (E) Date and time of last update.
- (F) Sublibrary from where the module is included or SYSLNK.
- (G) Phase statement. This statement defines the phase name, the load address (for example * to indicate relocatable) and for example, whether the phase has to be SVA eligible or the AUTOLINK feature has to be deactivated. The named phase is composed of the subsequent included modules.
- (H) Module autolinked, based on unresolved external reference 'MODD' in Phase 'PHASE2'.
- (I) List of named and unnamed Common Control Sections with name, load address, and length.
- (J) List of Pseudo Registers (External Dummy Sections) with name, origin (displacement within PR pool), and length. The total (cumulative) length indicates the amount of storage to be allocated during execution.
- (K) Error message from invalid ACTION statement.

(A)	(B)	(C)	(D)	(E)	(F)	(G)	(H)	(I)	(J)	(K)	(L)		
03/30/94	PHASE	XFR-AD	LOCORE	HICORE	CSECT/ ENTRY	LOADED AT	RELOC. FACTOR	PARTIT. OFFSET	PHASE OFFSET	TAKEN FROM	AMODE/RMODE		
-----										*P	ANY	24	RELOCATABLE

PHASE1	2AB3A8	2AB388	2AB3C3		CS1A	2AB388	2AB388	000310	000000	MODA	24	24	
					*LAB1A	2AB388							
					+LAB1B	2AB38C							
					CS1M	2AB398	2AB388	000320	000010	MODA	24	24	
					+LAB1M	2AB398							
					+LAB1N	2AB399							
					+LAB1P	2AB39A							
					CS1N	2AB3A0	2AB388	000328	000018	MODA	24	24	
					*LAB1Q	2AB3A0							
					+LAB1R	2AB3A1							
					*LAB1S	2AB3A2							
					CS1B	2AB3A8	2AB3A8	000330	000020	MODB	24	24	
					*LAB1C	2AB3AC							

PHASE2	2AB3C8	2AB3C8	2AB3F3		CS2A	2AB3C8	2AB3C8	000350	000000	MODC	24	24	
					*LAB2A	2AB3C8							
					+LAB2B	2AB3CC							
					MODD	2AB3D8	2AB3D8	000360	000010	MODD	24	24	
					*LAB3A	2AB3D8							
					+LAB3B	2AB3DC							

UNRESOLVED EXTERNAL REFERENCES
 UNRESOLVED ADCON AT OFFSET 002AB3C0
 001 UNRESOLVED ADDRESS CONSTANTS
 PHASE(S) CATALOGED INTO SUBLIBRARY IJSYSRS.SYSLIB VOLID= ESA132
 1S55I LAST RETURN CODE WAS 0004
 EOJ LINK MAX.RETURN CODE=0004

EXTRN LAB1T] >(M)

DATE 03/30/94,CLOCK 14/12/16,DURATION 00/00/04

- (A) Name of each phase.
- (B) Address where the phase is transferred to.
- (C) Lowest and highest virtual storage location of the phase.
- (D) Labels of all CSECTs which establish the phase in ascending order followed by the CSECT's entry labels. + indicates an entry label, which was referenced and * indicates an entry label which was not referenced.
- (E) CSECT load address / ENTRY address.
- (F) Difference between the start of virtual storage and the assembled CSECT start address.
- (G) Offset from the partition begin plus save area length to the CSECTs start location.
- (H) Offset from the phase begin to the CSECTs start location.
- (I) Name of the module from which the CSECT is taken.
- (J) Contains either *M or *P (or blank) of the phase:

*P indicates that the AMODE/RMODE assigned from ESD data is overridden by values from the PARM field of the EXEC LNKEDT control statement.

*M indicates that the AMODE/RMODE assigned from ESD data or from the PARM field is overridden by values from the MODE control statement.

The field is left blank, if neither the PARM field nor a MODE statement specifies AMODE/RMODE.

- (K) Contains AMODE and RMODE of phases and CSECTs.
- (L) Indicates loading characteristics of a phase.
- (M) Warning messages related to unresolved external references.

Figure 86. Sample: Linkage Editor Output (ACTION MAP) (Part 2 of 2)

DITTO: Dump a Disk, Diskette, or Tape



IBM Data Interfile Transfer, Testing and Operations/ESA for VSE (DITTO/ESA for VSE), an IBM program product, is a useful tool for the recovery of data that may have become inaccessible by VSE programs.

Other Diagnosis Tools


You can use DITTO/ESA for VSE to print, on a line printer, data stored on disk, diskette, or tape. The program dumps the data either in character-only format or in the character and vertical hexadecimal format.

For more detailed information about DITTO/ESA for VSE, refer to the *IBM DITTO/ESA for VSE Introducing* manual.

DSPLY/ALTER: Display or Alter Storage

The DSPLY command allows the console operator to display 16 bytes of virtual storage starting at the specified hexadecimal address on the device assigned to SYSLOG. Two characters (0-9, A-F) appear on SYSLOG for each byte of information; these characters represent the hexadecimal equivalent of the current information in virtual storage. You can alter this information either by the **ALTER** operator command or by using the hardware storage display feature as described under "Hardware Aids via the Operator Console" on page 235.

To request a **display** of storage, enter the command:



To **alter**, enter:



space-id

Indicates in which address space the specified address is to be displayed or altered. Valid specifications are:
0 through 9, A, B, R or S.

To display virtual storage in a **shared** area specify the space-id of any **existing** virtual address space.

BG, Fn

Indicates in which static or dynamic partition the specified address is to be displayed or altered. part can specify any of the static partitions BG, F1 through FB or a partition within a dynamic class, for example, P1.

address

Specifies the hexadecimal address at which the storage display or alteration is to start.

Figure 87 on page 227 shows an example of using the ALTER and DSPLY commands.

```

DSPLY 1,300
AR 015 90F21028 18215811 00185801 0014F9F9 *.2.....99*
AR 015 1I40I  READY

ALTER 1,300
AR 015 1I42D  ADDRESS WITHIN SUPERVISOR OR SVA
AR+015
15 IGNORE
AR 015 OLD DATA: 90F21028 18215811 00185801 0014F9F9 *.2.....99*
AR 015 ENTER HEX DATA (1-16 BYTES)
AR+015
15 FFF2
AR 015 1I40I

DSPLY 1,300
AR 015 FFF21028 18215811 00185801 0014F9F9 *.2.....99*
AR 015 1I40I  READY
    
```

Figure 87. Sample of the DSPLY and ALTER Commands

The example above shows the commands to display the contents of address X'300' in space 1 and to alter two bytes beginning with the same address to X'FFF2'. Message 1I42D is not displayed if the area you want to alter is in the user area.

For a detailed description of the ALTER or DSPLY command see *z/VSE System Control Statements*.

Restriction:


If the specified address is within an invalid address area, the command is ignored and a corresponding information message is issued.

If the 16 bytes to be displayed cross the boundary from a valid to an invalid address area, only the bytes in the valid address area are displayed, and a corresponding information message is issued.

Invalid addresses are:

- Locations beyond the end of virtual storage.
- Unused (not allocated) partition GETVIS space.
- A location in the page pool.
- A location within an unallocated area of the virtual address space.
- A location in the partition's virtual address area when a program in that partition is being executed in real mode, or vice versa.

LIBLIST: Display Library Chains



Library search chains established with the job control LIBDEF statement can be displayed either on the system console or on SYSLST with the LIBLIST job control statement.

- libtype = Corresponds to the type operand of the LIBDEF statement
- * = specifies that library definitions of all LIBDEF statements (except DUMP) are to be

Other Diagnosis Tools

displayed.

partition-id = static or dynamic partition whose library chains are to be listed.
* = The library chains of the partition processing the statement are to be listed (default).

SYSxxx = SYSST or SYSLOG device for output.
Default = SYSLOG if entered from SYSLOG,
= SYSST if entered from SYSRDR.

Figure 88 is an example of a listing of BG partition's active library search chains resulting from a LIBLIST command.

```
// LIBLIST *,BG
TYPE: PHASE
      BG-TEMP ** NO LIBRARY INFORMATION AVAILABLE
      BG-PERM LIBNAME SUBLIB      STATUS -PARTITIONS- +DYNPARTS
      SEARCH  PRVLIB  TCLIB          01 34
              IJSYSRS SYSLIB        0123456789AB DYNP
      CATALOG PRVLIB1 TCLIB          0 23
TYPE: OBJ
      BG-TEMP ** NO LIBRARY INFORMATION AVAILABLE
      BG-PERM LIBNAME SUBLIB      STATUS -PARTITIONS-
      SEARCH  PRVLIB  TCLIB          01 34
              IJSYSRS SYSLIB        0123456789AB DYNP
              PRVLIBS SLIB2          SEC SHR 0 4
TYPE: SOURCE
      BG-TEMP ** NO LIBRARY INFORMATION AVAILABLE
      BG-PERM LIBNAME SUBLIB      STATUS -PARTITIONS-
      SEARCH  PRVLIB1 TCLIB          01 34
              IJSYSRS SYSLIB        0123456789AB DYNP
              SERVLIB S1$XE8          0
TYPE: PROC
      BG-TEMP ** NO LIBRARY INFORMATION AVAILABLE
      BG-PERM LIBNAME SUBLIB      STATUS -PARTITIONS-
      SEARCH  PRVLIB1 TCLIB          01 34
              IJSYSRS SYSLIB        0123456789AB DYNP
```

No temporary search chain is defined in the above example.

The keyword DYNPARTS means that at least one dynamic partition has a LIBDEF to the sublibrary specified.

The device on which PRVLIBS SLIB2 resides is shared by two or more CPUs – indicated by SHR in the STATUS column. SLIB2 also is a secured (SEC) sublibrary.

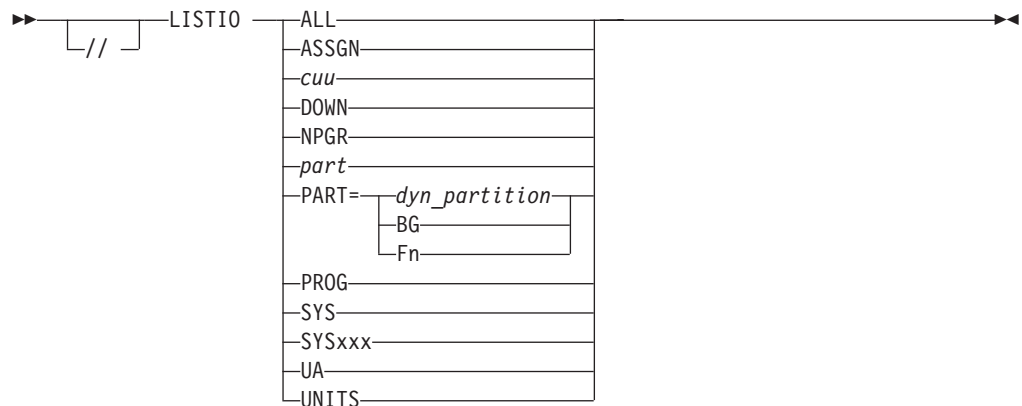
Figure 88. Example: Library Chain Listing

LIST: Print Language Translator Source Code

▶▶ `// OPTION LIST` ▶▶

Normally, a language translator source listing is requested for all language translator runs by specifying option **LIST** in the **STDOPT** command as part of the IPL procedure. Should **NOLIST** have been specified in that command, you can override it by including `// OPTION LIST` in the job control statements for your language-translator run.

LISTIO: List I/O Device Assignments



You can list I/O device assignments before execution of the program begins by inserting the job control **LISTIO** statement or command in the program. The list will appear on **SYSLOG** if you insert the **LISTIO** command without `//`, or on **SYSLST** with `//`. The listing represents the device assignment status at the time the statement or command is being executed and not when an error occurred during a previous run.

For a detailed description of the operands of the **LISTIO** command, refer to *z/VSE System Control Statements*.

LISTLOG: Display Console Communication

▶▶ `// EXEC LISTLOG` ▶▶

You can use the **LISTLOG** utility program to request a listing of all information collected for a specific job in the hard copy file. The program writes this list to the device assigned to **SYSLST**.

You invoke the program by inserting the statement `// EXEC LISTLOG` immediately following the `/&` statement for the job. *z/VSE* invokes the program automatically whenever a job is canceled.

The printout provided by the **LISTLOG** utility program lists:

- Job control statements submitted for the job,
- All messages displayed on the console for this job,

Other Diagnosis Tools

- Any attention routine messages and commands that occurred while the job was being executed,
- Operator responses.

LOG: Print Job Control Statements

The LOG command or statement causes the system to log all job control commands and statements.

```
▶▶ // LOG
```

```
▶▶ // OPTION LOG
```

The // LOG statement has the same effect as the // OPTION LOG statement.

```
▶▶ // NOLOG
```

The two statements are interchangeable, and each can be reset by either the // NOLOG or the // OPTION NOLOG statement.

For a detailed description of the LOG statement, refer to *z/VSE System Control Statements*.

