

Systems

MVS/Extended Architecture Conversion Notebook

Your conversion hints, please ...

This book will be updated as additional information becomes available. You can submit conversion hints for possible publication in this book. Use the reader's comment form or the conversion notebook input form at the back of this book or send your information to:

IBM Corporation
Publications Department
Department D58, Building 920-2
PO Box 390
Poughkeepsie, New York 12602
ATTN: MVS/Extended Architecture Conversion Notebook

When submitting hints, see the Preface for details.



Third Edition (May, 1984)

This is a major revision of, and obsoletes, GC28-1143-1. See the Summary of Amendments following the Contents for a summary of the changes made to this manual.

This edition applies to the following program products until otherwise indicated in new editions or technical newsletters:

- Assembler H Version 2 (5668-962)
- MVS/XA Data Facility Product (DFP) (5665-284)
- MVS/System Product - JES2 (MVS/SP) Version 2 (5740-XC6)
- MVS/System Product - JES3 (MVS/SP) Version 2 (5665-291)

Changes are continually made to the information herein; before using this publication in connection with the operation of IBM systems, consult the latest *IBM System Bibliography*, GC20-0001, for the editions and technical newsletters that are applicable and current.

It is possible that this material might contain references to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

Publications are not stocked at the address given below; requests for copies of IBM publications should be made to your IBM representative or to the IBM branch office serving your country.

A form for reader's comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Publications Development, Department D58, Building 920, PO Box 390, Poughkeepsie, New York 12602. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

Preface

This book is for two audiences:

- People responsible for converting to MVS/Extended Architecture (MVS/XA) from an OS/VS2 MVS Release 3.8 system with at least one of the following installed:
 - For JES2 users, MVS/SP Version 1 Release 3.0 or a later release
 - For JES3 users, MVS/SP Version 1 Release 3.1 or a later release

In this publication, MVS/370 refers to operating systems at one of those levels. Although it is possible to convert other releases of MVS/370 to MVS/XA, this book does not describe how to do so.

- People who have already converted to MVS/XA and are installing subsequent releases of MVS/SP Version 2 and MVS/XA Data Facility Product (DFP).

Readers are expected to have an in-depth knowledge of MVS/370, the configuration and procedures of the current installation, and the configuration of the target installation. They also need to be familiar with the MVS/XA overview information in the Licensed Programming Announcement letters and General Information Manuals for MVS/SP Version 2 and MVS/XA DFP. Reading the *MVS/XA Overview* is also helpful.

This book contains conversion information related to the program products required for MVS/XA:

- MVS/System Product - JES2 (MVS/SP) Version 2 (5740-XC6)
- MVS/System Product - JES3 (MVS/SP) Version 2 (5665-291)
- MVS/XA Data Facility Product (DFP) (5665-284)
- Assembler H Version 2 (5668-962)

Frequently, the book refers to different releases of MVS/SP Version 2 and MVS/XA DFP by the release number only:

- Release 1.0 refers to MVS/SP Version 2 Release 1.0 and MVS/XA DFP Release 1.0
- Release 1.1 refers to the Release 1.1 levels of those products
- Release 1.2 refers to their Release 1.2 levels

The information in this book is organized as follows:

“Chapter 1: Introduction” summarizes the work required to convert to MVS/XA, lists conversion tasks you can do on your MVS/370 system before receiving MVS/XA, and describes changes in the MVS/XA library.

“Chapter 2: Installation and Initialization” includes information related to installing and initializing an MVS/XA system and generating stand-alone dump.

“Chapter 3: Programming Considerations” describes changes that might affect user-written code, including changes to assembler language instructions and macros. It also describes new function available to programmers.

“Chapter 4: Operating Considerations” describes new and changed commands and operational procedures.

“Chapter 5: System Modifications” describes new and changed user exits and ways of tailoring the system.

“Chapter 6: Problem Determination” describes new and changed ways of tailoring and suppressing dumps, new and changed dump formats, trace facilities, and debugging considerations.

“Chapter 7: Accounting” describes changes that might affect your accounting procedures.

“Chapter 8: Measurement and Tuning” describes changes related to performance.

“Chapter 9: Coexistence Considerations” contains considerations for running both MVS/370 and MVS/XA in a single installation.

“Appendix A: Parameter Changes in Incompatible Macros” describes differences between the MVS/370 and MVS/XA parameters lists that downward incompatible macros pass to their service routines.

“Appendix B: Control Block Changes” lists control blocks that are new, changed, or deleted or that can reside anywhere in virtual storage (above or below 16 Mb).

The *Conversion Notebook* does not describe:

- How to install the program products. The *Program Directory* shipped with the product describes the installation procedure.
- JES conversion information other than “Installing the JES2 Component of MVS/SP - JES2 Version 2” in Chapter 2. *MVS/SP-JES2 1.3.3, MVS/SP-JES2 2.1.1 JES2 Migration Considerations*, a technical bulletin from the Washington Systems Center, provides an overview of new JES2 functions, conversion information, a sample migration plan, and results of field text experiences.
- The virtual storage constraint relief that MVS/XA provides. Technical bulletin, *Virtual Storage Tuning Cookbook*, provides that information. In addition, it describes how each subsystem uses virtual storage, tools available for measuring virtual storage use, rules for using virtual storage efficiently, and techniques for tuning subsystems for efficient storage use.
- How to write programs that execute in 31-bit addressing mode. *SPL: 31-bit Addressing* contains that information. The *Conversion Notebook* does, however, describe system changes that take advantage of or support 31-bit addressing and the impact the changes have on user-written programs. For example, the *Conversion Notebook* lists control blocks that have been moved to the extended

area of virtual storage and gives an example of how you can change programs to access them.

The phrase 'published external interfaces' appears several times in the book. It refers to interfaces documented in the following publications:

- *OS/VS2 MVS JCL*, GC28-0692
- *OS/VS2 Supervisor Services and Macro Instructions*, GC28-1114
- *OS/VS2 TSO Command Language Reference*, GC28-0646
- *OS/VS2 Guide to Writing a Command Processor or Terminal Monitor Program*, GC28-0648
- *MVS/370 Data Management Macro Instructions*, GC26-4057
- *MVS/370 Access Method Services Reference for the Integrated Catalog Facility*, GC26-4051
- *MVS/370 Access Method Services Reference for VSAM Catalogs*, GC26-4059
- *MVS/370 Virtual Storage Access Method (VSAM) User's Guide*, GC26-4066

Submitting Conversion Hints

This book will be updated as more information becomes available. You can submit conversion hints for possible publication in this book. Use the reader's comment form or the conversion notebook input form at the back of this book or send your information to:

IBM Corporation
Publications Department
Department D58, Building 920
PO Box 390
Poughkeepsie, New York 12601
ATTN: MVS/Extended Architecture Conversion Notebook

It is understood that IBM and its affiliated companies shall have the nonexclusive right, in their discretion, to use, copy, and distribute all submitted information or material, in any form, for any and all purposes, without any obligation to the submitter, and that the submitter has the unqualified right to submit such information or material upon such basis.

When submitting conversion hints, please indicate from what system you are converting, the program products installed on it, and the program products being installed on the MVS/XA system.

Related Publications

The *Conversion Notebook* highlights differences in MVS/XA to help you identify changes you need to make to existing procedures, programs, and system modifications. You will, however, need other books in the MVS/XA library to make any required changes to your procedures and programs. The following table lists all books referred to in the *Conversion Notebook*. The table shows the short title used in the reference, the corresponding full title, and the order number of the book. For a complete list of the publications that support MVS/XA, see the general information manuals.

Short Title Used in This Publication	Title	Order No.
<i>Principles of Operation</i>	<i>IBM System/370 Extended Architecture Principles of Operation</i>	SA22-7085
<i>Operator's Guide</i>	<i>IBM 3081 Operator's Guide for the System Console</i>	GC38-0034
<i>Operator's Guide</i>	<i>IBM 3083 Operator's Guide for the System Console</i>	GC38-0036
<i>Operator's Guide</i>	<i>IBM 3084 Operator's Guide for the System Console</i>	GC38-0037
<i>IOCP User's Guide and Reference</i>	<i>Input/Output Configuration Program User's Guide and Reference</i>	GC28-1027
****	<i>Device Support Facilities User's Guide and Reference</i>	GC35-0033
<i>Data Facility Product: Planning Guide</i>	<i>MVS/Extended Architecture Data Facility Product: Planning Guide</i>	GC26-4040
<i>DFP General Information Information</i>	<i>MVS/Extended Architecture Data Facility Product General Information</i>	GC26-4007
<i>Diagnostic Techniques</i>	<i>MVS/Extended Architecture Diagnostic Techniques</i>	LY28-1199
****	<i>MVS/Extended Architecture General Information Manual</i>	GC28-1118
****	<i>MVS/Extended Architecture Interactive Problem Control System Guide and Reference</i>	GC28-1297
<i>Linkage Editor and Loader</i>	<i>MVS/Extended Architecture Linkage Editor and Loader</i>	GC26-4011
<i>System Commands</i>	<i>MVS/Extended Architecture Operations: System Commands</i>	GC28-1206
<i>MVS/XA Overview</i>	<i>MVS/Extended Architecture Overview</i>	GC28-1348
<i>RMF Reference and User's Guide</i>	<i>MVS/Extended Architecture Resource Measurement Facility (RMF) Version 3 Reference and User's Guide</i>	LC28-1138
<i>Supervisor Services and Macro Instructions</i>	<i>MVS/Extended Architecture Supervisor Services and Macro Instructions</i>	GC28-1154
<i>System Generation Reference</i>	<i>MVS/Extended Architecture Installation: System Generation</i>	GC26-4009
<i>System-Data Administration</i>	<i>MVS/Extended Architecture System-Data Administration</i>	GC26-4010
<i>SPL: Initialization and Tuning</i>	<i>MVS/Extended Architecture System Programming Library: Initialization and Tuning</i>	GC28-1149
<i>SPL: JES2 Initialization and Tuning</i>	<i>MVS/Extended Architecture System Programming Library: JES2 Initialization and Tuning</i>	SC23-0065
<i>SPL: JES3 Initialization and Tuning</i>	<i>MVS/Extended Architecture System Programming Library: JES3 Initialization and Tuning</i>	SC23-0059

Short Title Used in This Publication	Title	Order No.
<i>SPL: Service Aids</i>	<i>MVS/Extended Architecture System Programming Library: Service Aids</i>	GC28-1159
<i>System Macros and Facilities</i>	<i>MVS/Extended Architecture System Programming Library: System Macros and Facilities Volumes 1 and 2</i>	GC28-1150 GC28-1151
<i>SPL: SMF</i>	<i>MVS/Extended Architecture System Programming Library: System Management Facilities</i>	GC28-1153
<i>SPL: System Modifications</i>	<i>MVS/Extended Architecture System Programming Library: System Modifications</i>	GC28-1152
<i>SPL: User Exits</i>	<i>MVS/Extended Architecture System Programming Library: User Exits</i>	GC28-1147
<i>SPL: 31-Bit Addressing</i>	<i>MVS/Extended Architecture System Programming Library: 31-Bit Addressing</i>	GC28-1158
<i>Utilities</i>	<i>MVS/Extended Architecture Utilities</i>	GC26-4018
<i>MSS Installation Planning and Table Create</i>	<i>OS/VS Mass Storage Subsystem (MSS) Installation Planning and Table Create</i>	GC35-0028
<i>Global Resource Serialization</i>	<i>OS/VS2 MVS Planning: Global Resource Planning</i>	GC28-1062
<i>RMF General Information</i>	<i>Resource Measurement Facility General Information</i>	GC28-1115
<i>*****</i>	<i>TSO Extensions (TSO/E) General Information Manual</i>	GC28-1061

Contents

Chapter 1. Introduction 1-1

- Conversion Tasks You Can Do Before Installing MVS/XA 1-2
- Publications Changes 1-3
 - New Books 1-3
 - Changes to SPL: Supervisor and SPL: Job Management 1-3
 - Changes to Diagnostic Techniques and Program Logic Manuals 1-4
 - MVS/XA DFP Information 1-4
 - Title Changes 1-5

Chapter 2. Installation and Initialization 2-1

- Installing an MVS/XA System 2-1
 - Performing a Full Sysgen 2-2
 - SMP/E APPLY Processing 2-2
 - Required Environment for Generating MVS/XA 2-2
 - Providing a Backup Copy of the Existing System 2-2
 - Defining Additional Devices 2-3
 - Creating a New IOCDs 2-3
 - New I/O Configuration Requirements 2-5
 - Coding Sysgen Macros 2-5
 - Devices Not Supported 2-6
 - Rebuilding Alternate Eligible Device Tables (EDTs) 2-7
 - Changes to the Program Properties Table (PPT) 2-7
 - Defining System Data Sets 2-8
 - Initializing DASD 2-10
 - Loading the IPL Programs 2-10
 - Loading the Microcode EC Tapes for Mass Storage Subsystems 2-10
 - Installing the JES2 Component of MVS/SP - JES2 Version 2 2-11
- System Parameter and SYS1.PARMLIB Considerations 2-12
 - Fixed Storage for SLIP Command Processors (IEACMD00) 2-12
 - Specifying the RSU Parameter (IEASYSxx) 2-12
 - Increasing the Minimum SQA Allocation (IEASYSxx) 2-13
 - Specifying the Size of CSA and SQA Above 16 Mb (IEASYSxx) 2-13
 - Minimizing Private Area Storage Lost Because of Rounding (IEASYSxx) 2-14
 - Specifying Dump Data Sets (IEASYSxx) 2-15
 - Requesting Storage for RMF I/O Measurements (IEASYSxx) 2-15
 - Controlling the Number of Available ASVT Entries (IEASYSxx) 2-15
 - Removing TRACE Commands from COMMNDxx PARMLIB Members 2-16
 - Updating the IEAFIXxx PARMLIB Member 2-16
 - Removing References to Device Allocation Tables (IEALPaxx) 2-16
 - Keeping RNLs in GRSRNLxx PARMLIB Members 2-16
 - Specifying MIH Intervals (IECIOSxx) 2-17
 - New, Changed, or Deleted PARMLIB Members 2-17
- SYS1.PROCLIB Changes 2-23
 - DUMPSRV Procedure 2-24
 - IEESYSAS Procedure 2-24
 - LLA Procedure 2-24
 - PRDMP Procedure 2-24
 - RMF Procedure 2-25
- Using Default RNLs 2-26
 - Duration of the RMF Initialization Process 2-26
 - Generating Stand-Alone Dump 2-26
 - Stand-Alone Dump Macro Instruction Changes 2-27

Chapter 3. Programming Considerations 3-1

- Changes that Might Affect Unauthorized Programs 3-1
- Changes that Might Affect Authorized Programs 3-1
- 31-bit Addressing Considerations 3-2
- New Function 3-2
- Macro Instructions Mentioned in This Publication 3-3
 - CHKPT Macro Instruction 3-6
 - IOHALT Macro Instruction (SVC 33) 3-7
 - IOGEN UCBL00K Macro Instruction 3-7
 - RESETPL (BTAM) Macro Instruction 3-8
 - Differences in SPIE Processing 3-8
 - STATUS STOP,SYNCH Macro Instruction 3-9

- SDWA Changes 3-9
- Differences in GETMAIN Processing 3-9
- TSO TEST Command 3-10
- Deleted Instructions 3-11
- Macro Expansions in JES Modifications 3-11
- Limiting Concurrent Global Resource Serialization Requests 3-11
- Format Changes to Hardcopy Log Records 3-12
- Link Editing Allocation User Routines 3-13
- Removal of the Interval Timer 3-13
- Checklist for Determining if Authorized Programs Must be Changed 3-13
- Changes to the SVC Table 3-17
- Changes to the Locking Structure 3-17
- Determining Which Locks a Processor Holds 3-17
- Page Protection 3-17
- PSA Low Address Protection 3-18
- Fetch-Protected PSA Areas 3-19
- Patch Areas in the PSA 3-19
- Real Addressing Considerations 3-19
 - Using the EXCPVR Macro Instruction 3-19
 - Changes in the Way RSM Backs Virtual Storage 3-20
 - DAT-off Restrictions 3-21
- Cross Memory Entry Table Entries 3-22
- Interfaces to System Services 3-22
 - Services Independent of Addressing Mode 3-23
 - Services with Some Restrictions on the Address Parameter Values 3-24
 - Services that Do Not Support 31-bit Addressing 3-24
- 31-bit Addressing Considerations 3-24
 - Impact of 31-bit Addressing on Programmers 3-25
 - Changing Addressing Mode 3-26
 - Establishing a Program's Addressing Mode 3-26
 - Restrictions on Using a Linkage Editor Overlay Structure 3-28
- Changed Instructions 3-28
 - BAL and BALR (Branch and Link) Instructions 3-28
 - BAS and BASR (Branch and Save) Instructions 3-29
 - CLCL, EDMK, MVCL, and TRT Instructions 3-29
 - LA (Load Address) Instruction 3-29
 - LRA (Load Real Address) Instruction 3-29
- New Instructions 3-30
 - BSM (Branch and Set Mode) Instruction 3-30
 - BASSM (Branch and Save and Set Mode) Instruction 3-31
- Modifying Programs that Invoke Modules Above 16 Mb 3-31
 - Using BASSM and BSM Instructions 3-32
 - Using Linkage Assist Routines 3-32
- Retrieving Data from a Control Block Above 16 Mb 3-34
- Performing I/O in 31-bit Addressing Mode 3-35
- Using the EXCP Macro 3-36
- Entry Points in IEFW21SD 3-37
- Summary of New and Changed Macros 3-37
- New Parameters on the GETMAIN Macro Instruction 3-42
 - VRC and VRU Parameters 3-42
 - LOC Parameter 3-43
- SDUMP Macro Instruction 3-43
- SETLOCK RELEASE,TYPE=(reg) | ALL Macro Instruction 3-43
- Using GTF to Trace User Events 3-44
- Unit Verification 3-44
 - IEFAB4UV 3-44
 - IEFEB4UV 3-45

Chapter 4. Operating Considerations 4-1

- Loading 370-XA Microcode at Power-on Reset Time 4-1
- SYSTL (SCP Manual CNTL) Console Frame 4-1
- Storing Status Before Taking a Stand-alone Dump 4-3
- Using Labeled Tapes for Stand-alone Dumps 4-3
- JCL Changes to Jobs that Allocate SYS1.DUMP Data Sets 4-3
- Processing Hot I/O Interrupts 4-3
- Extended Color Support on 3279 MCS Consoles 4-4
- Controlling Message Traffic on Operator Consoles 4-5
- New Response to Message IOS201E 4-6
- Summary of New, Changed, or Deleted Commands 4-6

Chapter 5. System Modifications 5-1

- Print Dump Exit Control Table (ECT) Modifications 5-1
- Updating SYSTEMS Exclusion RNLs 5-1
- Serializing VSAM Data Sets 5-1
- Limiting User Region Size using IEFUSI Instead of IEALIMIT 5-2
- Obtaining an Extended Region Size of More Than 32 Mb 5-3
- Bypassing the Storage Availability Check Before a Job Executes 5-3
- Changing the Hot I/O Threshold and Recovery Actions 5-4
- Pre-dump Exits 5-4
- Post-dump Exits 5-4
- RMF Exits 5-4
- JES2 User Exits 5-4
- JES2 Interfaces 5-5
- JES3 Dynamic Support Programs (DSPs) and User Exits 5-5
- PRDMP Exits 5-6
- PRDMP Header Exits 5-6
- SMF Exits 5-6
- New WTO/WTOR User Exits 5-7
- New Services for Dump Processing Exits 5-8
 - Exit Services Router 5-8
 - Format Model Processor Service 5-9
 - Control Block Formatter Service 5-9
 - ECT Service 5-10
 - GET Symbol Service 5-10
 - EQUATE Symbol Service 5-10
 - Select ASID Service 5-10

Chapter 6. Problem Determination 6-1

- New and Changed Dump Options 6-2
- New Symptom Dumps for Task-Mode Abends 6-4
- New User Summary Dumps 6-5
- Dump Format Changes 6-6
 - Changes to User Dump Headers 6-6
 - User Dump Indexes 6-6
 - Changes to SYSMDUMP and SVC Dump Formats 6-6
- Suppressing Dumps 6-7
 - New Operands on the SLIP Command for Suppressing Dumps 6-7
 - MVS/XA's Use of SLIP Commands 6-7
 - Dump Analysis and Elimination (DAE) 6-8
- New and Changed PRDMP Control Statements 6-10
- Print Dump Index 6-12
- Print Dump Requirements for Printers 6-13
- New and Changed IPCS Subcommands 6-13
- Accessing Additional Sources of Dump Data Using IPCS 6-16
- New IPCS Panels 6-16
- Changes to the IPCS BROWSE Panels 6-17
- Changes to the Titles of IPCS Print Files 6-18
- Using the MVS/XA Versions of IPCS and PRDMP on Other Systems 6-18
 - Copying Release 1.2 IPCS and PRDMP Modules and Data Sets 6-19
 - Copying Release 1.0 and 1.1 IPCS and PRDMP Modules and Data Sets 6-19
- Debugging Considerations 6-20
 - Changes to the System Trace Facility 6-20
 - SDWA Changes 6-22
 - Addressing Mode Reflected in Dumps 6-22
 - Specifying Reason Codes 6-23
 - System Termination Facility Wait State Codes 6-23
 - Exceeding the Region Limit 6-23
 - Diagnosing Checkpoint/Restart Errors 6-24

Chapter 7. Accounting 7-1

- Device Connect Time 7-1
- New Fields Measuring Virtual Storage Use 7-2
- SMF30PRV and SMF30SYS Fields 7-2
- Type 22 SMF Record Changes 7-3
- Increases in EXCP Counts for Program Fetch Activity 7-3
- Summary of SMF Record Changes 7-4
- SMF Compatibility Between Release 1.0 and Later Releases 7-5