LSERV: Display Label Information Area

LSERV, a system utility program, produces a printout of the label information area on the device assigned to SYSLST.

The job control statement for LSERV is as follows:

```
▶▶ // EXEC LSERV [ , PARM='STDLABEL  
PARSTD  
PARSTD=syslogid  
CLASSTD  
CLASSTD=class  
ALL  
FREE' ]
```

The following is a description of the parameters used in the job control statement:

STDLABEL

prints the system standard labels only.

PARSTD

prints the partition standard labels of all active partitions.

PARSTD=syslogid

prints the partition standard labels of the specified partition only.

CLASSTD

prints all class standard labels only.

CLASSTD=class

prints the class standard labels of the specified class only.

ALL prints all labels, including free-usage labels. In addition, label information for data secured files (see also the DSF operand in the DLBL statement) is displayed.

FREE prints all free-usage labels only.

If no parameter is specified, all labels (but not DSF labels and free-usage labels) are printed. User labels from static or dynamic partitions are only included when no parameter or ALL is specified. User labels change from job to job and, therefore, no special support is needed in the LSERV program.

A sample, partial output of an LSERV run for the above control statements, is shown in the example in Figure 89 on page 232. The output shows label-information records contained in different label groups, such as system standard labels, class standard labels, or partition labels. Each active partition (static or dynamic) can establish up to three label groups: free-usage, temporary, and permanent partition labels.

The output shows the relationship between job control DLBL and EXTENT statements. For further information, refer to the *z/VSE Guide to System Functions*, SC33-8233.

For VSAM files only

There is an additional label information record following the VSAM label record if, in the DLBL statement, at least one of the operands DISP, RECORDS, or RECSIZE is specified.

Note: A warning message is issued on SYSLST if you request LSERV while another partition is updating the label area. The free-usage label groups which are used internally by VSE/ICCF and vendor software, are only shown with PARM= 'ALL' or PARM= 'FREE'.

Other Diagnosis Tools

```
                                LABEL INFORMATION DISPLAY  PAGE nnn
EXAMPLE
  FILE IDENTIFIER                EXAMPLE
  FILE SERIAL NUMBER            OMITTED
  VOLUME SEQUENCE NUMBER        01
  CREATION DATE                 OMITTED
  RETENTION PERIOD (DAYS)       0999
  FILE TYPE                     SEQUENTIAL

  EXTENT INFORMATION
  EXTENT SEQUENCE NUMBER        000
  EXTENT TYPE                   1 (PRIME DATA)
  RELATIVE START ADDRESS        002
  NUMBER OF TRACKS/BLOCKS       045107
  SYMBOLIC UNIT                 SYSRES LOGICAL UNIT FORMAT
  VOLUME SERIAL NUMBER          OMITTED
                                TYP=00,NUM=06

  ADDITIONAL INFORMATION
  DISPOSITION                   (OLD,KEEP)
  RECORDS                       (0000000500,0000000100)
  RECORD SIZE                   0000000080

SALARY
  FILE IDENTIFIER                SALARY.1999.FILE
  FILE SERIAL NUMBER            DASD02
  VOLUME SEQUENCE NUMBER        01
  CREATION DATE                 OMITTED
  EXPIRATION DATE               90/365
  FILE TYPE                     SEQUENTIAL

  EXTENT INFORMATION
  EXTENT SEQUENCE NUMBER        000
  EXTENT TYPE                   1 (PRIME DATA)
  RELATIVE START ADDRESS IN TRACKS/BLOCKS 010000
  NUMBER OF TRACKS/BLOCKS       001000
  SYMBOLIC UNIT                 SYS019 LOGICAL UNIT FORMAT
  VOLUME SERIAL NUMBER          DASD02
                                TYP=01,NUM=13
```

Figure 89. Sample: LSERV Output

LVTOC: Display Volume Table of Contents

```
// ASSGN SYS004,cuu
// ASSGN SYS005,cuu
// EXEC LVTOC
```

A volume table of contents (VTOC) is an index of all files, and the remaining space, on a disk volume. A VTOC display can be requested by executing the LVTOC program with SYS004 assigned to the applicable disk drive and SYS005 to a printer. LVTOC lists the file labels contained in a VTOC in alphabetic sequence by file name. It also provides a listing of free space on the volume, with the start and end addresses and sizes of the unused space. The control statements needed to invoke that program may be submitted via SYSRDR or via the console as shown in Figure 90 on page 233.

A display of a VTOC can be requested also in response to messages. Such a response is CANCELV or DSPLYV. Use CANCELV if you intend to cancel the job,

or DSPLYV if the condition allows program execution to be continued after the VTOC display.

Submission via SYSRDR	Submission via the console
// JOB anyname	1. Press the Request key
// ASSGN SYS004, cuu (A)	2. Enter:
// ASSGN SYS005, cuu (B)	PAUSE p.-id, E0J (C)
// EXEC LVTOC	3. Wait for end-of-job in the specified partition.
/&	4. Enter:
	// ASSGN SYS004, cuu (A)
	// ASSGN SYS005, cuu (B)
	// EXEC LVTOC

(A) The unit address of the disk for which the VTOC is desired.
 (B) The unit address of the device on which the VTOC is to be listed (normally a line printer).
 (C) The identifier of the partition (F1 ...) in which LVTOC is to run.

Figure 90. Control Statements to Invoke LVTOC

STOP/PAUSE: Suspend Program Execution

STOP/PAUSE

Suspending program execution between job steps can be of much help during hands-on diagnosis.

You suspend program execution with the **STOP** command either via the console or, if you use a card reader, via SYSRDR. Another possibility is to submit the job control **PAUSE** statement or command. Both methods result in program execution to be suspended when job control executes the statement or command. The **PAUSE** command is used to interrupt the execution of the job. The operator may enter additional job control statements via SYSLOG at this time. The **STOP** command removes the partition from the system's task selection mechanism and no read is issued to the SYSLOG or SYSRDR device for that partition.

To resume program execution after a **STOP** command, issue an attention routine **START** command for the partition. To resume program execution after a **PAUSE** statement or command, simply press **END/ENTER**. Note that the **STOP** and **START** commands can be given in a static partition only. For dynamic partitions, use the **CANCEL** or **VSE/POWER PFLUSH** command.

Appendix C. Hardware Service Aids

Controlling the Recovery Management Support

The recording activity of RMS can be controlled via the operator command ROD.

Use the **ROD** command to:

- Add error statistics to the system recorder file.
- Have RMS write MDR records into the SYSREC file for those devices that are equipped with an internal error log.
- Have RMS build an end-of-day (EOD) record and write this record on SYSREC.
- Write the hardcopy buffer into the hardcopy file.

The ROD command is discussed in more detail in the following section.

The ROD Command

▶▶—ROD—▶▶

The ROD command has no operands. Issuing the ROD command causes the hardcopy buffer written to the hardcopy file and RMS to record, on SYSREC:

- Error statistics that were compiled for I/O devices (except telecommunication devices).
- An end-of-day record if RMS received an appropriate response to a prompting message via the console.

Retrieval and Analysis of RMS Information

The EREP Program

For the retrieval of information recorded by RMS on SYSREC, use the IBM EREP program. How to use this program is described in the separate publications, *EREP User's Guide* and *EREP Reference*.

For a number of VSE messages, the recommended response in *z/VSE Messages and Codes* includes instructions to run EREP.

Hardware Aids via the Operator Console

Current IBM processors provide a variety of hardware aids for hands-on diagnosis. The procedures for using these aids are, for the most part, processor-model dependent, and are described in detail in the operating procedures manuals for these processors. Therefore, this section discusses only aspects such as usefulness of the aids, when to use them, and requirements or precautions for their use.

The most important hardware aids available via the operator's console for system service and program diagnosis are:

- Alter/display feature.
- Instruction stepping feature.
- Stop on address compare feature.

Hardware Aids

CAUTION:

When using one of the above-mentioned hardware serviceability and debugging aids, you interfere with normal processing under VSE. Therefore, you should consider using these aids only (with your local management approval) in situations such as total system failure or a hard wait condition with no VSE-supported recovery possible.

Hardware Alter/Display

With the alter/display feature you can display the contents of storage areas and registers as indicated below. You can also alter any of these storage areas.

Note: The alter/display feature can be used only from the operator's console of your processor; the feature is not available, for example, from a channel-attached IBM 3277 that you use as an operator's console.

Following is a list of storage areas (and registers) that you can display and alter by using this hardware aid:

- Any selected area of real or virtual storage.
- Contents of the general purpose registers.
- Contents of the floating point registers.
- Contents of the control registers.
- Current PSW.
- Storage protection key.

For detailed information on how to use this feature and on the areas that you can display or alter from your processor's console, refer to IBM's operating procedures manual for your central processor.

Instruction Stepping Feature

With the instruction stepping feature you can check and record the address of each instruction that is executed during program operation. By combined application of this feature and the alter/display feature, you can trace, for example, a short program loop. This approach of tracing executable code of a program is indicated when only short sections of code are to be traced or if, for any reason, the SDAID tracing facility cannot be used.

Refer to the operations manual for your processor for details of this feature.

Stop-on-Address-Compare Feature

This feature is provided primarily for IBM service personnel. It enables one, for example, to stop all system activity at a selected instruction address within a program. In combination with the alter/display feature (or commands ALTER, DSPLY, or DUMP), the stop-on-address-compare feature can be used to display or alter the contents of storage at this selected address. The feature can be used, for example, if there is a need for a dump of a specific area of virtual storage at a specific point of program execution.

Another use of this feature is the generation of a sync signal at a certain instruction address (this is primarily a hardware service aid).

For more information refer to the operating procedures manual for your processor.

Glossary

This glossary defines technical terms and abbreviations used in the *z/VSE Diagnosis Tools*. If you do not find the term you are looking for, view the *IBM Dictionary of Computing* located at: www.ibm.com/ibm/terminology.

The glossary includes definitions with symbol * where there is a one-to-one copy from the IBM Dictionary of Computing.

* **abend.** 1. Abnormal end of task. 2. Synonym for *abnormal termination*.

access control. A function of VSE that ensures that the system and the data and programs stored in it can be accessed only by authorized users in authorized ways.

access method. A program, that is, a set of commands (macros), to define files or addresses and to move data to and from them; for example VSE/VSAM or VTAM.

access register (AR). A hardware register that a program uses to identify an address space or a data space. Each processor has 16 ARs, numbered 0 through 15, which are paired one-to-one with the 16 general-purpose registers (GPRs).

* **ACF.** Advanced Communications Function.

ACF/VTAM. See *VTAM*.

address space. A range of up to two gigabytes of contiguous virtual storage addresses that the system creates for a user. Unlike a data space, an address space contains user data and programs, as well as system data and programs, some of which are common to all address spaces. Instructions execute in an address space (not a data space). Contrast with data space.

addressing mode (AMODE). A program attribute that refers to the address length that a program is prepared to handle on entry. Addresses may be either 24 bits or 31 bits in length. In 24-bit addressing mode, the processor treats all virtual addresses as 24-bit values; in 31-bit addressing mode, the processor treats all virtual addresses as 31-bit values. Programs with an addressing mode of ANY can receive control in either 24-bit or 31-bit addressing mode.

* **Advanced Communications Function (ACF).** A group of IBM licensed programs, principally VTAM programs, TCAM, NCP and SSP that use the concepts of Systems Network Architecture (SNA), including distribution of function and resource sharing.

Advanced Function Printing (AFP). A group of IBM licensed programs that support APA printers.

* **AFP.** Advanced Function Printing.

ALET (access list entry token). A token that points to an entry in an access list. When a program is in AR mode and the ALET is in an access register (with the corresponding general-purpose register being used as base register), the ALET identifies the address space or data space that the system is to reference (while the GPR indicates the offset within the space).

alternate block. On an FBA disk, a block designated to contain data in place of a defective block.

* **alternate index.** In systems with VSE/VSAM, the index entries of a given base cluster organized by an alternate key, that is, a key other than the prime key of the base cluster data records; it gives an alternate directory for finding records in the data component of a base cluster. See also *path*.

* **alternate tape.** A tape drive to which the operating system switches automatically for tape read or write operations if the end of the volume has been reached on the originally used tape drive.

* **alternate track.** On a direct access device, a track designated to contain data in place of a defective primary track.

* **American National Standard Code for Information Interchange (ASCII).** The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity check), used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A)

AMODE. Addressing mode.

* **APA.** All points addressable.

* **APAR.** Authorized Program Analysis Report.

* **appendage routine.** Code physically located in a program or subsystem, but logically an extension of a VSE supervisor routine.

* **application profile.** A control block in which the system stores the characteristics of one or more application programs.

* **application program.** A program written for or by a user that applies directly to the user's work, such as a program that does inventory control or payroll. See also *batch program* and *online application program*.

* **ASCII.** American National Standard Code for Information Interchange.

ASI (automated system initialization) procedure. A set of control statements which specifies values for an automatic system initialization.

* **assemble.** To translate an assembly language program into an object program. (T)

* **assembler.** A computer program that converts assembly language instructions into object code.

assembler language. A programming language whose instructions are usually in one-to-one correspondence with machine instructions and allows to write macros.

attention routine (AR). A routine of the system that receives control when the operator presses the Attention key. The routine sets up the console for the input of a command, reads the command, and initiates the system service requested by the command.

* **authorized program analysis report (APAR).** A request for a correction of a problem caused by a defect in a current unaltered release of a program.

* **automated system initialization (ASI).** A function that allows control information for system startup to be cataloged for automatic retrieval during system startup.

auxiliary storage. Addressable storage that is not part of the processor, for example storage on a disk unit. Synonymous with *external storage*.

* **background partition.** An area of virtual storage in which programs are executed under control of the system. By default, the partition has a processing priority lower than any of the existing foreground partitions.

* **backup copy.** A copy, usually of a file or a library member, that is kept in case the original file or library member is unintentionally changed or destroyed.

* **base cluster.** In systems with VSAM, a key-sequenced or entry-sequenced file over which one or more alternate indexes are built.

basic telecommunications access method (BTAM). An access method that permits read and write communication with remote devices. Its current version is called BTAM-ES.

batch processing. 1. Serial processing of computer programs. 2. Pertaining to the technique of processing a set of computer programs in such a way that each is completed before the next program of the set is started. (A)

batch program. A program that is processed in series with other programs and therefore normally processes data without user interaction.

binary synchronous communication (BSC). Method of telecommunication using binary synchronous line discipline. Contrast with *SDLC*.

bits per second (bps). In serial transmission, the instantaneous bit speed with which a device or channel transmits a character.

block. Usually, a block consists of several records of a file that are transmitted as a unit. But if records are very large, a block can also be part of a record only. See also *control block*.

blocking. The process of combining (or cutting) records into blocks.

* **bps.** Bits per second.

* **bringup.** The process of starting a computer system or a subsystem that is to operate under control of the system.

BSC. Binary synchronous communication.

* **BTAM-ES (Basic Telecommunication Access Method Extended Storage).** An IBM supplied telecommunication access method. It permits read and write communication with remote devices.

B-transient. A phase with a name beginning with \$\$\$ and running in the Logical Transient Area (LTA). Such a phase is activated by special supervisor calls.

* **cache.** A high-speed buffer storage that contains frequently accessed instructions and data; it is used to reduce access time.

cache storage. A random access electronic storage in selected storage controls used to retain frequently used data for faster access by the channel. For example, the IBM 3990 Model 3 contains cache.

catalog. 1. A directory of files and libraries, with reference to their locations. A catalog may contain other information such as the types of devices in which the files are stored, passwords, blocking factors. (I) (A) 2. To store a library member such as a phase, module, or book in a sublibrary.

See also *VSAM master catalog*, *VSAM user catalog*.

* **cataloged procedure.** A set of control statements placed in a library and retrievable by name.

CCB. Command control block.

CCW. Channel command word.

* **central location.** The place at which a computer system's control device, normally the system console in the computer room, is installed.

central processing unit (CPU). The hardware component that interprets and executes instructions. Synonym for *processor*.

chained sublibraries. A facility that allows sublibraries to be chained by specifying the sequence in which they must be searched for a certain library member.

chaining. A logical connection of sublibraries to be searched by the system for members of the same type (phases or object modules, for example).

* **channel adapter.** A communication controller hardware unit used to attach the controller to a System/370 data channel.

channel-attached. Pertaining to the attachment of devices directly by data channels (I/O channels) to a computer. Contrast with link-attached. Synonymous with *locally attached*.

channel command word (CCW). One or more CCWs make up the channel program that directs data channel operations.

channel program. One or more channel command words that control a sequence of data channel operations. Execution of this sequence is initiated by a single start I/O (SIO) instruction.

* **channel scheduler.** The part of the supervisor that controls all input/output operations.

channel-to-channel attachment (CTCA). A function that allows data to be exchanged (1) under the control of VSE/POWER between two virtual VSE machines running under VM or (2) under the control of VTAM between two processors.

checkpoint. 1. A point at which information about the status of a job and the system can be recorded so that the job step can be restarted later. 2. To record such information.

CICS. Customer Information Control System.

CKD device. Count-key-data device.

class. In VSE/POWER, a group of jobs that either come from the same input device or go to the same output device.

CMS. Conversational monitor system.

COBOL. Common business-oriented language.

command control block (CCB). The name of a system control block to hold information about a specific instance of a command.

common business-oriented language (COBOL). A high-level programming language based on English used primarily for business application programs.

common library. A library that can be interactively accessed by any user of the (sub)system that owns the library.

* **communication adapter.** A circuit card with associated software that enables a processor, controller, or other device to be connected to a network.

* **communication controller.** 1. A device that directs transmission of data over the links of a network; its operation is controlled by a program executed in a processor to which the controller is connected or it may be controlled by a program executed within the device. (T)2. A type of communication control unit whose operations are controlled by one or more programs stored and executed in the unit. It manages the details of line control and the routing of data through a network.

communication line. See *telecommunication line*.

* **communication region.** An area of the supervisor that is set aside for transfer of information within and between programs.

* **compaction.** In SNA, the transformation of data by packing two characters in a byte so that only a subset of the allowable 256 characters is used. The most frequently sent characters are compacted.

* **compile.** To translate a source program into an executable program (object program). See also *assembler*.

compiler. A program used to compile.

component. 1. Hardware or software that is part of a computer system. 2. A functional part of an operating system, for example: job control program, VSE/POWER. 3. In VSE/VSAM, a named, cataloged group of stored records, such as the data component or index component of a key-sequenced file or alternate index.

conditional job control. The capability of the job control program to process or to skip one or more statements based on a condition that is tested by the program.

configuration. The devices and programs that make up a system, subsystem, or network.

connect. To authorize library access on the lowest level. A modifier such as "read" or "write" is required for the specified use of a sublibrary.

control block. An area within a program or a routine defined for the purpose of storing and maintaining control information.

control program. A program to schedule and supervise the running of programs in a system.

control unit. See *communication controller*. Synonymous with *controller*.

* **conversational monitor system (CMS).** A virtual machine operating system that provides general

interactive time sharing, problem solving, and program development capabilities and operates under the control of VM/SP.

* **corrective service.** The installation of a PTF or an APAR fix that corrects a specific problem.

count-key-data (CKD) device. A disk device that stores data in the record format: count field, key field, data field. The count field contains, among others, the address of the record in the format: cylinder, head (track), record number and the length of the data field. The key field, if present, contains the record's key or search argument. CKD disk space is allocated by tracks and cylinders. Contrast with *FBA disk device*. See also *extended count-key-data device*.

CPU. Central processing unit.

cross-partition communication control. A facility that enables VSE subsystems and user programs to communicate.

CTCA. Channel-to-channel attachment.

Customer Information Control System (CICS). An IBM licensed program that controls online communication between terminal users and a database. Transactions entered at remote terminals are processed concurrently by user-written application programs. The product includes facilities for building, using, and servicing databases.

Its current version is called CICS/VSE.

DASD. Direct access storage device.

DASD sharing. An option that lets independent computer systems use common data on shared disk devices.

database. A set of data available online that is organized by a common system and used for a common purpose.

* **data entry panel.** A panel in which the user communicates with the system by filling in one or more fields. See also *panel* and *selection panel*.

data file. See *file*.

data import. The process of reformatting data that was used under one operating system (for example, IBM System/3) such that it can subsequently be used under a different operating system (for example, the VSE system).

* **Data Interfile Transfer, Testing and Operations (DITTO) utility.** An IBM licensed program that provides file-to-file services for card I/O, tape, and disk devices.

Data Language/I (DL/I). A database access language used with CICS/VSE.

data link. In SNA, the combination of the link connection and the link stations joining network nodes, for example, a System/370 channel and its associated protocols. A link is both logical and physical.

In SNA, synonym for *link*.

* **data management.** A major function of the operating system. It involves organizing, storing, locating, and retrieving data.

data security. See *access control*.

data set. See *file*.

data space. A range of up to two gigabytes of contiguous virtual storage addresses that a program can directly manipulate through ESA/370 instructions. Unlike an address space, a data space can hold only user data; it does not contain shared areas, system data or programs. Instructions do not execute in a data space, although a program can reside in a data space as non-executable code. Contrast with address space.

deblocking. The process of making each record of a block available for processing. Contrast with *blocking*.

* **dedicated (disk) device.** A device that cannot be shared among users.

default value. A value assumed by the program when no value has been specified by the user.

* **device address.** 1. The identification of an input/output device by its channel and unit number. 2. In data communication, the identification of any device to which data can be sent or from which data can be received.

* **device class.** The generic name for a group of device types, for example, all display stations belong to the same device class. Contrast with *device type*.

* **Device Support Facilities.** An IBM-supplied SCP for performing operations on disk volumes so that they can be accessed by IBM and user programs. Examples of these operations are initializing a disk volume and assigning an alternate track.

* **device type code.** The four- or five-digit code to be used for defining an I/O device to a computer system.

* **dialog.** 1. In an interactive system, a series of related inquiries and responses similar to a conversation between two people. 2. For VSE/SP, a set of panels that can be used to complete a specific task, for example, defining a file.

dialog manager. The program component of VSE/SP that provides for ease of communication between user and system.

direct access. Accessing data on a storage device using their address and not their sequence. This is the typical

access on disk devices as opposed to magnetic tapes. Contrast with *sequential access*.

Direct access storage device. A device in which access time is effectively independent of the location of the data.

directory. 1. A table of identifiers and references to the corresponding items of data. (I) (A) 2. In VSE, specifically, the index for the program libraries. See also *library directory* and *sublibrary directory*.

disk operating system residence volume (DOSRES). The disk volume on which the system sublibrary IJSYSRS.SYSLIB is located including the programs and procedures required for system startup.

disk sharing. An option that lets independent computer systems use common data on shared disk devices.

display station. A display screen with attached keyboard for communication with the system or a network. See also *terminal*.

* **distribution tape.** A magnetic tape that contains, for example, a preconfigured operating system like z/VSE. This tape is shipped to the customer for program installation.

* **DITTO utility.** Data Interfile Transfer, Testing and Operations utility.

DL/I. Data Language/I.

DOSRES. Disk operating system residence volume.

DSF. Device Support Facilities.

dummy device. A device address with no real I/O device behind it. Input and output for that device address are spooled on disk.

dump. 1. Data that has been dumped. (I) (A) 2. To write at a particular moment some contents of storage to another data medium for the purpose of safeguarding or debugging the data. (T)

* **duplex.** Pertaining to communication in which data can be sent and received at the same time.

dynamic partition. A partition created and activated on an 'as needed' basis that does not use fixed static allocations. After processing, the occupied space is released. Dynamic partitions are grouped by class, and jobs are scheduled by class. Contrast with *static partition*.

* **dynamic partition balancing.** A VSE facility that allows the user to specify that two or more or all partitions of the system should receive about the same amount of time on the processor.

dynamic space reclamation. A librarian function that makes space freed by the deletion of a library member reusable.

EBCDIC. Extended binary-coded decimal interchange code.

ECKD device. Extended count-key-data device.

emulation. The use of programming techniques and special machine features that permit a computer system to execute programs written for another system or for the use of I/O devices different from those that are available.

end user. 1. A person who makes use of an application program. 2. In SNA, the ultimate source or destination of user data flowing through an SNA network. May be an application program or a terminal operator.

Enterprise Systems Architecture/370. See *ESA/370* and *ESA/390*.

environmental record editing and printing (EREP) program. The program that makes the data contained in the system recorder file available for further analysis.

EREP program. Environmental record editing and printing program.

error recovery procedures (ERP). Procedures to help isolate and, where possible, to recover from errors in equipment.

ESA/370. IBM Enterprise Systems Architecture/370. The extension to the IBM System/370 architecture which includes the advanced addressability feature that provides access registers.

ESA/390. IBM Enterprise Systems Architecture/390. The latest extension to the IBM System/370 architecture which includes the advanced addressability feature and advanced channel architecture.

* **escape.** To return to the original level of a user interface.

exit. A routine, normally user-supplied, that receives control from the system when a certain event occurs (abnormal-end exit, for example).

extended count-key-data (ECKD) device. A disk storage device that has a data transfer rate faster than some processors can utilize. A specialized channel program is needed to convert ordinary CKD channel programs for use with an ECKD device.

extent. Continuous space on a disk or diskette occupied by or reserved for a particular file or VSAM data space.

extended binary-coded decimal interchange code (EBCDIC). A coded character set consisting of 8-bit coded characters.

external storage. Storage that is not part of the processor.

Fast copy data set program (VSE/FCOPY). See *VSE/FCOPY*.

FBA disk device. Fixed-block architecture disk device.

* **FCB.** Forms control buffer.