Chapter 8. Measurement and Tuning 8-1

- Ensuring Optimal Program Fetch Performance 8-1
- Performance Related Changes to the Linkage Editor and IEBCOPY 8-2
- Performance Related Changes to Program Fetch 8-2
- Recommended Actions 8-3
- Maintaining Count Values and Optimal Block Sizes 8-4
- Factors Affecting Text Block Sizes 8-6
- Using a New Directory for LNKLST Data Sets 8-7
- Starting the LLA Function 8-8
- Including Data Sets that Are Not APF Authorized 8-8
- Updating the LLA Directory 8-8
- SMF Data Set Placement 8-9
- Using the ASM Backing Slot Function 8-9
- Using Residency Time to Calculate the Page-in Rate of an Address Space 8-9
- Changes to ASM's Paging Algorithms 8-9
- Changes to the Data Set Selection Algorithm 8-10
- Changes to the Slot Selection Algorithm 8-10

Chapter 9. Coexistence Considerations 9-1

- Maintaining Programs that Can Run on Both MVS/370 and MVS/XA Systems 9-1
- Assembling and Link Editing Programs 9-2
- Guidelines for Ensuring Program Compatibility 9-2
- Guidelines for Developing New Programs 9-3
- Handling Downward Incompatible Macros 9-4
- Downward Incompatible SYNCH Macros 9-7
- Backup Considerations 9-7
- Routing Jobs in a Mixed JES2 or JES3 Complex 9-8
- Using Global Resource Serialization 9-8
- System Data Sets that Cannot be Shared 9-9
- Using SYS1.PROCLIB in a Loosely-coupled JES3 Configuration 9-9
- DSI Procedures in a Loosely-coupled JES3 Configuration 9-9

Appendix A. Parameter Changes in Incompatible Macros A-1

- ATTACH Parameter List Changes A-1
- ESTAE Parameter List Changes A-2
- EVENTS Parameter Changes A-3
- SMFEXIT Parameter List Changes A-3
- STAX Parameter List Changes A-3
- STIMER Parameter Changes A-3
- SYNCH Parameter List Changes A-4
- WTOR Parameter List Changes A-4

Appendix B. Control Block Changes B-1

Index X-1

Figures

- 2-1. Obsolete Sysgen Macro, Keywords, and Options 2-6
- 2-2. Functionally Equivalent JES2 Components 2-11
- 2-3. New, Changed, or Deleted PARMLIB Members 2-18
- 2-4. Stand-Alone Dump Macro Instruction Changes 2-27
- 3-1. Unauthorized Macro Instructions Mentioned in This Publication 3-4
- 3-2. Authorized Macro Instructions Mentioned in This Publication 3-6
- 3-3. Example of Using BSM and BASSM 3-32
- 3-4. Example of a Linkage Assist Routine 3-34
- 3-5. Retrieving Data from Above 16 Mb 3-35
- 3-6. Summary of New and Changed Macros 3-38
- 4-1. Default Hot I/O Recovery Actions 4-4
- 4-2. Summary of New, Changed, or Deleted Commands 4-7
- 6-1. New, Changed, or Deleted Dump Options 6-3
- 6-2. New, Changed, or Deleted Print Dump Verbs 6-11
- 6-3. New and Changed IPCS Subcommands 6-14
- 7-1. SMF Record Changes 7-4
- 8-1. Processing Load Modules 8-5
- B-1. Control Block Changes B-1

Summary of Amendments

Summary of Amendments for GC28-1143-2 MVS/Extended Architecture

This major revision adds conversion information for Releases 1.1 and 1.2 of MVS/SP Version 2 and MVS/XA DFP. It also includes new Release 1.0 information and several editorial changes. Bars (|) in the left-hand margin highlight the new information. Editorial changes are not barred.

The new and changed topics are, by chapter:

Chapter 2, "Installation and Initialization"

- "Installing an MVS/XA System"
- "Performing a Full Sysgen"
- "SMP/E APPLY Processing"
- "Creating a New IOCDs"
- "New I/O Configuration Requirements"
- "Devices Not Supported"
- "Changes to the Program Properties Table (PPT)"
- "Defining System Data Sets"
- "Installing the JES2 Component of MVS/SP - JES2 Version 2"
- "System Parameter and SYS1.PARMLIB Considerations"
- "SYS1.PROCLIB Changes"
- "Using Default RNLs"
- "Stand-Alone Dump Macro Instruction Changes"

Chapter 3, "Programming Considerations"

- "Macro Instructions Mentioned in This Publication"
- "Deleted Instructions"
- "Macro Expansions in JES Modifications"
- "Limiting Concurrent Global Resource Serialization Requests"
- "Format Changes to Hardcopy Log Records"
- "Link Editing Allocation User Routines"
- "Removal of the Interval Timer"
- "Checklist for Determining if Authorized Programs Must be Changed"
- "Determining Which Locks a Processor Holds"
- "Interfaces to System Services"
- "Changed Instructions"
- "New Instructions"
- "Performing I/O in 31-bit Addressing Mode"
- "Entry Points in IEFW21SD"
- "New Instructions"
- "Summary of New and Changed Macros"

Chapter 4, "Operating Considerations"

- "JCL Changes to Jobs that Allocate SYS1.DUMP Data Sets"
- "Processing Hot I/O Interrupts"
- "Controlling Message Traffic on Operator Consoles"

“New Response to Message IOS201E”
“Summary of New, Changed, or Deleted Commands”

Chapter 5, “System Modifications”

“Print Dump Exit Control Table (ECT) Modifications”
“Updating SYSTEMS Exclusion RNLs”
“Serializing VSAM Data Sets”
“Obtaining an Extended Region Size of More Than 32 Mb”
“Bypassing the Storage Availability Check Before a Job Executes”
“Changing the Hot I/O Threshold and Recovery Actions”
“JES2 User Exits”
“PRDMP Header Exits”
“SMF Exits”
“New WTO/WTOR User Exits”
“New Services for Dump Processing Exits”

Chapter 6, “Problem Determination”

“Dump Analysis and Elimination (DAE)”
“New and Changed PRDMP Control Statements”
“Print Dump Requirements for Printers”
“New and Changed IPCS Subcommands”
“Accessing Additional Sources of Dump Data Using IPCS”
“New IPCS Panels”
“Changes to the IPCS BROWSE Panels”
“Changes to the Titles of IPCS Print Files”
“Using the MVS/XA Versions of IPCS and PRDMP on Other Systems”
“Exceeding the Region Limit”
“Diagnosing Checkpoint/Restart Errors”

Chapter 7, “Accounting”

“Increases in EXCP Counts for Program Fetch Activity”
“Summary of SMF Record Changes”
“SMF Compatibility Between Release 1.0 and Later Releases”

Chapter 8, “Measurement and Tuning”

“Ensuring Optimal Program Fetch Performance”
“Using a New Directory for LNKLST Data Sets”
“SMF Data Set Placement”
“Using the ASM Backing Slot Function”
“Using Residency Time to Calculate the Page-in Rate of an Address Space”
“Changes to ASM’s Paging Algorithms”

Chapter 9, “Coexistence Considerations” (previously Chapter 10)

“Handling Downward Incompatible Macros”
“Downward Incompatible SYNCH Macros”
“Using SYS1.PROCLIB in a Loosely-coupled JES3 Configuration”

Appendix A, “Parameter Changes in Incompatible Macros”

“SYNCH Parameter List Changes”

Appendix B, "Control Block Changes"

This chapter now includes control block changes for all releases of MVS/XA.

Two chapters are deleted: "Chapter 10: Incompatibilities" and "Chapter 11: Optional Program Products." The information from Chapter 10 is incorporated elsewhere in the book, as is the information about BTAM/SP (5665-279), RMF Version 3 (5665-274), and TSO Extensions for MVS/XA (5665-285). For additional information about optional program products, see announcement letters and general information manuals for the program products.

Summary of Amendments for GC28-1143-1 MVS/Extended Architecture

This major revision incorporates additional conversion considerations related to MVS/SP Version 2 Release 1.0, MVS/XA DFP Release 1.0, CICS/VS Version 1 Releases 5 and 6, and IMS/VS Version 1 Releases 2 and 3. Most of the information was gathered from early installation experience.

The following topics, listed by chapter, contain new information:

Chapter 2: System Generation and Initialization

- "Loading the IPL Programs"
- "Loading the Microcode EC Tapes for Mass Storage Subsystems"
- "Specifying the RSU Parameter (IEASYSxx)"
- "Duration of the RMF Initialization Process"
- "Using the MVS/XA Versions of IPCS and PRDMP on Other Systems"

Chapter 3: Programming Considerations

- "Differences in GETMAIN Processing"
- "Macro Expansions in JES Modifications"

Chapter 5: System Modifications

- "PRDMP Exits"

Chapter 8: Measurement and Tuning (new chapter)

Chapter 11. Optional Program Products

- "CICS/VS Version 1 Releases 5 and 6 (5740-XX1)"
- "IMS/VS Version 1 Releases 2 and 3 (5740-XX2)"

The following topics, listed by chapter, contain technical changes:

Chapter 2: System Generation and Initialization

- "Installing an MVS/XA System" -- The topic includes SMP/E as an alternative to using SMP Release 4 to generate MVS/XA.
- "Creating a New IOCDS" -- The last paragraph is new.

“Defining System Data Sets” -- The information about defining the SYS1.NUCLEUS data set is new.

“Fixed Storage for SLIP Command Processors (IEACMD00)”

“RMF Procedure”

Chapter 3: Programming Considerations

“Macro Instructions Mentioned in This Publication” -- SYNCH is now in the unauthorized rather than the authorized macro table.

“IOGEN UCBLook Macro Instruction” -- The topic includes using IOSVSUCB to obtain UCB addresses in MVS/XA.

“Interfaces to System Services” -- The lists of examples is changed.

“Summary of New and Changed Macros” -- A description of RACROUTE is included in the table.

“New Parameters on the GETMAIN Macro Instruction” -- The topic describes two additional GETMAIN parameters, VRC and VRU.

Chapter 4. Operating Considerations

“Processing Hot I/O Interrupts” -- The default recovery action for recursive hot I/O conditions on reserved DASD is changed.

“CONFIG Command” -- The sample response from a CONFIG ONLINE command is changed.

Chapter 5. System Modifications

“JES2 User Exits” -- The topic now describes how to obtain the correct level of downward incompatible macros in JES2 user exits.

“JES3 Dynamic Support Programs (DSPs) and User Exits” -- The topic now describes how to obtain the correct level of downward incompatible macros in JES3 user exits.

Chapter 6. Problem Determination

“Print Dump Index” -- The topic describes changes to two additional PRDMP verbs, ASMDATA and DISPLAY.

Chapter 7. Accounting

The entire chapter is rewritten.

Chapter 9. Incompatibilities

“Features Not Supported” and “Programs and Functions Not Supported” are deleted. (That information is included in the Licensed Programming Announcement letter titled “Programs Supported in an MVS/Extended Architecture Environment.”)

“Devices Not Supported” -- The list includes additional devices.

Chapter 10. Coexistence Considerations

“Assembling and Link Editing Programs”

Appendix B. Control Block Changes

Chapter 1. Introduction

Conversion to MVS/Extended Architecture (MVS/XA) is the process of installing the program products that will comprise your MVS/XA system, making any required changes to existing programs and procedures, and running and testing the new system as the production system. At a minimum, you must install MVS/System Product (MVS/SP) Version 2 and MVS/XA Data Facility Product (DFP). Beyond that, the work required to convert to MVS/XA varies greatly from one installation to another and depends on:

- The level of the MVS/370 system to be converted. The more your current system resembles your target system, the less work you have to do at the same time you install the MVS/XA components. The next topic describes several ways you can prepare your MVS/370 system for conversion to MVS/XA.
- The number of programs that must be modified. Early installers reported that none of their high-level (assembler) language programs had to be changed. About fifteen percent of their authorized assembler language programs required modification.

With few exceptions, user-written assembler language programs that use only unauthorized services and published external interfaces will run unchanged. Many programs that use authorized services or undocumented interfaces will also work unchanged, but you might have to modify some. Specifically, you need to modify programs that depend on the structure and content of system control blocks or interfaces that are changed. The changed interfaces are almost exclusively authorized, internal interfaces.

- The number and type of modifications your installation has made to MVS/370 that must be adapted to MVS/XA, and which components your installation has modified. Some components are changed more than others.

In general, there is a high degree of compatibility between MVS/370 and MVS/XA:

- Exit interfaces, in general, are unchanged or compatibly expanded.
- You do not have to recompile or relink edit existing programs, unless, of course, you change them.
- JCL and JES control statements are not changed. In some instances, however, you might have to change JCL specifications, including:
 - DD statements for SYS1.DUMPnn data sets. The DD statements must specify `DISP=SHR`.
 - The `REGION` parameter on a linkage editor job.
 - JCL that specifies programs not supported in MVS/XA (for example, `IEHDASDR`).
 - JCL that specifies unsupported devices.

- MVS/XA uses the same system data sets as MVS/370. Some changes have been made to SYS1.PARMLIB members.
- Most operator procedures and commands remain the same, although some exceptions are noted in Chapter 4.

Conversion Tasks You Can Do Before Installing MVS/XA

You can stage the conversion to MVS/XA by performing many of the conversion tasks on your MVS/370 system before installing the MVS/XA components. Moving in the direction of MVS/XA as early as possible has several advantages. The most obvious is that it minimizes the activities you must perform at the same time you install MVS/SP Version 2 and MVS/XA DFP. In addition, you become familiar with the new environment gradually and have less to learn all at once. Finally, the MVS/370 system will be in the best position for coexistence and back-up. MVS/370 and MVS/XA can operate and share data in the same system complex.

To prepare for MVS/XA, you can:

- Upgrade your system to at least MVS/SP - JES3 Version 1 Release 3.1 or MVS/SP - JES2 Version 1 Release 3.0 with PTFs installed. If your MVS/370 system is at one of these levels, you do not have to install the JES component of MVS/SP Version 2.

The JES3 component at that level is functionally equivalent to the JES3 component in MVS/SP Version 2 (all releases). The JES2 component is functionally equivalent to the JES2 in MVS/SP Version 2 Release 1.0. Releases 1.1 and 1.2 can run with that JES2 as well, although both include enhanced JES2 components that are functionally equivalent to the JES2 in later releases of MVS/SP Version 1. (See “Installing the JES2 Component of MVS/SP - JES2 Version 2” in Chapter 2.)

- Install the MVS/XA-compatible levels of other program products that your installation needs in MVS/XA. Ideally, you will then have to change only the base control program (BCP) of MVS at the time you install MVS/SP Version 2. The Licensed Program Announcement letter, “Programming Support in an MVS/XA Environment,” lists the program products, IBM Field Developed Programs (FDPs), and Installed User Programs (IUPs) that can be installed on and will support both MVS/370 and MVS/XA.

When initially installing an MVS/XA-compatible program product on MVS/370, check the RETAIN Preventive Service Planning (PSP) bucket for that product. Some products might require PTFs to ensure compatibility with MVS/XA. Check RETAIN again just before testing the product under MVS/XA.

- Install products whose functions are included in MVS/XA. Such products include Data Facility Device Support, Data Facility Extended Function, and SAM-E, as well as MVS/370 Data Facility Product (MVS/370 DFP), which includes each of the previous three products.
- Review the devices and functions that are not supported in MVS/XA. If you are currently using any of them, migrate to the successor product or function.

- Find out which of the FDPs, IUPs, and other program products you have installed will work in MVS/XA. (Licensed Programming Announcement letters contain that information.) If a product your installation uses is not supported, migrate to a successor product.
- Check the RPQ devices or features you have on your system to determine if they will work in MVS/XA.
- Install compatibility PTFs on your MVS/370 system and reassemble the affected programs. See the following topics in Chapter 3:
 - “IOHALT Macro Instruction (SVC 33)” on page 3-7
 - “IOGEN UCBLINK Macro Instruction” on page 3-7
- Identify and make required programming changes that can be made on your MVS/370 system.

Publications Changes

Most MVS/XA publications are technically updated versions of their MVS/370 counterparts, reissued with new order numbers. Most title pages of MVS/SP Version 2 and MVS/XA DFP publications include “MVS/Extended Architecture” to allow you to easily distinguish between MVS/370 and MVS/XA publications. The following topics describe specific differences between the MVS/370 and MVS/XA libraries.

New Books

The MVS/XA Library includes two new books: *SPL: 31-bit Addressing* and *SPL: User Exits*.

SPL: 31-bit Addressing describes system changes that support 31-bit addressing, and how to change existing programs or to develop new programs to execute in 31-bit addressing mode.

SPL: User Exits contains reference information for coding exits provided by the base control program (BCP) of MVS/SP Version 2, with the exception of SMF exits (which are still documented in *SPL: System Management Facilities*) and TSO exits (which remain in *SPL: TSO*). In some cases, consolidating this information involved removing it from other books in the library.

Changes to *SPL: Supervisor* and *SPL: Job Management*

SPL: Supervisor and *SPL: Job Management* have been redesigned.

SPL: Supervisor has been retitled *SPL: System Macros and Facilities* and documents macros and facilities that can be used in any installation-written extension to the system, regardless of the part of the system being extended.

SPL: Job Management is now titled *SPL: System Modifications* and documents ways an installation can use IBM-provided interfaces other than initialization parameters to modify the BCP (non-JES) part of MVS/SP Version 2. *SPL: System Modifications* contains planning information on using exit routines (such as the new dumping services exits); *SPL: User Exits* contains reference information for coding the exits. *SPL: System Modifications* also describes facilities (such as virtual fetch) that are related to subsystems.

Changes to the content of these books are not as radical as the changed titles might imply. Examples of changes include:

- Information on dynamic allocation has been moved to *SPL: System Macros and Facilities*.
- Information on intercepting hot I/O has been moved to *SPL: System Modifications*.

Changes to Diagnostic Techniques and Program Logic Manuals

Section 5, "Component Analysis," and Appendix A, "Process Flows," have been deleted from the MVS/XA version of *Diagnostic Techniques* and incorporated into the component introductions of program logic manuals (PLMs, specifically *System Logic Library* and *JES2 Logic*).

In addition to information formerly in *Diagnostic Techniques*, *System Logic Library* includes logic documentation for global resource serialization and the I/O supervisor. These components are documented in separate PLMs in the MVS/370 library.

MVS/XA DFP Information

Some of the MVS/XA DFP books consolidate information from more than one book in the previous library, as follows:

Previous Titles	MVS/XA Title
<i>DFEF: Administration and Services</i>	<i>Data Facility Product: Planning Guide</i>
<i>OS/VS Planning for Enhanced VSAM</i>	
<i>OS/VS VSAM Programmer's Guide</i>	
<i>OS/VS VSAM Programmer's Guide</i>	<i>VSAM Administration: Macro Instruction Reference</i>
<i>OS/VS VSAM Options for Advanced Applications</i>	
<i>DFEF: Administration and Services</i>	
<i>OS/VS Planning for Enhanced VSAM</i>	<i>VSAM Administration Guide</i>
<i>OS/VS VSAM Options for Advanced Applications</i>	
<i>OS/VS VSAM Programmer's Guide</i>	
<i>OS/VS AMS Cryptographic Option</i>	

*OS/VS2 CVOL Processor
DFEF: Administration and
Services*

*OS/VS Planning for Enhanced
VSAM*

*OS/VS VSAM Options for
Advanced Applications*

*OS/VS VSAM Programmer's
Guide*

Catalog Administration Guide

Title Changes

Minor changes have been made to several titles, either to shorten titles or to achieve more consistency among titles of books with similar information. For example:

MVS/370 Title

*SPL: Initialization and
Tuning Guide*

*JES3 SPL: Installation Planning
and Tuning*

*SPL: JES2 Installation,
Initialization, and Tuning*

MVS/XA Title

*SPL: Initialization and
Tuning*

*SPL: JES3 Initialization
and Tuning*

*SPL: JES2 Initialization
and Tuning*

Chapter 2. Installation and Initialization

This chapter contains information related to installing an MVS/XA system, initializing it, and generating stand-alone dump. Topics related to installing MVS/XA are grouped under “Installing an MVS/XA System” and include:

- “Performing a Full Sysgen” on page 2-2
- “SMP/E APPLY Processing” on page 2-2
- “Required Environment for Generating MVS/XA” on page 2-2
- “Providing a Backup Copy of the Existing System” on page 2-2
- “Defining Additional Devices” on page 2-3
- “Creating a New IOCDS” on page 2-3
- “New I/O Configuration Requirements” on page 2-5
- “Coding Sysgen Macros” on page 2-5
- “Devices Not Supported” on page 2-6
- “Rebuilding Alternate Eligible Device Tables (EDTs)” on page 2-7
- “Changes to the Program Properties Table (PPT)” on page 2-7
- “Defining System Data Sets” on page 2-8
- “Initializing DASD” on page 2-10
- “Loading the IPL Programs” on page 2-10
- “Loading the Microcode EC Tapes for Mass Storage Subsystems” on page 2-10
- “Installing the JES2 Component of MVS/SP - JES2 Version 2” on page 2-11

Topics related to initializing MVS/XA are:

- “System Parameter and SYS1.PARMLIB Considerations” on page 2-12
- “SYS1.PROCLIB Changes” on page 2-23
- “Using Default RNLs” on page 2-26
- “Duration of the RMF Initialization Process” on page 2-26

Topics related to generating stand-alone dump include:

- “Generating Stand-Alone Dump” on page 2-26
- “Stand-Alone Dump Macro Instruction Changes” on page 2-27

Installing an MVS/XA System

You can install the MVS/XA products several ways. Your options depend on whether you are:

- Initially installing MVS/XA (that is, installing Release 1.0, or installing Release 1.1 or 1.2 concurrently with earlier MVS/XA releases)
- Installing Release 1.1 or 1.2 on a system with earlier MVS/XA releases installed

To initially install MVS/XA, you can either perform a full sysgen or use SMP/E APPLY processing. To install Releases 1.1 or 1.2 on top of earlier releases, use either SMP Release 4 or SMP/E. The program directories for the products separately describe each method.

Performing a Full Sysgen

If your MVS/XA system needs optional products that have sysgen support, accept those products into your DLIBs before performing the sysgen. Install optional products that do not have sysgen support and any related service after performing the sysgen. *System Generation Reference* describes how to perform a sysgen.

SMP/E APPLY Processing

SMP/E APPLY processing replaces and deletes the same base control program and DFP-related products and service as a full sysgen. ("Providing a Backup Copy of the Existing Program" lists the parts of MVS Release 3.8 that are replaced or deleted.) However, APPLY processing leaves intact all other products and related service. Therefore, you do not need to re-install optional products or modifications that do not have sysgen support.