FCOPY. See *VSE/FCOPY*.

fetch. 1. To locate and load a quantity of data from storage. (A) 2. To bring a program phase into virtual storage from a sublibrary and pass control to this phase. 3. The name of the macro instruction (FETCH) used to accomplish 2. See also *loader*.

file. A named set of records stored or processed as a unit. (T) Synonymous with *data set*.

fixed-block architecture (FBA) disk device. A disk device that stores data in blocks of fixed size. These blocks are addressed by block number relative to the beginning of the file. Contrast with *CKD device*.

* **foreground partition.** A space of virtual storage in which programs are executed under control of the system. By default, a foreground partition has a higher processing priority than the background partition.

* **forms control buffer (FCB).** In the 3800 Printing Subsystem, a buffer for controlling the vertical format of printed output.

* **fragmentation (of storage).** Inability to allocate unused sections (fragments) of storage in the real or virtual address range of virtual storage.

GB. Gigabyte.

generation. See *macro generation*.

generation feature. An IBM licensed program order option used to tailor the object code of a program to user requirements.

* **GETVIS space.** Storage space within a partition or the shared virtual area, available for dynamic allocation to programs.

gigabyte (GB). 1024MB of storage (see MB). One gigabyte equals 1 073 741 824 bytes, which is 2 to the thirtieth power.

guest system. A data processing system that runs under control of another (host) system.

* **half-duplex.** In data communication, pertaining to transmission of data in only one direction at a time. Contrast with *duplex*.

hardcopy file. A system file on disk, used to log all lines of communication between the system and the operator at the system console, to be printed on request.

hard wait. The condition of a processor when all operations are suspended. System recovery from a hard wait is impossible without performing a new system startup.

hardware. Physical equipment used in data processing, as opposed to programs, procedures, rules, and associated documentation. (I) (A) Contrast with *software*.

help panel. A display of information provided by the system in response to a user's help request.

* **host system.** The controlling or highest level system in a data communication configuration.

ICA. Integrated communication adapter.

ICCF. See *VSE/ICCF*.

include function. Retrieves a library member for inclusion in program input.

index. In data management, a table used to locate the records of a file.

* **initial program load (IPL).** The process of loading system programs and preparing the system to run jobs.

input/output control system (IOCS). A group of IBM supplied routines that handle the transfer of data between main storage and auxiliary storage devices.

intelligent workstation. Replaced by *programmable workstation*.

integrated communication adapter (ICA). The part of a processor where multiple lines can be connected.

interactive. A characteristic of a program or system that alternately accepts input and then responds. An interactive system is conversational, that is, a continuous dialog exists between user and system. Contrast with *batch*.

Interactive Computing and Control Facility (ICCF). An IBM licensed program that serves as interface, on a time-slice basis, to authorized users of terminals linked to the system's processor.

interactive interface. A system facility which controls how different users see and work with the system by means of user profiles. When signing on, the interactive interface makes available those parts of the system authorized by the profile. The interactive interface has

sets of selection- and data-entry panels through which users communicate with the system.

interactive partition. An area of virtual storage for the purpose of processing a job that was submitted interactively via VSE/ICCF.

interface. A shared boundary between two hardware or software units defined by common functional or physical characteristics. It might be a hardware component or a portion of storage or registers accessed by several computer programs.

* **intermediate storage.** Any storage device used to hold data temporarily before it is processed. See also *buffer storage*.

I/O (input/output). See *input* and *output*.

IOCS. Input/output control system.

IPL. Initial program load.

* **irrecoverable error.** An error for which recovery is impossible without the use of recovery techniques external to the computer program or run. (T)

JCL. Job control language.

JECL. Job entry control language.

JHR. Job header record.

job accounting. A system function that lists how much every job step uses of the different system resources.

job accounting interface. A function that accumulates accounting information for each job step, to be used for charging the users of the system, for planning new applications, and for supervising system operation more efficiently.

* **job accounting table.** An area in the supervisor where accounting information is accumulated for the user. in the respective DLBL statement.

job control language (JCL). A language that serves to prepare a job or each job step of a job to be run. Some of its functions are: to identify the job, to determine the I/O devices to be used, set switches for program use, log (or print) its own statements, and fetch the first phase of each job step.

job control statement. A particular statement of JCL.

job entry control language (JECL). A control language that allows the programmer to specify how VSE/POWER should handle a job.

job step. One of a group of related programs complete with the JCL statements necessary for a particular run.

Every job step is identified in the job stream by an EXEC statement under one JOB statement for the whole job.

job stream. The sequence of jobs as submitted to an operating system.

KB. Kilobyte (KB equals 1024 bytes).

key. In VSE/VSAM, one or several characters taken from a certain field (key field) in data records for identification and sequence of index entries or of the records themselves.

key sequence. The collating sequence either of records themselves or of their keys in the index or both. The key sequence is alphanumeric.

* **kilobyte (KB).** 1024 bytes of storage. One kilobyte equals 1024 bytes, which is 2 to the twelfth power.

label. 1. An identification record for a tape, disk, or diskette volume or for a file on such a volume. 2. In assembler programming, a named instruction generally used for branching.

label information area. An area on a disk to store label information read from job control statements or commands. Synonymous with *label area*.

language translator. A general term for any assembler, compiler, or other routine that accepts statements in one language and produces equivalent statements in another language.

* **librarian.** The set of programs that maintains, services, and organizes the system and private libraries.

library. See *VSE library* and *VSE/ICCF library*.

* **library block.** A block of data stored in a sublibrary.

* **library directory.** The index that enables the system to locate a certain sublibrary of the accessed library.

* **library member.** The smallest unit of data to be stored in and retrieved from a sublibrary.

* **licensed program.** A separately priced program and its associated materials that bear an IBM copyright and are offered to customers under the terms and conditions of either the Agreement for IBM Licensed Programs (ALP) or the IBM Program License Agreement (PLA).

line. Short for telecommunication line. Any physical medium such as a wire or microwave beam, that is used to transmit data. Synonymous with *transmission line*.

line printer. A device that prints a line of characters as a unit. (I) (A) Contrast with *character printer* or *page printer*.

link. To connect items of data or portions of programs, for example linking of object programs by the linkage editor or linking of data items by pointers.

linkage editor. A program to build a phase (executable code) from one or several independently translated object modules or existing phases or both. In creating the phase, the program resolves cross references among the modules and phases available as input. The program can catalog the newly built phases.

* **link-attached.** Pertaining to devices connected to a control unit by a data link. Synonymous with *remote*. Contrast with *channel-attached*.

link-edit. To create a loadable computer program by having the linkage editor process compiled (assembled) source programs.

loader. A routine, commonly a computer program, that reads data or a program into processor storage. See also *relocating loader*.

local shared resources (LSR). A VSE/VSAM option activated by three extra macros to share control blocks among files.

* **lock file.** In a shared disk environment under VSE, a system file on disk used by the sharing systems to control their access to shared data.

* **logging.** The recording of data about specific events.

logical record. A user record, normally pertaining to a single subject and processed by data management as a unit. Contrast with *physical record* which may be larger or smaller.

logical unit (LU). A name used in programming to represent an I/O device address. *physical unit (PU)*, *system services control point (SSCP)*, *primary logical unit (PLU)*, and *secondary logical unit (SLU)*.

logical unit name. In programming, a name used to represent the address of an input/ output unit.

logo. A trademark or other art work that is associated with a firm or product. A logo often appears as the first screen of an interactive program.

LSR. Local shared resources.

LU. Logical unit.

macro (instruction). 1. In assembler programming, a user-invented assembler statement that causes the assembler to process a set of statements defined previously in the macro definition. 2. A sequence of VSE/ICCF commands defined to cause a sequence of certain actions to be performed in response to one request.

macro definition. A set of statements and instructions that defines the name of, format of, and conditions for

generating a sequence of assembler statements and machine instructions from a single source statement.

macro expansion. See *macro generation*.

macro generation. An assembler operation by which a macro instruction gets replaced in the program by the statements of its definition. It takes place before assembly. Synonymous with *macro expansion*.

* **main task.** The main program within a partition in a multiprogramming environment.

* **Maintain system history program (MSHP).** A program used for automating and controlling various installation, tailoring, and service activities for a VSE system.

* **MB.** Megabyte (MB equals 1 048 576 bytes).

* **megabyte (MB).** 1024KB of storage (see KB). One megabyte equals 1 048 576 bytes, which is 2 to the twentieth power.

* **member.** The smallest unit of data that can be stored in and retrieved from a sublibrary.

message. 1. In VSE, a communication sent from a program to the operator or user. It can appear on a console, a display terminal or on a printout. 2. In telecommunication, a logical set of data being transmitted from one node to another.

* **microcode.** 1. A code written using the instructions of a specific instruction set and implemented in a part of storage that is not program-addressable. 2. To design write, and test one or more micro instructions.

* **migrate.** To move to a changed operating environment, usually to a new release or version of a system.

* **module.** A program unit that is discrete and identifiable with respect to compiling, combining with other units, and loading; for example, the input to, or output from, an assembler, a compiler, linkage editor, or executive routine. (A)

* **MSHP.** Maintain system history program.

* **multiprogramming.** 1. A mode of operation that provides for interleaved execution of several programs by a single processor. (I) (A) 2. Pertaining to concurrent execution of several programs by a computer. (A)

multitasking. Concurrent running of one main task and one or several subtasks in the same partition.

* **nest.** To incorporate a structure or structures of some kind into a structure of the same kind. For example, to nest one loop (the nested loop) within another loop or to nest one subroutine (the nested subroutine) within another subroutine. (T)

NetView. An IBM licensed program to monitor a network, manage it, and diagnose its problems.

network. 1. An arrangement of nodes (data stations) and connecting branches. 2. The assembly of equipment through which connections are made between data stations.

networking. Making use of the services of a network program.

* **object code.** Output from a compiler or assembler which is itself executable machine code or is suitable for processing to produce executable machine code. (A)

object module (program). A program unit that is the output of an assembler or compiler and is input to a linkage editor.

OCCF. See VSE/OCCF.

online processing. Processing by which the input data enters the computer directly from a display station and the output data is transmitted directly to the display station.

* **operating system.** Software that controls the running of programs; an operating system may provide services such as resource allocation, scheduling, input/output control, and data management. (I) (A)

* **operator command.** A statement to a control program, issued via a console or terminal. It causes the control program to provide requested information, alter normal operations, initiate new operations, or end existing operations.

Operator Communication Control Facility (OCCF). An IBM licensed program that helps reduce operator interaction in the operation of a VSE controlled installation and helps centralize data processing skills.

optical reader/sorter. A device that reads hand written or machine printed symbols on a voucher and, after having read the voucher, can sort it into one of the available stacker-select pockets.

optional licensed program. An IBM licensed program that a user can install on VSE by way of available installation-assist support.

page. 1. In a virtual storage system, the unit of code or data or both which is transferred between processor storage and the PDS as needed for processing. 2. To transfer pages between processor storage and the page data set.

page data set (PDS). One or more extents of disk storage in which pages are stored when they are not needed in processor storage.

page fault. A program interruption that occurs when a program page marked "not in processor storage" is referred to by an active page.

* **page fixing.** Marking a page so that it is held in processor storage until explicitly released. Until then, it cannot be paged out.

page frame. An area of processor storage that can contain a page.

page-in. The process of transferring a page from the PDS to processor storage.

page I/O. Page-in and page-out operations.

page-out. The process of transferring a page from processor storage to the PDS.

* **page pool.** The set of page frames available for paging virtual-mode programs.

panel. The complete set of information shown in a single display on a terminal screen. Scrolling back and forth through panels is like turning manual pages. See also *selection panel* and *data entry panel*.

partition. A division of the virtual address area available for running programs. See also *dynamic partition*, *static partition*.

* **partition balancing, dynamic.** A VSE facility that allows the user to specify that two or more or all partitions of the system should receive about the same amount of time on the processor.

PDS. Page data set.

* **phase.** The smallest unit of executable code that can be loaded into virtual storage.

* **physical record.** The amount of data transferred to or from auxiliary storage. Synonymous with *block*.

* **physical unit (PU).** In SNA, the component that manages and monitors the resources of a node, such as attached links and adjacent link stations, as requested by an SSCP via an SSCP-SSCP session.

PL/I. A programming language designed for use in a wide range of commercial and scientific computer applications.

PNET. Programming support available with VSE/POWER; it provides for the transmission of selected jobs, operator commands, messages, and program output between the nodes of a network.

POWER. See VSE/POWER.

* **pregenerated operating system.** An operating system such as VSE/SP that is shipped by IBM mainly in object code. IBM defines such key characteristics as the size of the main control program, the organization and size of libraries, and required system areas on disk. The customer does not have to generate an operating system.

* **preventive service.** The installation of one or more PTFs on a VSE system to avoid the occurrence of anticipated problems.

* **primary library.** A VSE library owned and directly accessible by a certain terminal user.