To use SMP/E APPLY processing, SMP/E must be installed on the system you are using to create the MVS/XA system. Also, the required SMP data sets in the target system must be in SMP/E format. The *SMP/E User's Guide* describes SMP/E APPLY processing in detail.

Required Environment for Generating MVS/XA

You can generate an MVS/XA system on either:

- An MVS/370 system that is at least at the OS/VS2 Release 3.8 level. The MVS/370 system must also support the device types on which the MVS/XA system libraries are to reside.
- An MVS/XA system.

The following program products must be installed on the system used to build the MVS/XA system:

- Assembler H Version 2.
- The linkage editor in MVS/XA DFP.
- SMP Release 4 or SMP/E.
- Device Support Facilities Release 6, which is required to write the IPL text and to initialize the volumes on which the new system will reside.

In addition, DFDSS (Data Facility Data Set Services) or an equivalent dump/restore product is recommended to make a backup copy of the new system. IEHDASDR does not work in MVS/XA. Furthermore, DFDSS cannot restore data dumped using IEHDASDR.

DFDSS 1.2 runs on both MVS/370 and MVS/XA. Using DFDSS you can dump data on one system and restore it on the other.

Providing a Backup Copy of the Existing System

Before using the SMP ACCEPT function to incorporate the MVS/XA products into your DLIBs, copy the DLIBs using DFDSS or an equivalent product. MVS/SP Version 2 completely replaces (and, therefore, deletes from the existing

DLIBs), the base control program (BCP) in MVS Release 3.8. MVS/XA DFP completely replaces all MVS/370 modules containing the functions that MVS/XA DFP provides. The ACCEPT function also deletes:

- System Activity Measurement Facility (MF/1)
- Display Exception Monitor Facility (DEMF)
- The External Writer
- TSO TEST
- The TSO command package. The functions in that package are:
 - Support for running terminal sessions as batch jobs
 - Automatic saving of data
 - Accounting facilities enhancements
 - Defaults for the user attribute data set
 - ATTRIB and FREE subcommands
 - ALL keyword for the FREE command and subcommand
 - Eight-character station ID
- TSO/E for MVS/370, which includes the TSO command package
- All service information for the deleted modules

MF/1 and DEMF are not replaced. TSO TEST, the TSO command package, and the TSO/E functions are included in TSO/E for MVS/XA (5665-285). If your installation requires these functions, install TSO/E for MVS/XA. The External Writer function is incorporated into the MVS/SP Version 2 BCP. The program directories for MVS/XA DFP and MVS/SP Version 2 list the FMIDs that MVS/XA replaces.

Once the DLIBs are updated, there is no simple way to restore them to the MVS/370 level unless you have a backup copy. The SMP RESTORE function cannot restore DLIBs.

In addition, if you are updating your existing sysgen and IOCP deck to use for the MVS/XA sysgen, first copy the deck. If your installation intends to use the existing master catalog in MVS/XA, provide backup for its contents also.

Defining Additional Devices

MVS/XA supports a maximum of 4096 devices. However, the number that your installation can actually connect depends on the processor model. The number is usually a few less than 4096.

MVS/370 allows no more than 1917 devices because UCB pointers are only 2-bytes long. MVS/XA removes that limitation by using 3-byte UCB pointers.

Creating a New IOCDS

You must create a new IOCDS for MVS/XA. To create a new IOCDS, execute the 370/370-XA version of IOCP (either the stand-alone IOCP or the MVS IOCP that is shipped in MVS/SP Version 2). The 370/370-XA MVS IOCP executes on an MVS/370 or an MVS/XA system. If creating the new IOCDS on an MVS/370 system, obtain the 370/370-XA MVS IOCP by copying it from the MVS/XA DLIBs after accepting MVS/SP Version 2. IOCP is a set of fifteen CSECTS in the SYS1.AOSC5 data set of the DLIB. The sysgen macro SGICP400 contains the linkage editor control statements required to link edit IOCP. You must supply the JCL.

Although you can run IOCP before or after sysgen is completed, running it before sysgen is preferable. IOCP performs many validity checks on the input deck. If any input statements contain errors, rerunning IOCP is quicker than rerunning sysgen.

As long as you provide 370-dependent information (such as channel numbers), the 370/370-XA IOCP creates an IOCDS that can be used in either 370 or 370-XA mode. The 370 level of IOCP creates an IOCDS that can be used only in 370 mode. Also, you cannot use the 370 level of IOCP to read or print reports from a 370/370-XA IOCDS. If you attempt to, the system issues message ICP404I, which indicates that the level of the IOCDS directory is invalid.

During the migration period, it is important that you create an IOCDS that can be used in both 370 and 370-XA modes. Therefore, you must run the 370/370-XA IOCP to create the IOCDS. In addition, to ensure that the IOCDS works for both 370 and 370-XA, you can use the same IOCP macro specifications to create it as used to create a 370 IOCDS. Although the macro instructions are upward compatible between 370 and 370-XA modes, be aware of differences in the way IOCP treats the macro specifications when defining a 370-XA I/O configuration. Most of the differences support the new I/O architecture:

Macro MVS/XA Differences

CHPID IOCP requires channel numbers and channel sets only for devices to be used in 370 mode. The 370-XA architecture does not use channel numbers and channel sets.

IODEVICE **ADDRESS keyword.** IOCP treats the ADDRESS keyword value as a device address in 370 mode and as a device number in 370-XA mode. During the conversion period, specify the device numbers exactly the way you specify 370 device addresses.

From a users point of view, MVS/XA device numbers are equivalent to MVS/370 device addresses (sometimes referred to as CUAs in MVS/370). Both uniquely identify a device. In some publications and messages, you might still see device numbers referred to as device addresses.

UNITADD keyword. UNITADD is a new optional keyword that specifies the two-digit physical unit address of the device being described. UNITADD provides an alternative to specifying the unit address on the ADDRESS keyword (the last two digits of ADDRESS=xxx). If you specify UNITADD, the last two digits of ADDRESS=xxx need not be the device's actual physical unit address, as previously required. Instead, they can be any value that: (a) makes the device number unique, and (b) follows the rules listed in the *IOCP User's Guide and Reference*.

UNITADD allows you to assign the same unit address to more than sixteen devices. Without UNITADD, the limit is sixteen because the first digit of ADDRESS=xxx must be 1-F, the last two digits must be the device's unit address, and the three-digit combination must be unique. The first two restrictions allow only sixteen unique combinations (for example, 1FF-FFF for devices having unit address FF).

You cannot use UNITADD on IODEVICE macros used to generate an MVS/370 system. The MVS/370 sysgen program does not recognize UNITADD on IODEVICE and fails.

PATH keyword. PATH is a new optional keyword that has meaning only in 370-XA mode. It specifies a preferred path, which the channel subsystem tries first when initiating I/O to the device. You cannot include PATH on IODEVICE macros used to generate an MVS/370 system. As with UNITADD, the MVS/370 sysgen program does not recognize the PATH keyword on IODEVICE and fails.

WARNING: When coding the CNTLUNIT macro, remember to specify on the UNITADD parameter all unit addresses that the control unit can address, regardless of whether a device is actually attached. This rule is not new. However,

not specifying all unit addresses can degrade performance more in MVS/XA. If all unit addresses are not specified, channel processing is less efficient.

New I/O Configuration Requirements

Before running the 370/370-XA IOCP, review your current I/O configuration to ensure that it is valid in 370-XA mode, **even if you plan to run only in 370 mode**. The 370/370-XA IOCP imposes new restrictions on the I/O configuration (for example, no physical control unit can share devices with more than three other physical control units). Current I/O configurations might violate the new restrictions, in which case the 370/370-XA IOCP issues error messages.

The new restrictions are related to logical control units. A logical control unit represents a group of physical control units that physically and logically attach I/O devices in common. The 370-XA channel subsystem uses logical control units when queuing I/O requests and establishing path selection orders. The *IOCP User's Guide and Reference* describes logical control units in more detail and lists the new I/O configuration requirements they must satisfy.

Coding Sysgen Macros

Incompatibilities you need to check for when coding sysgen macros are:

- Macros that specify unsupported device types. (See "Devices Not Supported.") The macros you need to check include **CONSOLE**, **DATASET**, **GENERATE**, **IODEVICE**, and **SCHEDULR**. If a specified device is unsupported (for example, **IODEVICE UNIT=2314**), sysgen processing identifies the invalid device, issues an error message indicating the quit switch has been set, and does not produce a Stage II jobstream.
- Invalid specifications on the lock (L) parameters of **SVCTABLE** macros. MVS/XA changes to the locking structure might affect which locks the system must obtain before calling the SVC routine. Sysgen processing, however, does not identify invalid lock requests. See "Changes to the Locking Structure" for more detail.

Figure 2-1 summarizes the changes to sysgen macros. Most of the changes are compatible. MVS/XA sysgen processing generally ignores MVS/370 macros, keywords, and options that have no meaning when generating an MVS/XA system. In a few cases, it accepts them and issues an informational or warning message.

Obsolete Sysgen Macro, Keywords, and Options

Macro	Description of Changes
CHANNEL	Obsolete. If sysgen processing encounters a CHANNEL macro, it issues an informational message and continues.
CTRLPROG	The following keywords and options are obsolete: <ul style="list-style-type: none"> - ACRCODE - STORAGE - WARN - CRH
DATASET	A new DUMPDSN keyword specifies the suffixes or ranges of suffixes for dump data set names. You can define up to 100 dump data sets. Earlier releases allow a maximum of 10. You can specify secondary extents for the PARMLIB data set.
GENERATE	The default name for the INDEX keyword is SYSX, instead of SYS1.
IODEVICE	The following keywords are obsolete: <ul style="list-style-type: none"> - AP - GCU - OPTCHAN <p>The following options on the FEATURE keyword are obsolete:</p> <ul style="list-style-type: none"> - ABSLTVEC - BUFFER4K - BUFFER8K - CHARGNTR - DATACONV - DEKYB2260 - DESIGNFEAT - LIGHTPEN - LINEADDR - MDECOMPAT - NMKEYB2260 - NODESCUR - READWRITE - 2-CHANSW <p>Sysgen processing ignores all options except DATACONV, MDECOMPAT, and READWRITE. If those are specified, sysgen processing issues a warning message.</p>
SCHEDULR	TAVR=200 is obsolete.
UNITNAME	The maximum number of unique groups allowed is increased from 100 to 256. The maximum number of device numbers allowed is 4112 minus the number of unique groups. Earlier releases allow 2056. You can include a maximum of 4111 device numbers in one group.

Figure 2-1. Obsolete Sysgen Macro, Keywords, and Options

Devices Not Supported

The devices listed below are ones you might be using on your MVS/370 system but that you cannot use when running MVS/XA. Either the devices cannot be attached to a 308x processor, or MVS/XA does not support them. An asterisk (*) next to a device number indicates an RPQ has been approved for that device. To obtain RPQ approvals, descriptions, and details, see your IBM marketing representative.