Print Services Facility/VSE. An access method that provides support for the advanced function printers.

priority. A rank assigned to a partition or a task that determines its precedence in receiving system resources.

* **private library.** A user-owned library that is separate and distinct from the system library.

* **private partition.** Any of the system's partitions that are not defined as shared. See also *shared partition*.

procedure. See *cataloged procedure*.

* **processing.** The performance of logical operations and calculations on data, including the temporary retention of data in processor storage while this data is being operated upon.

processor. The hardware component that interprets and executes instructions. (I) (A)

processor storage. The storage contained in one or more processors and available for running machine instructions. Synonymous with *real storage*.

* **production library.** 1. In a pre-generated operating system (or product), the program library that contains the object code for this system (or product). 2. A library that contains data needed for normal processing. Contrast with *test library*.

profile. A description of the characteristics of a user or a computer resource.

* **programmer logical unit.** A logical unit available primarily for user-written programs. See also *logical unit name*.

program product. See *licensed program*.

program service. The customer- or program-related IBM service of correcting design or implementation errors via APARs and PTFs.

program temporary fix (PTF). A solution or by-pass of one or more problems documented in APARs. PTFs are distributed to IBM customers for preventive service to a current release of a program.

prompt. To issue messages to a terminal or console user, requesting information necessary to continue processing.

PSF/VSE. Print Services Facility/VSE.

PTF. Program temporary fix.

PU. Physical unit.

punch. 1. To make holes in some data medium according to a signal code and thus save data on that medium. 2. A machine (output device) to punch 80-column punch cards.

* **punch card.** A card into which hole patterns can be punched; normally, it is characterized by 80 columns and 12 rows of punch positions.

* **queue.** 1. A line or list formed by items in a system waiting for service; for example, tasks to be performed or messages to be transmitted in a network. 2. To arrange in, or form, a queue.

queue file. A disk file maintained by VSE/POWER that holds control information for the spooling of job input and job output.

queue record. A record in the queue file containing descriptive information about a job or job output.

* **random processing.** The treatment of data without respect to its location on disk storage, and in an arbitrary sequence governed by the input against which it is to be processed.

real address. The address of a location in processor storage.

* **real address area.** In VSE, the area of virtual storage where virtual addresses are equal to real addresses.

* **real address space.** The address space whose addresses map one to one to the addresses in processor storage.

real mode. In VSE, a processing mode in which a program may not be paged. Contrast with *virtual mode*.

real storage. See *processor storage*.

* **record.** A collection of related data or words, treated as a unit. See *logical record*, *physical record*.

record formatted maintenance statistics (RECFMS). In NetView, a statistical record built by an SNA controller and usually solicited by the host.

recovery management support (RMS). System routines that gather information about hardware failures and that initiate a retry of an operation that failed because of processor, I/O device, or channel errors.

* **reentrant.** The attribute of a program or routine that allows the same copy of the program or routine to be used concurrently by several tasks.

refresh release. An upgraded VSE system with the latest level of maintenance for a release.

relocatable module. In VSE, a library member of type object. It consists of one or more control sections cataloged as one member.

relocating loader. A function that modifies addresses of a phase, if necessary, and loads the phase for running into the partition selected by the user.

* **remote job entry (RJE).** Submission of jobs through an input unit that has access to a computer through a data link.

residency mode (RMODE). A program attribute that refers to the location where a program is expected to reside in virtual storage. RMODE 24 indicates that the program must reside in the 24-bit addressable area (below 16 megabytes), RMODE ANY indicates that the program can reside anywhere in 31-bit addressable storage (above or below 16 megabytes).

* **restore.** To write back on disk data that was previously written from disk to an intermediate storage medium such as tape.

RJE. Remote job entry.

RJE workstation. Any workstation that is used for remote job submission and for the remote retrieval of output.

RMODE. Residency mode.

RMS. Recovery management support.

* **routine.** Part of a program, or a sequence of instructions called by a program, that may have some general or frequent use. (I) (A)

* **routing.** The assignment of the path by which a message will reach its destination.

RPG II. A commercially oriented programming language suitable for writing application programs that meet common business data processing requirements.

* **run.** 1. A performance of one or more jobs. (I) (A) 2. A performance of one or more programs. (I) (A) 3. To cause a program or job to be performed.

SAM. Sequential access method.

SAM ESDS file. A SAM file managed in VSE/VSAM space, so it can be accessed by both SAM and VSE/VSAM macros.

schedule. To select a program or task for getting control over the processor.

SCP. System control programming.

SDL. System directory list.

SDLC. Synchronous data link control.

* **search chain.** The order in which chained sublibraries are searched for the retrieval of a certain library member of a specified type.

second-level directory. A table in the SVA containing the highest phase names found on the directory tracks of the system sublibrary.

security. See *access control*.

* **selection panel.** A displayed list of items from which a user can make a selection. Synonymous with *menu*.

sense. Determine, on request or automatically, the status or the characteristics of a certain I/O or communication device.

sequential access. The serial retrieval of records in their entry sequence or serial storage of records with or without a premeditated order. Contrast with *direct access*.

sequential access method (SAM). A data access method that writes to and reads from an I/O device record after record (or block after block). On request, the support performs device control operations such as line spacing or page ejects on a printer or skip a certain number of tape marks on a tape drive.

sequential file. A file in which records are processed in the order in which they are entered and stored.

shared disk option. An option that lets independent computer systems use common data on shared disk devices.

* **shared partition.** In VSE, a partition allocated for a program (VSE/POWER, for example) that provides services for and communicates with programs in other partitions of the system's virtual address spaces.

* **shared spooling.** A function that permits the VSE/POWER account file, data file, and queue file to be shared among several computer systems with VSE/POWER.

* **shared virtual area (SVA).** In VSE, a high address area that contains a list system directory list (SDL) of frequently used phases, resident programs shared between partitions, and an area for system support.

skeleton. A set of control statements and/or instructions that requires user-specific information to be inserted before it can be submitted for processing.

SNA. System Networks Architecture.

SNA network. The part of a user-application network that conforms to the formats and protocols of SNA.

* **software.** Programs, procedures, rules, and any associated documentation pertaining to the operation of a computer system.

source member. A library member containing source statements in any of the programming languages supported by VSE.

* **source program.** A computer program expressed in a source language. (I) (A) Contrast with *object module*.

source statement. A statement written in symbols of a programming language.

spanned record. A record that extends over several blocks.

SQL/DS. Structured Query Language/Data System.

SS system. Start-stop system.

stand-alone program. A program that runs independently of (not controlled by) the VSE system.

* **standard label.** A fixed-format record that identifies a volume of data such as a tape reel or a file that is part of a volume of data.

start-stop (SS) system. A data transmission system in which each character is preceded by a start signal and is followed by a stop signal. (T)

startup. The process of performing IPL of the operating system and of getting all subsystems and application programs ready for operation.

static partition. A partition, defined at IPL time and occupying a defined amount of virtual storage that remains constant. Contrast with *dynamic partition*.

storage dump. See *dump*.

storage fragmentation. Inability to allocate unused sections (fragments) of storage in the real or virtual address range of virtual storage.

Structured Query Language/Data System. An IBM licensed program for using a database in an online, interactive, or batch environment.

sublibrary. In VSE, a subdivision of a library. Members can only be accessed in a sublibrary.

sublibrary directory. An index for the system to locate a member in the accessed sublibrary.

submit. A VSE/POWER function that passes a job to the system for processing.

* **subsystem.** A secondary or subordinate system or program, usually capable of operating independently of, or asynchronously with, the operating system.

subtask. A task that is initiated by the main task or by another subtask.

* **supervisor.** The part of a control program that coordinates the use of resources and maintains the flow of processor operations.

SVA. Shared virtual area.

switched line. A telecommunication line in which the connection is established by dialing.

SYSRES. System residence volume.

* **system console.** A console, usually equipped with a keyboard and display screen for control and communication with the system.

system control programming (SCP). IBM-supplied, nonlicensed program fundamental to the operation of a system or to its service or both.

system directory list (SDL). A list containing directory entries of frequently-used phases and of all phases resident in the SVA. The list resides in the SVA.

* **system file.** In VSE, a file used by the operating system, for example, the hardcopy file, the recorder file, the page data set.

system logical unit. A logical unit available primarily for operating system use. See also *logical unit name*.

Systems Network Architecture (SNA). The description of the logical structure, formats, protocols, and operational sequences for transmitting information units through and controlling the configuration and operation of networks.

system recorder file. The file that is used to record hardware reliability data. Synonymous with *recorder file*.

system refresh release. See *refresh release*.

system residence volume (SYSRES). The disk volume on which the system sublibrary is stored and from which the hardware retrieves the initial program load routine for system startup.

system sublibrary. The sublibrary that contains the operating system. It is stored on the system residence volume (SYSRES).

* **tailor.** A process that defines or modifies the characteristics of the system.

* **task.** The basic unit of synchronous program execution. A task competes with other tasks for system resources such as processing time and I/O channels.

task management. The functions of a control program that control the use, by tasks, of the processor and other resources (except for input/output devices).

* **telecommunication.** Transmission of data between computer systems and between such a system and remote devices.

telecommunication line. Any physical medium such as a wire or microwave beam, that is used to transmit data. Contrast with *data link*.

terminal. A point in a system or network at which data can either enter or leave. (A) Usually a display screen with a keyboard.

* **throughput.** 1. A measure of the amount of work performed by a computer system over a given period of time, for example, jobs per day. (I) (A) 2. In data communication, the total traffic between stations per unit of time.

trace. 1. To record a series of events as they occur. 2. A record of specified events during the run of a program. 3. A program to produce such a record.

* **track.** A circular path on the surface of a disk or diskette. Smallest unit of physical disk space.

track hold. A function that protects a track while it is being updated by one program from being accessed by another program.

* **transient area.** An area within the control program used to provide high-priority system services on demand.

transmission line. Synonym for *telecommunication line*.

* **transmit.** To send data from one place for reception elsewhere. (A)

UCB. Universal character set buffer.

* **UCS.** Universal character set.

universal character set buffer (UCB). A buffer to hold UCS information.

* **utility program.** 1. A program in general support of computer processes, for example, a diagnostic program, a trace program, or a sort program. (T) Synonymous with *service program*. 2. A program that performs an everyday task such as copying data from one storage device to another. (A)

VAE. Virtual addressability extension.

virtual address. An address that refers to a location in virtual storage. It is translated by the system to a processor storage address when the information stored at the virtual address is to be used.

virtual addressability extension (VAE). A storage management support that gives the user of VSE multiple address spaces of virtual storage.

virtual address area. The virtual range of available program addresses.

virtual address space. In VSE, a subdivision of the virtual address area available to the user for the allocation of private (non-shared) partitions.

* **virtual I/O area (VIO).** An extension of the page data set; used by the system as intermediate storage, primarily for control data.

* **virtual machine.** A functional simulation of a computer system and its associated devices.

* **virtual mode.** The operating mode of a program which may be paged.

* **virtual partition.** In VSE, a division of the dynamic area of virtual storage.

virtual storage. Addressable space image for the user from which instructions and data are mapped into processor storage locations.

volume. A data carrier that is mounted and demounted as a unit, for example, a reel of tape or a disk pack. (I) Some disk units have no demountable packs. In that case, a volume is the portion available to one read/write mechanism.

volume ID. The volume serial number, which is a number in a volume label assigned when a volume is prepared for use by the system.

volume table of contents (VTOC). A table on a disk volume that describes every file on it.

VSAM. See *VSE/VSAM*.

VSE (Virtual Storage Extended). A system that consists of a basic operating system and any IBM-supplied and user-written programs required to meet the data processing needs of a user. VSE and the hardware it controls form a complete computing system. Its current version is called z/VSE.

VSE/Advanced Functions. The basic operating-system component of z/VSE.

VSE/DITTO (VSE/Data Interfile Transfer, Testing, and Operations Utility). An IBM licensed program that provides file-to-file services for disk, tape, and card devices.

* **VSE/FCOPY (VSE/Fast Copy Data Set program).** An IBM program for fast copy data operations from disk to disk and dump/restore operations via an intermediate dump file on magnetic tape or disk.

* **VSE/ICCF (VSE/Interactive Computing and Control Facility).** An IBM program that serves as interface, on a time-slice basis authorized users of terminals linked to the system's processor.

VSE/ICCF library. A file composed of smaller files (libraries) including system and user data which can be accessed under the control of VSE/ICCF.

VSE library. A collection of programs in various forms and storage dumps stored on disk. The form of a program is indicated by its member type such as source

code, object module, phase, or procedure. A VSE library consists of at least one sublibrary which can contain any type of member.

VSE/OCCF (Operator Communication Control Facility). An IBM licensed program that helps reduce operator interaction in the operation of a VSE-controlled installation and helps centralize data processing skills.

* **VSE/OLTEP (VSE/Online Test Executive Program).** An IBM program for managing the online tests that are available for preventive service for I/O devices. Normally, only IBM service representatives use this program.