Card I/O and Printers	OCR/MICR
1053 Model 4	1275
1403 Model 3	1287*
1442	1288*
1443	1419
2520	3881
2596	3886
	3895*
Consoles	Disks
1052 Model 7	2305 Model 1
2150	2311
2250 Model 1	2314
2260*	2319
3036	
3066	Transmission Controllers
3158	2702*
3210	2703*
3213	
3215	
2715	
Tapes	Other
2401	1012
2402	1066
2403	2671*
2404	2790
2420	2816
3410	2955
3411*	3540*
Cartridge Readers	3670
	7443
2495	7770*

Rebuilding Alternate Eligible Device Tables (EDTs)

EDTs are not compatible between MVS/370 and MVS/XA. Neither an MVS/XA nor an MVS/370 system can use EDTs verified on the other system. If your installation uses alternate EDTs, you must rebuild them using the EDTGEN macro in the MVS/XA SYSGEN macro library. You can build the EDTs on either an MVS/370 or an MVS/XA system, but you must verify them on an MVS/XA system.

Changes to the Program Properties Table (PPT)

Installing Release 1.1 replaces the program properties table (PPT). The updated PPT contains two new entries: one for IFASMF and one for CSVLLCRE. IFASMF is a new SMF module that is required to start the new SMF address space. CSVLLCRE creates and maintains a new directory of modules in the LNKST concatenation. (See "Using a New Directory for LNKST Data Sets" for more information.)

If your installation has added PPT entries, either include them in the new PPT table, or copy the new Release 1.1 entries into your version of the PPT. *System Modifications* describes how to update the PPT.

Note: ACF/VTAM (5665-280) also replaces the PPT table. If you install that product after Release 1.1, you need to make your PPT updates again.

Defining System Data Sets

An MVS/XA system requires data sets with the same names and characteristics as an MVS/370 system. Additional information related to defining system data sets is described below.

Device Types Allowed

Except for page data sets, you can place system data sets on all devices that were previously allowed, provided MVS/XA supports those device types. The next paragraph describes the page data set exception. You must move data sets on unsupported devices to supported devices. See "Devices Not Supported."

Page Data Sets

You can use the same page data sets as used in MVS/370 with one exception. Release 1.0 does not support page data sets that reside on a 3350 device attached to a 3880 Model 11 (UNIT=3350P on the IODEVICE macro). That support is added in Release 1.1.

You need to evaluate the number and size of page data sets you have defined. Both the system's and users' virtual storage requirements might increase enough to require additional external page space.

Swap Data Sets

You need to evaluate the number and size of swap data sets defined. As virtual storage requirements increase, you might need to define additional swap space.

Dump Data Sets

Installations can now define up to 100 SYS1.DUMPnn data sets. A maximum of 10 dump data sets are allowed in MVS/370.

Your installation might want to increase the number and size of dump data sets defined during the migration period. Allocate dump data sets large enough to contain the maximum size SVC dump expected. The size of the dump depends on the dump options. Early test experience indicates that SYS1.DUMPnn data sets used with MVS/370 are probably large enough for MVS/XA SVC dumps.

A new command, DUMPDS, allows installations to add and delete SYS1.DUMPnn data sets after IPL/NIP time. See "Summary of New, Changed, or Deleted Commands" in Chapter 4. Also, SYS1.DUMPnn data sets must be allocated DISP=SHR instead of DISP=OLD. See "JCL Changes to Jobs that Allocate SYS1.DUMP Data Sets" in Chapter 4.

SYS1.DAE Data Set

To start dump analysis and elimination (DAE), Release 1.1 requires that a new system data set, SYS1.DAE, be allocated at IPL time. DAE, a new function in Release 1.1, stores in SYS1.DAE symptom information from dumps it identifies as unique. It uses that information when determining if subsequent dumps are duplicates. "Dump Analysis and Elimination (DAE)" in Chapter 6 gives an overview of DAE and describes in more detail how and when SYS1.DAE is used.

You can create SYS1.DAE using JCL in the DAEALLOC member of SYS1.SAMPLIB. (The DATASET sysgen macro does not support SYS1.DAE.) For instructions, see *System Modifications*.

New MSTRJCLxx Members in the SYS1.LINKLIB Data Set

In Release 1.1, the JCL for starting the master scheduler address space is contained in MSTRJCLxx members of SYS1.LINKLIB. MSTRJCL, the member that earlier releases use, is deleted. IBM supplies default JCL in MSTRJCL00 instead. To change the JCL, create additional members. Use the new MSTRJCL system parameter to specify which member the system is to use. The default is MSTRJCL=00.

Concatenating Data Sets to the SYS1.LPALIB Data Set

If Release 1.1 is installed, you can concatenate data sets to the SYS1.LPALIB data set. The system uses the modules in the concatenated data sets, as well as the SYS1.LPALIB data set, to build the PLPA, the MLPA, and the FLPA. Earlier releases of MVS use only the modules in the SYS1.LPALIB. LPALIB concatenation allows you to share a single SYS1.LPALIB data set among several systems, yet still tailor the PLPA, MLPA, and FLPA of each system by varying the concatenation.

To concatenate data sets:

- List in a new LPALSTxx PARMLIB member which data sets are to be concatenated. The data sets must be included in the master catalog and must be APF authorized.
- Specify on the new LPA system parameter which LPALSTxx members are to be processed. You can include the LPA parameter in IEASYSxx, or an operator can specify it when prompted for system parameters. If you omit the LPA parameter, the system uses SYS1.LPALIB only, as in previous releases of MVS.

See *System Initialization and Tuning* for more detail on creating LPALSTxx members and specifying the LPA system parameter.

SYS1.LOGREC Data Sets

If you install Release 1.1, you can place SYS1.LOGREC data sets on a volume other than the SYSRES volume. Several systems can then share a SYSRES volume and still have separate SYS1.LOGREC data sets. To use an alternate SYS1.LOGREC data set, simply include a data set named SYS1.LOGREC in the master catalog. The system searches for a data set with that name first in the master catalog, then in the SYSRES volume.

The sysgen process requires that the SYS1.LOGREC data set reside on the SYSRES volume. Therefore, do not include a data set named SYS1.LOGREC in the master catalog until after the sysgen process is completed.

You might want to increase the size of your SYS1.LOGREC data set because MVS/XA can produce more diagnostic information.

SYS1.NUCLEUS Data Set

As in MVS/370, the MVS/XA SYS1.NUCLEUS must be a single extent. If you attempt to allocate a multiple extent data set, MVS/XA enters a restartable wait state (wait state code x'081'). MVS/370 takes different actions.

SYS1.PARMLIB Data Set

In MVS/XA, the SYS1.PARMLIB data set can be blocked and can have multiple extents. In MVS/370, the PARMLIB has to be unblocked and a single extent. Also see "New, Changed, or Deleted PARMLIB Members."

System Data Set Qualifiers

You cannot specify a system data set qualifier of SYS1 on the INDEX parameter of the GENERATE macro. You can either specify some other high level qualifier or let sysgen processing assign the default (SYSX). Sysgen Stage II processing changes the high level qualifiers to SYS1. This restriction is not new. It always applies when using a system other than a starter system to do a complete sysgen. Note that the default high level qualifier is changed from SYS1 to SYSX.

Initializing DASD

You must use Device Support Facilities to initialize DASD volumes. IEHDASDR is no longer supported. See the *Device Support Facilities User's Guide* for directions.

Loading the IPL Programs

To IPL MVS/XA, you must use the IPL text distributed with MVS/SP Version 2. The MVS/370 and MVS/XA IPL programs are not compatible. You must use Device Support Facilities Release 6 to write the IPL text to DASD.

Device Support Facilities loads the second MVS/XA bootstrap record into the frame at main storage absolute address 8 K. It loads the IPL text into main storage frame 0.

If you write your own bootstrap programs, ensure that the addresses used to load the bootstrap records are in storage that will not be taken offline. In MVS/XA, your choices are:

- Main storage frame 0, 2, 4, 8, and so on, throughout the first main storage range.
 - The low end of the highest storage range specified on the CONFIG frame on the system console. That storage range always remains online.
-

Loading the Microcode EC Tapes for Mass Storage Subsystems

To load MSS microcode in an MVS/XA environment, use the MSC Table Create (MSCTC) utility with PTF UZ09020 installed, instead of IEHDASDR. IEHDASDR does not work in MVS/XA.

The MSCTC control statement to use is CREATE, the required parameter is RESTOREC. For more information, see *MSS Installation Planning and Table Create*.

Installing the JES2 Component of MVS/SP - JES2 Version 2

If you are converting to MVS/XA from an MVS/370 system that has at least MVS/SP Version 1 Release 3.0 with PTFs installed, you do not have to install the JES2 component of MVS/SP Version 2, regardless of which release you install. The Release 3.0 JES2 with PTFs installed is functionally equivalent to the JES2 component in MVS/SP Version 2 Release 1.0. Furthermore, it will run on all releases of MVS/SP Version 2, although Releases 1.1 and 1.2 include enhanced levels of JES2.

The JES2 component of Release 1.1 is functionally equivalent to the JES2 component of MVS/SP Version 1 Release 3.3. The JES2 component of Release 1.2 is functionally equivalent to the JES2 component of MVS/SP Version 1 Release 3.4.

If your current and target systems have equivalent JES2 components, you need not reinstall JES2. To install an enhanced JES2 component, perform either a warm or a cold start, depending on the level of your current JES2 component.

Note that the JES2 components mentioned are upward, but not downward, compatible. That is, they work only on systems that have functionally equivalent or enhanced JES2 components.

The following chart summarizes the JES2 correlations between MVS/SP Versions 1 and 2. It also indicates whether you perform a cold or warm start to upgrade the level of JES2.

Level of your current system's JES2 component	Level of the target MVS/XA system		
	Release 1.0	Release 1.1	Release 1.2
MVS/SP Version 1 Release 3.0 with PTFs Release 3.1 Release 3.2	The JES2 components are functionally equivalent.	The target system can run with the current level of JES2. Upgrading JES2 requires a cold start.	The target system can run with the current level of JES2. Upgrading JES2 requires a cold start.
MVS/SP Version 1 Release 3.3 MVS/SP Version 2 Release 1.1	The current JES2 component will not work on the target system.	The JES2 components are functionally equivalent.	The target system can run with the current level of JES2. Upgrading JES2 requires a warm start.
MVS/SP Version 1 Release 3.4 MVS/SP Version 2 Release 1.2	The current JES2 component will not work on the target system.	The current JES2 component will not work on the target system.	The JES2 components are functionally equivalent.

Figure 2-2. Functionally Equivalent JES2 Components

System Parameter and SYS1.PARMLIB Considerations

The topics in this section contain information related to specifying system parameters. “New, Changed, or Deleted PARMLIB Members” on page 2-17 summarizes changes to SYS1.PARMLIB members.

Fixed Storage for SLIP Command Processors (IEACMD00)

A new PARMLIB member, IEACMD00, contains IBM-supplied SLIP commands to suppress dumps that are not usually required for problem determination. (See “Suppressing Dumps” in Chapter 6 for detail.)

When the system processes IEACMD00 at IPL time, it allocates fixed storage for the SLIP action processors and the control blocks they use. The action processors require approximately 30 K bytes of fixed storage in extended LPA. The control blocks require approximately 1 K bytes of fixed storage in extended SQA.

If your installation does not use SLIP commands in MVS/370, the fixed storage requirement is new. Because the fixed storage is allocated above the 16 Mb line, installations that previously used SLIP commands will gain approximately 31 K bytes of virtual storage below 16 Mb.

Specifying the RSU Parameter (IEASYSxx)

If you specified an RSU parameter in the past, when initializing Release 1.1, you need to review that specification. Release 1.1 satisfies the RSU request in a different way than previous releases do. The same RSU value might result in less reconfigurable storage.

In Release 1.1, the RSU parameter specifies the *total number of storage units MVS is to mark reconfigurable*. The system attempts to satisfy the request using offline storage units. It marks online storage reconfigurable only if there is not enough offline storage. After satisfying the RSU request, the system marks all remaining storage units (both online and offline) as preferred. If the system cannot satisfy the RSU request, the operator receives message IAR004I, as in previous releases.

In Release 1.0 and earlier releases, the RSU parameter specifies the *number of online storage units to be marked reconfigurable when initializing the system*. Those releases use only online storage to satisfy the request. However, they also automatically mark storage that is offline at IPL time as reconfigurable when bringing it online. Thus, the total amount of reconfigurable storage is the amount marked reconfigurable when satisfying the RSU parameter, plus the amount brought online after the IPL. As a result, the RSU parameter in earlier releases can be less than the resulting amount of reconfigurable storage.

On a 3084 processor, the RSU value needs to be equivalent to at least the amount of storage you plan to take offline before the next IPL. Some installations specify one additional storage unit to increase the probability that storage can be taken offline later. Remember that you can lose reconfigurable storage during normal processing in two ways:

- If the system runs out of preferred storage frames, it dynamically converts some reconfigurable storage to preferred storage.
- The system might not be able to reconfigure storage that contains storage errors that cannot be corrected.

Do not, however, specify more reconfigurable storage than you anticipate needing. Specifying too much can negatively affect performance.

If performing an IPL on a 3081 processor, you do not need to specify an RSU value. Storage is not reconfigurable, and the RSU default value is zero.

System Initialization and Tuning contains more detail on specifying the RSU parameter.

Increasing the Minimum SQA Allocation (IEASYSxx)

If you changed the NVTNVSQA field in module IEAVNIPO to increase the minimum SQA allocation during previous system initializations, you need to read this topic. During system initialization, if the PAGE parameter specifies a large number of page data sets or if several 2305 Model 2 page data sets are active, the system's minimum allocation for SQA and extended SQA (seven 64 K blocks) might be depleted before the system processes the SQA parameter. In MVS/370, you can solve that problem by changing the contents of the NVTNVSQA field. That method does not work in MVS/XA. You can, however, increase the minimum allocations by changing the halfwords NVSQA and/or NVESQA in module IEAIPL04. Consult microfiche for the locations of these fields. If you increase the minimum SQA and/or extended SQA allocations and you want the total SQA size to remain the same, decrease the corresponding value on the SQA parameter.

The following Release 1.2 changes cause more SQA and extended SQA to be available earlier in the initialization process. As a result, you might not need to change the minimum SQA allocation:

- The minimum SQA allocation is increased to four 64 K blocks. In earlier releases, it is three.
 - The system processes the SQA parameter earlier during system initialization.
-

Specifying the Size of CSA and SQA Above 16 Mb (IEASYSxx)

You can use the CSA and SQA parameters on the CTRLPROG macro to specify the size of CSA and SQA below, but not above, 16 Mb. MVS/XA assigns default sizes for extended CSA and extended SQA (CSA and SQA above 16 Mb). To override those defaults, use new options on the CSA and SQA system parameters.

The default sizes are:

CSA/extended CSA - 100 K/100 K
SQA/extended SQA - 256 K/256 K plus approximately 8 Mb

The CSA and SQA system parameters each have an additional option for specifying the size of extended CSA and extended SQA:

CSA = (a,b) where "a" specifies the size of CSA or SQA,
SQA = (a,b) and "b" specifies the size of extended CSA or SQA.

The CSA values indicate the number of 1 K units to be reserved. The values override the default specifications. For example:

CSA = (100,200) results in:

Reserved CSA = 100 K
Reserved extended CSA = 200 K.

The SQA values specify the number of 64 K blocks MVS/XA is to reserve in addition to the minimum amount of storage it allocates for SQA and extended SQA. The minimum amounts are 256 K for SQA and 256 K plus approximately 8 Mb for extended SQA. For example:

SQA = (3,5) results in:

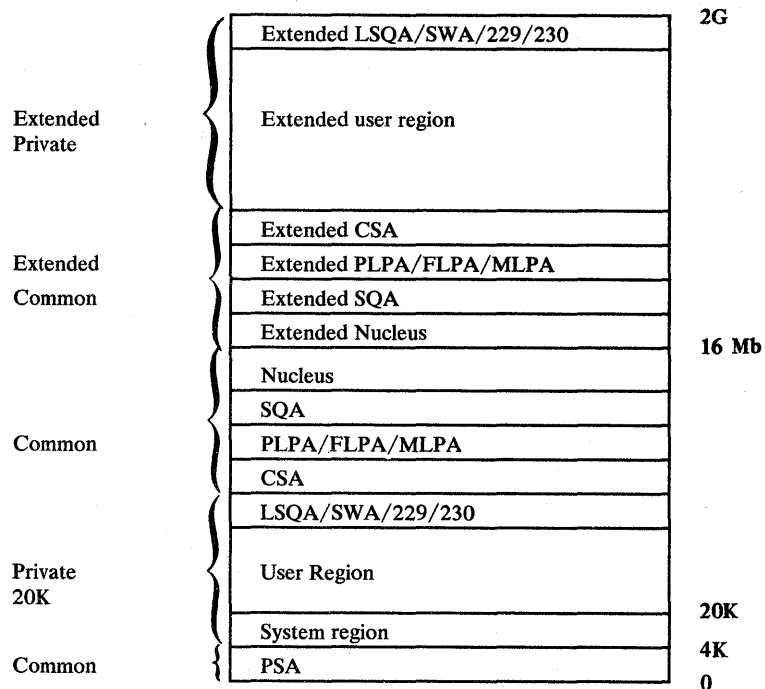
Reserved SQA = 3 x 64 K + 256 K
Reserved extended SQA = 5 x 64 K + 256 K + approximately 8 Mb

The default SQA values are SQA=(1,0).

See *Initialization and Tuning* for more information.

Minimizing Private Area Storage Lost Because of Rounding (IEASYSxx)

Because the segment size in MVS/XA increases from 64 K to 1 Mb, you need to pay closer attention to the amount of private area storage lost to CSA at initialization time because of rounding. During IPL processing, MVS/XA builds the common area below 16 Mb beginning at the 16 Mb address and working down.



To determine the lower boundary of the common area (which is the upper boundary of the private area), MVS/XA rounds the bottom CSA address to the the next lowest 1 Mb. Thus, as much as 1020 K bytes (1 Mb minus 4 K) of virtual storage can be added to the CSA, and consequently lost from the private area, because of rounding. Although it is not expected that your installation will lose

that much private area, choose your CSA and SQA parameters carefully. If the size of the private area falls below 8 Mb, MVS/XA informs the operator.

MVS/XA also builds a common area above 16 Mb. Private area storage above 16 Mb can also be lost because of rounding. However, users are not expected to have virtual storage constraint problems above 16 Mb.

Specifying Dump Data Sets (IEASYSxx)

An additional operand of the DUMP parameter in IEASYSxx, 'DASD,xx-yy', allows an installation to request that the system being initialized use a specific range of DASD dump data sets. The DASD operand was redesigned specifically for systems that use SYS1.DUMP data sets that can be accessed by other systems. An installation can specify unique dump data sets for each system, which prevents SDUMP routines from using a dump data set concurrently for different systems.

The DASD operand also shortens dump data set catalog processing at initialization time. When specific dump data sets are indicated, the initialization routines do not have to issue all 100 locates to determine which dump data sets are cataloged.

Requesting Storage for RMF I/O Measurements (IEASYSxx)

If your installation wants RMF I/O data for device classes other than tape or DASD, you must request storage at initialization time for the control blocks in which the data is to be collected. SRM collects the I/O data that RMF uses in new control blocks, one per device. To request control block storage, specify on the new CMB parameter in IEASYSxx the non-tape and non-DASD device classes for which I/O data is to be collected.

Controlling the Number of Available ASVT Entries (IEASYSxx)

A system with Release 1.2 installed creates and manages the address space vector table (ASVT) differently. The changes are designed to prevent the system from running out of ASVT entries.

When creating the ASVT, the system adds extra entries and reserves them for use when no unreserved entries are available. It uses one group of reserved entries only for address spaces being created in response to a START command. It uses a second group as replacements for entries that are not reusable because of latent cross-memory binds.

Two new system parameters allow your installation to specify the number of entries to be reserved for each purpose:

RSVSTRT Specifies the number of entries to be reserved for address spaces created in response to a START command. The default is five.

RSVNONR Specifies the number of entries to be used as replacements for entries that are not reusable. The default is also five.

The system still uses the MAXUSER parameter to limit the number of jobs and started tasks that can execute concurrently under normal conditions. However, MAXUSER no longer specifies the maximum number of jobs or started tasks the system allows. That number is usually the MAXUSER value plus the RSVSTRT value. If supervisor recovery reconstructs the ASVT, the maximum number might be the sum of the MAXUSER, RSVSTRT, and RSVNONR values. The default MAXUSER value is still 256.

Removing TRACE Commands from COMMNDxx PARMLIB Members

You might want to remove or update any TRACE operator commands in the COMMNDxx PARMLIB member. The syntax of the TRACE command is changed. MVS/370 TRACE commands do not work in MVS/XA. Also, the MVS/XA system trace remains active after system initialization time. No TRACE ON command is required to keep it active, as in MVS/370. Issuing MVS/370 TRACE commands, however, does not prevent MVS/XA system trace from being initialized or activated. "Summary of New, Changed, or Deleted Commands" in Chapter 4 describes the TRACE command changes.

Updating the IEAFIXxx PARMLIB Member

Remove from IEAFIXxx the names of modules that have been moved from LPA to the nucleus. Module names to be removed include:

- IGC0004F (the TTIMER service routine, which is renamed IGC046 in MVS/XA)
- IGC0004G (the STIMER service routine, which is renamed IGC047 in MVS/XA)
- IEWFETCH (program fetch, aliases IEWMBOSV and IEWMSEPT)
- IGC0001F (the PURGE service routine)

If those modules are in IEAFIXxx, the operator receives a message indicating that the modules could not be found. The PARMLIB member is not rejected.

Removing References to Device Allocation Tables (IEALPAxx)

The DEVNAMET, IEFDEVPT, and DEVMASKT device allocation tables are deleted in MVS/XA. Remove any references to these tables in PARMLIB members (for example, in the MLPA list).

Keeping RNLs in GRSRNLxx PARMLIB Members

If Release 1.2 is installed, you can keep global resource serialization resource name lists (RNLs) in new GRSRNLxx PARMLIB members instead of in the ISGGRNLO load module in SYS1.LINKLIB. RNLs are easier to update when kept in PARMLIB members. You can, however, continue using the RNLs in SYS1.LINKLIB.

Regardless of where the RNLs are located, if your system is to be part of a global resource serialization complex (GRS=START or GRS=JOIN), you must have at least one GRSRNLxx member. Use the new system parameter, GRSRNL=, to specify which members the system is to use.

If you keep RNLs in SYS1.LINKLIB, the GRSRNLxx member must begin with a statement that tells the system to use the RNLs in SYS1.LINKLIB (and ignore the rest of the statements in the member). That statement is RNLDEF LINKLIB(YES).