* **VSE/POWER.** An IBM program primarily used to spool input and output. The program's networking functions enable a VSE system to exchange files with or run jobs on another remote processor.

VSE/SP Unique Code. A component of z/VSE.

VSE/VSAM (VSE/Virtual Storage Access Method). An IBM access method for direct or sequential processing of fixed and variable length records on disk devices.

* **VSE/VSAM managed space.** A user-defined space on disk placed under the control of VSE/VSAM.

VTAM (Virtual Telecommunications Access Method). An IBM licensed program that controls communication and the flow of data in an SNA network. It provides single-domain, multiple-domain, and interconnected network capability; it supports application programs and subsystems (VSE/POWER, for example).

VTOC. Volume table of contents.

wait state. The condition of a processor when all operations are suspended. System recovery from a hard wait is impossible without performing a new system startup. Synonym for *hard wait*.

Workstation File Transfer Support. Enables the exchange of data between IBM Personal Computers linked to a VSE host system where the data is kept in intermediate storage. PC users can retrieve that data and work with it independently of VSE.

z/VSE (z/Virtual Storage Extended). The most advanced VSE system currently available.

Index

Numerics

3211
use of with indexing xi

A

AB/S 186
ABEND
 definition of 3
ABEND dump
 activation via
 OPTION statement 5
 STDOPT command 5
 contents
 data space dump 6
 partition dump 6
 symptom part 5
 system dump 5
 contents controlled via
 OPTION statement 20
 STDOPT command 20
 definition of 4
 options for
 partition or system 20
 output contents 5
 output destination 20
 output device 20
 output written into
 a dump sublibrary 20
ABEND trace (interactive trace program) 35
absolute addresses
 See addresses
access register 67
access registers, displaying 78
ACF/VTAM traces (see VTAM traces) 68, 107
activating dump writing 14
adding dumps 178
additional output for tracing events 160
ADDRESS definition 111
ADDRESS parameter (interactive trace program) 37
addresses
 of phases in storage
 on linkage editor map 225
 to define trace range 158
ADRS 186
aids
 See serviceability aids
ALET 67
alter
 low address storage
 with DSPLY/ALTER commands 226
 virtual storage 226
ALTER command
 interactive trace program 38
alter/display feature 236
 areas accessible by 236

altering
 with the ALTER command 226
analysis routines
 IJBXCSMG 191
 IJBXDEBUG 192
 IJBXSDA 198
 invocation 191
 selecting 190
area definition
 by partition-ID 157
 for tracing events
 by absolute addresses 158
 by partition-ID 157
 by phase name 158
 by relative addresses 157
 for address spaces 158
 in the LTA 158
 space ID 158
 in a DUMP command 21
AREA definition
 ALL 70, 111
 for tracing events
 by partition-ID 157
 general description 70
 in direct input mode 111
 partition_id 70
 partition-id 111
 SUP 70, 111
array extensions 219
assignments, I/O devices
 See device assignments
audience of this manual xi

B

batch
 See also batch mode
 batch mode 169
 reader mode 169
batch mode 169
branch execution trace 56, 89
branch trace (interactive trace program) 35
branch trace (SDAID) 145
 general description 56
 in direct input mode 89
 in procedure mode 121
buffer 57, 90
 contents of
 on event 74, 114
 tracing usage of by VTAM 68, 107
buffer sizes for SDAID 143
buffer trace 145
 general description 57
 in direct input mode 90
buffer-overflow trace
 See wraparound buffer
BUFFOUT keyword operand 129
BUFFOUT=CANCEL 129
BUFFOUT=EXT 129
BUFFOUT=PGMC 129

BYPASS
 statement 199

C

CALL statement 190
calling an analysis routine 191
CANCEL command 21
 dump stored in
 dump sublibraries 12
 output
 as partition dump 21
 as system dump 21
cancel trace 57, 90, 145
 general description 57
 in direct input mode 90
CCB
 See command control block
CCW
 See channel command word
CCWD (CCW with data), display of on event 74, 115
chain extensions 219
chaining
 of control blocks 220
channel command word
 display of on event 74, 114, 160
channel command word (CCW) 74, 114, 160
CHannel definition 113
CICS/VSE dump 175
command control block (CCB)
 contents of on event 74, 160
command symbols xv
commands 163
 ALTER 38, 226
 CANCEL 21
 DISPLAY 38
 DSPLY 226
 DUMP 21
 ENDSD 163
 for tracing events
 method of syntax presentation 134
 prompting for 143
 sequence of 135
 summary 131
 syntax summary 135
 GO 39
 LISTIO 229
 OUTDEV 53, 142
 PAUSE 233
 QUERY 37
 SDAID
 See tracing events
 STOP 233
 STOPSD 163
 TRACE 143
 tracing events
 syntax diagrams 134

COMREG
See partition communication region
 console
 console 169
 console communication per job 229
 continuation
 EXEC PROC statement 117
 control block
 display contents of
 on event 74, 114
 formatting 221
 linkages 220
 locating 219
 control blocks
 dumped by ABEND dump 6
 control information
 OUTDEV command
 prompting for 142
 control information input
 DUMP command 21
 invoking DOSVSDMP 27
 LVTOC program 233
 SDAID program 143
 sequence of commands 131
 tracing events
 additional output 160
 TRACE command 143
 control record
 See dump management
 control registers (CREG)
 display of on event 75, 160
 control statements
 ACTION 223
 invoking DOSVSDMP 27
 invoking LVTOC program 233
 LISTIO 229
 PAUSE 233
 stand-alone dump program 25
 conventions
 notational, SDAID commands 134
 conventions, command xv
 creating a stand-alone dump
 program 25
 CREG
 See control registers
 CU definition 82, 113
 current dump
 selecting 178

D

DATA
 for PRINT statement 181, 184
 data recovery
 See recovery
 data space
 dump file 8
 dumping priority (OPTION
 SADUMP) 8
 data space dump 6
 output contents 6
 DDT
 See device definition table
 deactivating dump writing 15
 debugging
 hardware aids for 235

default values
 set by procedures 117
 set by SDAID 117
 defining
 dump area
 See dump area definition
 dump sublibraries
 example 13
 I/O device range for tracing
 events 158
 occurrence of events 160
 output device for tracing events 53,
 142
 type of a trace 144
 DELETE
 Info/Analysis control statement 180
 DELETE statement 180
 deleting a dump 180
 from a dump sublibrary 180
 device assignments
 DOSVSDMP 27
 for printing from tape 27
 listing of 229
 device definition table 6
 device specification
 in an OUTDEV command 53, 142
 device status information
 saving disk usage statistics 235
 writing of onto SYSREC 235
 devices
 input 169
 output 169
 DIBEXT
 See disk information block extension
 DIBTAB
 See disk information block
 direct input mode
 trace initialization 83
 disk data, dumping of 225
 disk information block 6
 disk information block extension 6
 diskette data, dumping of 225
 display 226, 227
 console communication per job 229
 label information area 230
 low address storage
 on event 79, 160
 of current event
 XPCCB (XPCC communication
 control block) 80
 XPDATABU (XPCC data buffer
) 81
 of current LOCK or UNLOCK event
 LOCKTE (Locktable area) 79
 on event
 CCB (command control block) 74,
 160
 CCW (channel command
 word) 74, 114, 160
 CCW with data 74, 115
 COMREG (partition
 communication region) 75, 160
 CREG (control registers) 75, 160
 GReg (general registers) 78
 GREG (general registers) 160
 IORB (I/O request block) 74, 160
 IOTAB (I/O tables and blocks) 78

display (*continued*)
 on event (*continued*)
 IOTAB (I/O tables and
 blocks) 160
 low address storage 79, 160
 LTA (logical transient area) 79,
 160
 volume table of contents 232
 with the DSDPLY command 226
 DISPLAY command
 interactive trace program 38
 display feature
 See alter/display feature
 documenting
 system problems 15
 DOSVSDMP utility
 control statements for 27
 functions 25
 printing stand-alone dump 29
 SDAID tape printing 31
 DSDPLY/ALTER commands 226
 dump 1, 11
 ABEND dump 3
 adding 178
 analysis routine call 191
 CANCEL command dump 21
 CICS/VSE 175
 contents 172
 overview 3, 183
 contents of
 for a DUMP command 21
 creation 172
 current 178
 data from disk, diskette, or tape 225
 data records 3
 data space 6
 deleting 180
 DUMP command 21
 DUMP command dump
 printed output 210
 DUMP NAME
 statement 178
 event triggered 76, 115
 external routines file 171
 files to process 11
 identification of 15
 invoking from within a program 23
 library 172
 life cycle 172
 listing of
 examples 181
 methods of requesting 19
 name format 15
 onloading 203
 overview 1
 partition 3
 PRINT control statement 181
 printed from tape 204
 printing 183
 formatted areas 190
 selective areas 189
 symptoms 184
 requested by macros 23
 SDAID dump trace 23
 selecting 178
 sending to IBM Support 16
 stand-alone 172, 218

- dump (*continued*)
 - stand-alone dump
 - formatted output 208
 - stand-alone dump program 7
 - stored in
 - dump sublibraries 12
 - summary 19
 - symptom records 3
 - system 3
 - types 185
 - virtual storage dump, on event 115
 - written into sublibraries
 - prerequisites 20
 - written to SYSDUMP 11
- DUMP (virtual storage dump) output
 - definition 76, 115
- dump area definition
 - by OPTION specification 20
 - in a DUMP command 21
 - in a TRACE command 76, 115
- DUMP command 21
 - area definition 21
 - control information 21
 - dump stored in
 - dump sublibraries 12
 - formatted output 210, 211
- DUMP command dump
 - formatted printout 210
- dump contents 6
 - DDT 6
 - DIBEXT 6
 - DIBTAB 6
 - EDT 6
 - LDT 6
 - LOADLS 6
 - LPT 6
 - LUBEXT 6
 - LUBTAB 6
 - PUB2TAB 6
 - PUBOWN 6
 - PUBTAB 6
 - SDT 6
 - SYSFIL 6
 - TCB 6
 - TIB 6
- dump library
 - See also* dump sublibraries
 - defining 13
 - example 13
 - labels needed 13
 - purpose of 11
 - requirements 13
 - security of 13
- DUMP macro 23
- dump management
 - See also* dump management file
 - adding dump 178
 - deleting dump 180
 - file 169, 179
 - control record 180
 - initialization 180
 - selecting dump 178
- dump management file 169
 - initialization 170
 - labels 171
- dump management library 11
- DUMP NAME 170

- DUMP NAME (*continued*)
 - control statement 170
 - statement 178
- dump offload 198
 - volume 199
- dump onload 202, 203
 - volume 202
- dump option 39
- dump output definition
 - SDAID 23
- dump sublibraries 11
 - clear space 18
 - concept 12
 - contents 12
 - defining 13
 - example 13
 - deleting a dump 180
 - dump storing
 - deactivation 15
 - full condition
 - handling of 16
 - identify dumps 15
 - load dump into 203
 - required JCL options 14
 - required LIBDEF statement 13
 - requirements 13
 - used to store dumps
 - requirements 13
 - writing dumps into 13
- dump symptoms 185
 - See also* symptom record
 - description 185
 - dump type 185
 - dump types 185
 - environment 185
 - function 184
 - interactive mode 173
 - optional symptoms 187
 - panel 173
 - printing 184
 - summary 185, 186
- dump tape
 - sending to IBM Support 17
- dump types
 - in dump symptoms 185
- dump utilities 25
- dump viewing 188
- dumps
 - DOSVSDMP utility 25
 - virtual storage dump, on event 76
- dumps requested by
 - CANCEL command 21
 - DUMP command 21
 - macros 23
 - SDAID dump 23
 - stand-alone dump 21

E

- EDT
 - See* extent definition table
- electronic
 - sending of a dump to IBM Support 16
- end SDAID trace initialization
 - in direct input mode 86
- end-of-input 176, 211

- end-of-job 176, 211
- ending
 - SDAID control information input 163
 - SDAID execution
 - by an ENDSD command 163
 - on event 160
- ending a session 211
- ENDSD command
 - description 163
 - summary of 49
- environment
 - for Info/Analysis 169
- ENVIRONMENT 185
- environment section 174, 185, 217
- EOJ (end of job) trace 57, 90
- ERASE
 - statement 200
- EREP program
 - use of
 - in response to VSE messages 235
- event record 30
 - instruction-execution trace 60
 - program-load trace 65
 - storage-alter trace 66
- events
 - limiting number of to be traced 160
 - number of, traceable per session 52
 - traceable by SDAID
 - See* trace type definition
 - tracing of
 - See* tracing events
- examples
 - call an analysis routine 191
 - delete a dump job 181
 - DOSVSDMP prints SDAID tape 31
 - DOSVSDMP utility
 - prints stand-alone dump 30
 - DUMP command dump
 - symptoms 211
 - dump symptom output 185
 - file labels for
 - dump management file 172
 - external routines file 172
 - SYSDUMP library 172
 - generate the stand-alone dump program 26
 - IJBXDEBUG output 194
 - Info/Analysis control statements 211
 - initialization of
 - dump management file 170
 - instruction-execution trace event record 60
 - interactive trace program 41
 - invoke Info/Analysis 176
 - label information area, display of 231
 - library chain listing 228
 - linkage editor output 225
 - list managed dumps 181
 - loading the
 - external routines file 171
 - map of virtual storage
 - by the linkage editor 225
 - offload a dump job 200
 - onload a dump job 204
 - print active partition of stand-alone dump tape 207

examples (*continued*)
 PRINT control statement
 output 182
 print DUMP command dump
 formatted 210
 print dump symptoms 184
 print formatted dump 190
 print main dump file of stand-alone
 dump tape 206
 print selective dump areas 189
 program-load trace event record 65
 prompting sequence
 program-load trace request 152
 storage-alter trace event record 66
 symptoms in a
 stand-alone dump output 208

extensions
 array 219
 chain 219
 hexadecimal 220
 keyfield 220
 text 219

extent definition table 6
 external interrupt trace 58, 145
 external interrupt 91
 general description 58
 in direct input mode 91

external routines file 169
 labels 171
 loading 171

F

file 169
 See also dump management
 external routines 169
 sequence number 203

FILE
 statement 203

floating point registers (FREG), display of
 on event 160

floating-point registers (FReg), display of
 on event 77

format
 DUMP command 21
 of a dump name 15

FORMAT
 for PRINT statement 189

formatted
 control blocks 221

formatted dump print 190

formatting descriptor 221

free-form symptoms 174, 187

freeform symptoms 218

functions
 dump management 179
 dump offload 198
 dump onload 202
 dump symptoms 184
 interactive mode 173
 dump viewing 188
 overview 172

G

generating a stand-alone dump
 program 25

GETVIS / FREEVIS trace (SDAID)
 general description 58

GETVis (getvis / freevis request)
 trace 146

GETVIS trace (SDAID)
 in direct input mode 93

glossary 236

GO command
 interactive trace program 39

GReg (general registers)
 display of on event 78

GREG (general registers)
 display of on event 160

H

halt (processing) on event 160

Halt option 81

halt temporarily
 program execution 233
 tracing of events 163

hardware aids 235
 alter/display feature 236
 instruction stepping feature 236
 overview 235
 stop-on-address-compare feature 236

header record 169
 symptom 217

help function, event tracing 143

hex addresses, to define trace range 158

hexadecimal
 extensions 220

high-speed dump
 See dump

high-speed dumps
 See dumps

I

I/O device range definition 158
 by channel address 159
 by control unit address 159
 by unit address 159

I/O interrupt trace
 device range definition 158
 general description 60
 in direct input mode 96
 in procedure mode 123

I/O request block (IORB)
 display of on event 74, 160

I/O tables and blocks (IOTab)
 display of on event 78

I/O tables and blocks (IOTAB)
 display of on event 160

IBM Support
 mailing a dump stored on tape 17
 sending a dump to, electronically 16

ICCF
 interfaces 169

IJBXCSMG 191
 activation 191
 functions 191
 invocation 191

IJBXDEBUG

activation 192
 invocation 191
 name loaded into
 external routines file 171

output 192
 address validation 192
 from hard wait dumps 193
 from loop dumps 194
 from soft wait dumps 194
 general dump information 192
 header information 192
 specific analysis information 193

output examples 194

IJBXSDA 191, 198
 activation 198
 functions 198
 invocation 191

Info/Analysis
 functions 10
 print dumps from tape 205
 prints DUMP command dump 210
 prints stand-alone dump 206

Info/Analysis control statements
 BYPASS 199
 CALL 190
 common 177
 DELETE 180
 DUMP NAME 178
 entering 176
 ERASE 200
 examples 211
 FILE 203
 PRINT 189
 dump management 181
 dump symptoms 184
 dump viewing 188
 reference 213
 RETURN 178
 SELECT 177
 summary 213
 syntax 176
 UTILITY 170, 180
 VOLID
 for dump offload 199
 for dump onload 202

Info/Analysis functions
 See functions

INFOANA
 invocation 176

information gathering
 aids for
 dump of disk, diskette, or
 tape 225
 LISTLOG utility 229
 LSERV (label service)
 program 230
 LVTOC program 232

information retrieval
 See printing

initialization of
 dump management file 170

input/output
 See I/O

instruction stepping feature 236

instruction trace (interactive trace
 program) 35

instruction trace (SDAID)
 general description 59
 in direct input mode 95
 in procedure mode 122

instruction-execution trace 59, 95, 147
 event record 60

integrated console 7

interactive trace commands
 ALTER 38
 DISPLAY 38
 GO 39
 QUERY 37
 TRACE definition 37
 TRACE END 37

interactive trace examples 41

interactive trace program 39
 ABEND trace 35
 activation 36
 branch trace 35
 commands
 See interactive trace commands
 examples 41
 instruction trace 41
 interactive trace program 35
 restrictions 40
 scope of tracing 40
 storage alteration trace 41
 trace activation 36

Interactive Trace Program
 tracing in a user partition with
 subtasks attached 39

interfaces 169

invalid address space
 display of virtual storage 227

invocation
 of Info/Analysis 176

invoking 163
 dump
 by the DUMP command 21
 from within a program 23
 dumps from within a program 23
 LVTOC program 233
 SDAID program 163
 stand-alone dump program 7

IO (I/O) interrupt trace 60, 96, 148
 for VTAM 154

IORB
 See command control block

IOTab (I/O tables and blocks), display of
 on event 78

IOTAB (I/O tables and blocks), display of
 on event 160

IPL load parameter 8

J

JDUMP macro 23

job cancel trace 57, 90

job control
 for Info/Analysis invocation 175

job control statements
 logging of on SYSLST 230

JOBNAME definition
 general description 70
 in direct input mode 111

JOBNUM definition 70

K

keyfield
 extension 220

L

label information
 for SYSDUMP library 13
 example 13

label information area
 display of 230, 231

label service (LSERV) program 230

language translator source code 229

LBD 219
 See locating block descriptor

LDT
 See library definition table

LIBDEF statement
 to establish the
 dump sublibraries 13

LIBR program
 used to define
 the dump sublibraries 13

library chain listing 228

library chains, display of 227

library definition table 6

library pointer block 6

life cycle of a dump 172

limiting occurrence of events to be
 traced 160

line mode 169, 175

linkage
 descriptor 220

linkage editor
 obtaining a map of virtual
 storage 223
 output of, example 225

linkage editor map
 example 225

list log utility (LIST LOG) 229

list option 229

listing
 device assignments 229
 job related SYSLOG
 communication 229

LISTIO command/statement 229

load a dump
 into a dump sublibrary 203

load parameter 8

LOADLS
 See phase load trace table

locating block descriptor 217
 See LBD

locator 219

LOCK / UNLOCK trace (SDAID)
 general description 61
 record/print Locktable entry 79

LOCK (lock / unlock of resources)
 trace 148

LOCK trace (SDAID)
 in direct input mode 97

Locktable entry (LOCKTE)
 display or print 79

LOCKTE
 See Locktable area

logging job control statements on
 SYSLST 230

logical transient area (LTA)
 display of on event 79, 160

logical unit block 6

logical unit block extension 6

loop
 prompting
 program-load trace request 151
 tracing events for 164

lost characters on IBM 3211 printouts xi

low address range of storage (LOWcore),
 contents of on event 79

low address range of storage
 (LOWCORE), contents of on event 160

low address storage
 display of
 on event 79, 160
 display/alter
 with DSPLY/ALTER
 commands 226

LOWcore
 See low address storage

LOWCORE
 See low address storage

LPT
 See library pointer block

LSERV (label service) program 230

LTA
 See logical transient area

LUBEXT
 See logical unit block extension

LUBTAB
 See logical unit block

LVTOC (list volume table of contents)
 program 232

LVTOC program
 invoking of 233

M

macros, dump invoking 23

mapping virtual storage
 by the linkage editor 223
 example 225
 map of virtual storage 223

modes of operation 169

MON (monitor call) trace 99, 150

monitor call trace
 general description 62
 in direct input mode 99

MS 186

multiple tape support (stand-alone
 dump) 22

N

name of dumps 15

no-dump option 20

NOJCL option 81

NOJCL specification 160

NOSource option 81

NOTarget option 81

notational conventions 134

notations, command xv

number of
 events per trace type 160
 trace types per session 52

O

OBJMAINT
 used to load
 external routines file 171
occurrence definition 160
OCCurrence option 81
offload
 See dump offload
OFFS 186
OFFset definition 112
onload
 See dump onload
OPCS 186
operating
 system 169
operating environment 169
operational hints
 dump sublibraries 11
 invoking
 dumps from within a program 23
 the SDAID program 163
 tracing events
 for a loop 164
 traces per SDAID session 52
operator console 169
OPTION definition
 Halt 81
 HALT 116
 in direct input mode 116
 NOJCL 81, 116
 NOSource 81, 116
 NOTarget 81, 116
 OCCurrence 81
 OCCurrence 116
 SUPervisor 81, 116
 TERminate 81
 TERMinate 116
OPTION statement
 for ABEND dump function 20
 for stand-alone dumps 9
 SADUMP option 9
 to store dumps
 into dump sublibraries 14
optional symptoms 174
non-SDB section 187, 218
SDB section 187, 218
options
 controlling the ABEND dump
 function 20
 invoking
 partition or system dumps 20
 listing language translator source
 code 229
 to activate dump storing
 OPTION SYSDDUMP 14
 STDOPT SYSDDUMP=YES 14
 to deactivate dump storing
 // OPTION NOSYSDDUMP 15
 LIBDROP DUMP,PERM 15
 STDOPT SYSDDUMP=NO 15
 UNBATCH command to deactivate
 the partition. 15

OUTDEV
 general description 53
OUTDEV command
 description 142
 prompting for 142
 summary of 131
 syntax diagram 135
OUTDEV specification
 description 53
OUTDEV statement
 in direct input mode 86
output
 ABEND dump function 20
 DUMP command dump
 formatted output 210
 from IJBXDEBUG 194
 language translator 229
 linkage editor 225
 lost characters on IBM 3211 xi
 of IJBXDEBUG 192
 of the
 stand-alone dump program 8
 routing 169
 SDAID
 from tape 30
 SDAID program
 definition of 160
 device for 53, 142
 to magnetic tape 53, 142
 to printer 53, 142
 to wraparound buffer 53, 142
 stand-alone dump
 printing of tape 29
output definition 160
OUTPut definition 113
output definition, tracing events 160
 buffer 74, 114
 CCB (command control block) 74,
 160
 CCW (channel command word) 74,
 114, 160
 CCW with data 74, 115
 COMREG 75, 160
 control registers 75, 160
 FReg (floating-point registers) 77
 FREG (floating-point registers) 160
 GReg (general registers) 78
 GREG (general registers) 160
 I/O tables and blocks 78, 160
 IORB (I/O request block) 74, 160
 LOCKTE (Locktable area) 79
 LOWcore (low address storage) 79
 LOWCORE (low address
 storage) 160
 LTA (logical transient area) 79, 160
 partition communication region 75,
 160
 partition-related control blocks 79
 PIB (program information block) 160
 PTA (physical transient area) 79, 160
 supervisor area 80, 160
 SYSCom 80
 SYSCOM 160
 task-related control blocks 80, 160
 time of day 80, 160
 virtual storage dump 115
 virtual storage dumps 76

output definition, tracing events
 (*continued*)
 XPCCB (XPCC communication control
 block) 80
 XPDATABU (XPCC data buffer) 81
output device 169
output device (SDAID)
 direct input mode 86
 general description 53, 73
 overview 50
 procedure mode 129
output device specification
 for tracing events
 See OUTDEV command

P

page manager address space (PMRAS)
 dump file 8
panels
 dump symptoms 173
partition communication region
 display of on event 75, 160
partition dump 3
 option for 20
 output contents 6
 produced by
 the CANCEL command 21
partition information block extension 6
pattern specification
 for storage-alter trace 66, 103, 153
PAUSE command/statement 233
PDUMP macro 23
PER
 See Program Event Recording
performance considerations, tracing of
 events 50
PHase definition 112
phase load trace table 6
physical transient area (PTA), display of
 on event 79, 160
physical unit block (PUB) 6
physical unit block table 6
PIB
 See program information block
PIB2TAB
 See partition information block
 extension
PIDS 186
plus sign 173
PMRAS (page manager address space)
 dump file 8
pointer section 217
prefix 174
print
 DUMP command dump
 formatted 210
 dump symptoms 184
 dumps from tape 204
 steps 205
 with Info/Analysis 205
 selective dump areas 189
 stand-alone dump formatted 206
 stand-alone dump unformatted 29
PRINT
 statement
 for dump management 181

PRINT (*continued*)
 statement (*continued*)
 for dump symptoms 184
 for dump viewing 188
 PRINT statement 189
 printed output
 See output
 printer definition
 procedure mode 129
 printing
 error information
 from SYSREC 235
 printing dumps
 sample job to
 invoke Info/Analysis 176
 printing the stored dump 183
 problem
 log 15
 problem data collection 172
 procedure
 creation for SDAID trace 119
 statement for SDAID trace 117
 to initialize SDAID traces 121
 procedure mode
 trace initialization 117
 program check 100
 trace 150
 program check trace 63
 general description 63
 in direct input mode 100
 in procedure mode 125
 Program Event Recording (PER)
 interactive trace program 40
 synchronization with the interactive
 trace program 40
 program information block
 display of on event 160
 program load trace 64, 101, 151
 event record 65
 general description 64
 in direct input mode 101
 in procedure mode 124
 prompting sequence for 152
 prompting control information 143
 additional output 160
 creating a stand-alone dump
 program 25
 event occurrence definition 160
 event type 144
 halt on event 160
 I/O definition in TRACE
 command 158
 OUTDEV command 53, 142
 output definition 160
 TRACE command 143
 trace type 144
 tracing events 143
 number of event occurrences 160
 PTA (physical transient area), display of
 on event 79, 160
 PTAB (partition-related control
 blocks) 79, 160
 PTRACE 40
 PUB (physical unit block) 6
 PUB ownership table 6
 PUBOWN
 See PUB ownership table

PUBTAB
 See physical unit block table

Q

QUERY command
 interactive trace program 37
 question mark 134
 single, request help function 143

R

range of trace
 defining of
 by occurrence definition 160
 of devices, for I/O events 158
 reader mode 169, 175
 READY command
 summary of 131
 record
 See also header record
 See also symptom records
 header 169
 recording hardware failures
 controls for
 via the console 235
 recording time of day, on event 80, 160
 recovery management support
 control of
 via the console 235
 recovery of data 225
 REGS 186
 relative addresses, to define the trace
 range 157
 required symptoms 186
 required symptoms section 174, 186, 217
 restrictions
 display of virtual storage 227
 program-load trace request prompting
 loop 151
 SDAID tape output 53
 using an IBM 3211 with indexing xi
 retrieval of error information
 See printing
 RETURN
 statement 178
 RIDS 186
 ROD command 235
 routines
 See analysis routines

S

SADUMP option 9
 scope of tracing 40
 SDAID
 general description 53
 overview 45
 session 47
 trace options 81
 SDAID buffer formatting 191
 SDAID buffer sizes 143
 SDAID commands
 prompting for 143
 SDAID commands, summary 131
 SDAID dump trace 23

SDAID program 47, 131
 ending execution 163
 invoking of 163
 output from tape 30
 prompting control information 143
 storage requirements 52
 SDAID trace
 additional keywords 128
 SDAID trace initialization
 direct input mode 48
 in direct input mode 83
 in procedure mode 117
 overview 47
 procedure mode 48
 SDAID trace procedures
 summary 121
 SDAID trace types
 summary 49, 88, 89, 110, 111, 144,
 145
 summary(general description) 56
 SDAID wait states 165
 SDAID wraparound
 See wraparound buffer
 SDB
 See also required symptoms section
 See structured data base
 SDT
 See sublibrary definition table
 SDUMP macro 23
 SDUMPX macro 23
 section 6 218
 SELECT statement 177
 selection
 level 177
 sequence of commands, SDAID
 program 131
 serviceability aids
 aids for
 DSPLYV message reply 232
 CANCELV message reply 232
 display of
 library chains 227
 dump invoking macros 23
 dump libraries 11
 dump sublibraries 11
 dump utility DOSVSDMP 25
 hardware 235
 library chain listing 228
 list log utility 229
 listing I/O device assignments 229
 LSERV program 230
 map of virtual storage
 by the linkage editor 223
 stand-alone dump 7
 single question mark 134
 SSCH (start subchannel) trace 65, 103
 device range definition 158
 for VTAM 154
 SSCH (start subchannel) Trace 152
 SSCH instruction trace
 general description 65
 in direct input mode 103
 in procedure mode 123
 stand-alone dump 25, 172, 218
 dumping priority (OPTION
 SADUMP) 9
 formatted output 208

- stand-alone dump (*continued*)
 - formatted print 206
 - multiple tape support 22
 - OPTION SADUMP 9
 - printed output 208
 - requesting a 21
 - SADUMP option 9
 - saving machine information 22
 - STORE STATUS command 22
 - unformatted output 29
- stand-alone dump program 7
 - creating 26
 - output 8
 - unformatted output 29
 - when to use 7
- start I/O trace 152
- start SDAID after Halt 81
- start SDAID trace initialization
 - in direct input mode 85
 - overview 49
- start subchannel trace 65, 103
- start tracing of events 163
- STARTSD command
 - STARTSD 163
 - summary of 49
- statements
 - See also* control statements
 - See also* end-of-input
 - See also* end-of-job
 - See* job control
- STDOPT command
 - for ABEND dump function 20
 - SADUMP option 9
 - to store dumps
 - into dump sublibraries 14
- STOP command 233
- stop temporarily
 - program execution 233
 - tracing of events 163
- stop-on-address-compare feature 236
- STOPSD command
 - description 163
 - summary of 49
- storage alteration trace (interactive trace program) 35
- storage alteration trace (SDAID)
 - general description 66
 - in direct input mode 103
 - in procedure mode 126
- storage requirements
 - SDAID program 52
- storage-alter trace 66, 103, 153
 - event record 66
 - pattern for 66, 103, 153
- structured data base 218
- structured data base (SDB) 187
- sublibrary definition table 6
- sublibrary name 15
- SUP (supervisor area), dumping on
 - event 80
- supervisor area, contents of on event 80, 160
- SUPervisor option 81
- Supvr (supervisor area), dumping on
 - event 160
- suspend program execution 233
- suspend tracing of events 163

- SVC (supervisor call) trace 67, 105, 153
 - for VTAM 154
- SVC trace
 - general description 67
 - in direct input mode 105
 - in procedure mode 127
- symptom part
 - DUMP command dump 211
 - of an ABEND Dump 5
 - stand-alone dump 208
- symptom records 3
 - creation 218
 - environment section 174, 185, 217
 - formatting descriptors 221
 - introduction 172
 - linkage descriptor 220
 - locator 219
 - optional symptoms 174
 - non-SDB section 187, 218
 - SDB section 187, 218
 - overview 217
 - pointer section 217
 - required symptoms section 174, 217
 - section 6 218
- symptoms 173
 - See also* dump symptoms
 - See* symptom record
- syntax
 - for Info/Analysis control statements 176
- syntax diagram 134
 - OUTDEV command 135
 - TRACE command 136
- syntax symbols xv
- syntax, of commands xv
- SYSCOM
 - See* system communication region
- SYSCOM
 - See* system communication region
- SYSDUMP
 - label information 11
 - library 11
 - option 11, 14
- SYSDUMP library 11
 - concept 11
 - labels 13
 - example 13
 - requirements 13
- SYSDUMP used by
 - ABEND dump function 20
 - CANCEL command 21
- SYSFIL
 - See* system file buffer
- SYSLST 169
- SYSREC file, printing of 235
- system
 - input device 169
 - operating 169
 - output device 169
- system areas
 - dumped 208
- system communication region, display of
 - on event 80, 160
- system dump 3
 - option for 20
 - output contents 5

- system dump (*continued*)
 - produced by
 - the CANCEL command 21
- system dump library
 - See* dump library
- system file buffer 6

T

- tables
 - SDAID command summary 131
- tape
 - mailing to IBM Support (containing a dump) 17
- tape data, dumping of 225
- tape definition
 - procedure mode 129
- task control block 6
- task information block 6
- TCB
 - See* task control block
- temporarily stop
 - program execution 233
 - tracing of events 163
- TERM keyword operand 129
- TERM=CANCEL 129
- TERM=EXT 129
- TERM=PGMC 129
- Terminate option 81
- termination 211
 - See* ending
- text
 - extension 219
- TIB
 - See* task information block
- time of day, recording of on event 80, 160
- trace buffer definition
 - procedure mode 129
- TRACE command 143
 - prompting for
 - output definition 160
 - trace type 144
 - summary of 131
 - syntax diagram 136
- TRACE command (interactive) 36
- trace commands (interactive trace program)
 - ALTER 38
 - DISPLAY 38
 - GO 39
 - QUERY 37
 - TRACE 37
- TRACE commands (SDAID)
 - prompting for 143
- trace examples (interactive trace program)
 - trace initialization 41
- trace output, overview 50
- TRACE statement
 - in direct input mode 88
- trace type definition 91, 144
 - branch 145
 - branch-execution 89
 - buffer 145
 - buffer-overflow 90
 - cancel 145
 - cancel (end-of-job) 90

trace type definition (*continued*)

- external interruption 145
- GETVis 146
- GETVIS 93
- instruction execution 95, 147
- IO (I/O) interrupt 96, 148
- LOCK 97, 148
- MON (monitor call) 99, 150
- program check 100, 150
- program load 101, 151
- SSCH (start subchannel) 103, 152
- storage alter 103, 153
- summary 144
- SVC (supervisor call) 105, 153
- VTAM I/O 107
- VTAMBU 154
- VTAMBU (VTAM buffer usage) 107
- XPCC 108, 154

trace type description

- branch-execution 56
- buffer-overflow 57
- cancel (end-of-job) 57
- external interrupt 58
- GETVIS / FREEVIS 58
- instruction execution 59
- IO (I/O) interrupt 60
- LOCK / UNLOCK 61
- program check 63
- program load 64
- SSCH (start subchannel) 65
- storage alter 66
- SVC (supervisor call) 67
- VTAM I/O 68
- VTAMBU (VTAM buffer usage) 68
- XPCC 69

tracing branch instructions

- using SDAID program 145

tracing buffer

- using SDAID program 145, 154

tracing cancel events

- using SDAID program 145

tracing events 47, 131

- additional output for 160
- command summary 131
- command syntax summary 134
- control information for
 - syntax diagrams 134
- defining the I/O range 158
- defining type of trace 144
- dump on event 76, 115
- ending control information input 163
- ending of 160
- for a loop 164
- information input
 - See* control information input
- invoking the SDAID program 163
- number of trace types per session 52
- output device for 53, 142
- overview 47, 131
- performance considerations 50
- prompting loop
 - program-load trace request 151
- start after control-information input 163
- stopping of temporarily 163
- storage requirements for 52

tracing external interruptions

- using SDAID program 145

tracing getvis / freevis request

- using SDAID program 146

tracing in a user partition with subtasks attached 39

tracing instruction execution

- using instruction stepping feature 236
- using SDAID program 59, 95, 147

tracing interactively 35

tracing lock / unlock of resources request

- using SDAID program 148

tracing partition communication request

- using SDAID program 154

tracing, scope of 40

TTAB (task-related control blocks) 80, 160

types of traces

- See* trace type definition

U

UNit definition 82, 113

UTILITY statement 170, 180

V

viewing dumps

- See* dump viewing

virtual storage

- altering
 - See* alter
- display of
 - See* display
- mapping of
 - by the linkage editor 223

virtual storage dump

- See* dump

virtual storage dumps

- See* dumps

virtual storage map

- See* example, map of virtual storage

VOLID

- statement
 - for dump offload 199
 - for dump onload 202

volume

- device for offload 199
- device for onload 202

volume table of contents

- display of
 - by LVTOC 233
 - by LVTOC program 232
 - by response to messages 232
- free space
 - by LVTOC program 232

VSE/Advanced Functions 169

VSE/ICCF 169

VTAM buffer trace 154

- general description 68
- in direct input mode 107

VTAM traces

- device range definition 158
- VTAM buffer usage 68, 107
- VTAM I/O events 68, 107

VTAMBU (VTAM buffer usage) trace 68, 107

VTAMIO trace

- general description 68
- in direct input mode 107

VTOC

- See* volume table of contents

W

wait state

- caused by SDAID 165

wraparound buffer 90

- contents of on overflow 57, 90
- tracing of overflow 57

X

XPCC (partition communication)

- trace 154

XPCC communication control block (XPCCB)

- display or print 80

XPCC data buffer (XPDATABU)

- display or print 81

XPCC trace (SDAID)

- general description 69
- in direct input mode 108

XPCCB

- See* XPCC communication control block

XPDATABU

- See* XPCC data buffer

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

IBM z/VSE
Diagnosis Tools
Version 3 Release 1

Publication No. SC33-8229-00

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

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

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

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

Please tell us how we can improve this book:

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

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

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Deutschland Entwicklung GmbH
Department 3248
Schoenaicher Strasse 220
D-71032 Boeblingen
Federal Republic of Germany



Fold and Tape

Please do not staple

Fold and Tape



File Number: S370/S390-37
Program Number: 5609-ZVS

Printed in USA

SC33-8229-00



Spine information:



z/VSE

Diagnosis Tools

Version 3 Release 1