To keep RNLs in a GRSRNLxx member, you need to include in the member one statement for each RNL entry. Each statement begins with RNLDEF, specifies the

resource name, and indicates the RNL to which it belongs. *Initialization and Tuning* describes how to write RNLDEF statements.

One member, GRSRNL00, is shipped with Release 1.2. GRSRNL00 contains entries for the same resources as the default RNLs in SYS1.LINKLIB (which are also shipped in Release 1.2). In addition, it begins with the statement RNLDEF LINKLIB(YES), which causes the system to use the RNLs in SYS1.LINKLIB. IEASYS00 contains the parameter GRSRNL=00, so the system uses GRSRNL00 by default.

Because of the defaults, if using RNLs in SYS1.LINKLIB, you need not do anything. To use the RNLs in GRSRNL00, you need to:

- Remove the first statement: RNLDEF LINKLIB(YES)
- Add, delete, or modify RNLDEF statements to match your installation's resource serialization requirements.

You can also create and use other GRSRNLxx members.

Systems in the same global resource serialization complex can use different methods of defining RNLs (either statements in GRSRNLxx PARMLIB members or the ISGGRNL0 LINKLIB module). However, as before, the RNLs for all systems in the complex must be identical. The resources identified in the RNLs must be the same and they must appear in the same order.

Specifying MIH Intervals (IECIOSxx)

Installations can specify by device the time intervals at which MIH scans for missing interrupts. A new MIH statement in the IECIOSxx PARMLIB member allows installations to specify separate time intervals for:

- All DASD except 3330V devices. The default is 15 seconds.
 - 3330V devices (MSS virtual DASD). The default is 12 minutes.
 - 3851 devices (mass storage controller). The default is 12 minutes.
 - Specific devices identified by device number. There is no default. Installations can bypass MIH processing for specific devices by setting a time interval of zero.
 - All other devices. The default is 3 minutes.
-

New, Changed, or Deleted PARMLIB Members

Figure 2-3 summarizes SYS1.PARMLIB members that are new, changed, or deleted in MVS/XA. Most of the changes are compatible, some are not. For example, if the MVS/370 version of IEASYSxx specifies the ALT parameter, you cannot use it in place of the MVS/XA version of IEASYSxx. (See the entry for IEASYSxx.) In other cases, MVS/XA ignores parameters that it no longer supports and uses defaults for new parameters. If you use the MVS/370 IEAIPSxx member, you need to review the specifications to ensure optimal performance.

Initialization and Tuning describes the PARMLIB members in more detail.

Member	Release			Description of Change
	1.0	1.1	1.2	
ADYSETxx		x		<p>A new member. It contains records that control dump analysis and elimination (DAE). Each record can specify:</p> <ul style="list-style-type: none"> - Whether or not DAE is to be active - The functions DAE is to perform - The number of symptom records to be kept in the SYS1.DAE data set <p>See "Dump Analysis and Elimination (DAE)" in Chapter 6 for more information about DAE processing.</p>
CONFIGxx	x			<p>New and changed parameters:</p> <ul style="list-style-type: none"> - CHAN and CHANNEL are deleted. MVS/XA does not use channel set information. If CHAN and CHANNEL are specified, the operator receives a message indicating that they are ignored. MVS/XA continues processing CONFIGxx. - CHP specifies the configuration of channel paths. It replaces CHAN and CHANNEL. - CPU specifies a processor and can be used in place of CPUAD. You can continue to specify the CPUAD parameter, however. - The syntax of the DEV parameter is changed. It specifies channel path identifiers instead of channel set IDs. <p>The operator can use the CONFIGxx member when reconfiguring the system. The new CONFIG command has an operand, MEMBER, which specifies a CONFIGxx member. In response to a CONFIG MEMBER command, the system logically and physically reconfigures processors, storage, and channel paths as defined by the CPU, STOR, and CHP parameters in the specified CONFIGxx member.</p>
GRSRNLxx			x	<p>A new member that contains either global resource serialization resource name lists (RNLs), or a statement indicating the system is to use the RNLs in SYS1.LINKLIB. (In Release 1.2, you can keep RNLs in GRSRNLxx members or in SYS1.LINKLIB.) IBM provides one default member, GRSRNL00. See "Keeping RNLs in GRSRNLxx PARMLIB Members" earlier in this chapter for more information.</p>
GTFPARM	x			<p>Contains new options for requesting I/O event recording. USRP is also a new option, which prompts for specific USR events to be recorded.</p>

Figure 2-3 (Part 1 of 6). New, Changed, or Deleted PARMLIB Members

Member	Release			Description of Change
	1.0	1.1	1.2	
IEAABD00 (SYSABEND)	x			<p>New and changed options on the SDATA keyword:</p> <ul style="list-style-type: none"> - ALLSDATA includes all of the SDATA options except ALLVNUC and NOSYM. - ALLVNUC requests the PSA, CVT, and DAT-on nucleus. - NOSYM requests that symptom dumps be suppressed. Unless NOSYM is specified, the system produces a symptom dump each time a task abends, even if the user did not specify a dump DD statement. - NUC requests only the DAT-on, non-page protected section of the nucleus and the PSA and CVT. In MVS/370, NUC causes the entire nucleus to be dumped. - SUM requests a summary dump. - TRT requests trace data from the active trace facilities, as in MVS/370. However, only an MVS/XA dump can include both system trace and GTF data. Both trace facilities can be active at the same time in MVS/XA, but not in MVS/370. <p>SUBTASKS, a new option on the PDATA keyword, specifies that the requested PDATA information be dumped for all subtasks of the abending task. When a task receives ABEND code 'x22', the system dumps SUBTASKS data, regardless of whether the SUBTASKS option is specified.</p> <p>The PDATA default options specified in the IBM-supplied member are changed. The MVS/XA default options are all of the PDATA options except SUBTASKS. In MVS/370, ALLPDATA is the default.</p> <p>The IBM-supplied member includes an additional SDATA default option, SUM. It also includes the options specified in the MVS/370 member (LSQA, CB, ENQ, TRT, ERR, DM, and IO).</p>

Figure 2-3 (Part 2 of 6). New, Changed, or Deleted PARMLIB Members

Member	Release			Description of Change
	1.0	1.1	1.2	
IEABLDxx	x			Systems with Release 1.0 installed process any IEABLDxx member you supply, but provide no default members. Before using existing BLDL lists, ensure their accuracy. Some system modules are moved to different libraries in MVS/XA.
		x		Systems with Release 1.1 installed do not process IEABLDxx members. The LNKST lookaside (LLA) function in Release 1.1 provides a directory of modules in the LNKST concatenation. The new directory eliminates the need for the BLDL table and, thus, the need for IEABLDxx members. The system ignores the BLDL and BLDLF system parameters. For more information about the LLA function, see "Using a New Directory for LNKST Data Sets" in Chapter 8.
IEACMD00	x			A new member that contains IBM-supplied commands. Except for one CHNGDUMP command, all of the commands in the Release 1.0 member are SLIP commands. The CHNGDUMP command adds trace table and LSQA information to SVC dumps. The SLIP commands suppress dumps that are normally not required for problem determination. See "Suppressing Dumps" in Chapter 6.
		x		Release 1.1 adds two new commands: <ul style="list-style-type: none"> - SET DAE=00, which causes the system to process the ADYSET00 PARMLIB member. ADYSET00 starts DAE processing. For more information, see "Dump Analysis and Elimination (DAE)" in Chapter 6. - START LLA, which starts the LLA procedure in SYS1.PROCLIB. The LLA procedure in turn starts the LNKST lookaside (LLA) function. See "Using a New Directory for LNKST Data Sets" for a description of the LLA function.
IEADMP00 (SYSUDUMP)	x			The new and changed parameters for IEADMP00 are the same as those for IEAABD00. See the IEAABD00 entry in this table. The only default option specified in the IBM-supplied member is SUM. In most cases, the summary dump will be sufficient to debug user program checks and ABEND dumps.
IEADMR00 (SYSMDUMP)	x			New and changed SDATA options are: <ul style="list-style-type: none"> - ALLNUC, which requests that the entire nucleus be included in the dump. <p>Note: You can also request that the entire nucleus be included in SYSMDUMPs via the SNAP parameter list and the CHNGDUMP command. The required SNAP parameter list option is ALLVNUC, not ALLNUC. The CHNGDUMP option is ALLNUC.</p> <ul style="list-style-type: none"> - ALLSDATA includes all of the SDATA options except ALLNUC and NOSYM. - NOSYM, NUC, and TRT request the same data as when specified in IEAABD00. See the IEAABD00 entry in this table. - SUM requests summary dump information like that included in SVC dumps. The default options specified in the IBM-supplied member include SUM, in addition to the SDATA options specified in the MVS/370 default member (NUC, SQA, LSQA, SWA, TRT, and RGN).
IEAFIXxx	x			Unless the NOPROT option is specified on the FIX parameter in the IEASYSxx member, the system page-protects the modules listed in IEAFIXxx. See "Page Protection" in Chapter 3.
IEAIPSxx	x			A new parameter, IOSRVC, specifies whether SRM is to base I/O service on EXCP counts or device connect time. The default is EXCP counts.
			x	A new parameter, PPGRTR, requests that SRM use residency time instead of execution time when it calculates the page-in rate for address spaces in the specified performance group. PPGRTR specifies the high or low limit the rate must exceed before SRM adjusts the address space's working set size. See "Using Residency Time to Calculate the Page-in Rate of an Address Space" in Chapter 8 for more detail.

Figure 2-3 (Part 3 of 6). New, Changed, or Deleted PARMLIB Members

Member	Release			Description of Change
	1.0	1.1	1.2	
IEALPAxx	x			Unless the new NOPROT option is specified on the MLPA parameter in the IEASYSxx member, the system page-protects the modules listed in IEALPAxx. See "Page Protection."
IEALOD00	x			If IEALOD00 is specified, MVS/XA ignores it. Unlike MVS/370, MVS/XA does not build contents directory entries (CDEs) for PLPA modules. It uses the LPA directory entries (LPDEs) instead.
IEAOPTxx	x			New parameters: <ul style="list-style-type: none"> - CPENABLE specifies upper and lower thresholds for the percentage of I/O interrupts that occur while the channel subsystem is processing another I/O interrupt. SRM uses the thresholds to control I/O interrupt processing. - ICCLPB(TAPE), ICCLPB(NDPSDASD), and ICCLPB(DPSDASD) specify logical path utilization thresholds for tape, DASD without the dynamic path selection feature, and DASD with the dynamic path selection feature, respectively. SRM uses the thresholds for I/O load balancing and for non-specific device allocation.
IEAPAKxx	x	x		MVS/SP Version 2 does not provide PAK lists. However, it processes IEAPAKxx members that you supply. Release 1.0 and earlier releases recognize only IEAPAK00. Release 1.1 and later releases allow multiple IEAPAKxx members. By using more than one member, you can vary the LPALST concatenation from system to system or from IPL to IPL without changing IEAPAK00. The PAK=xx system parameter specifies which members the system is to use. IEAPAK00 is the default, which is consistent with earlier releases. Do not use existing PAK lists without first reevaluating their usefulness. Unless all modules in a group have the same RMODE (that is, they all reside in virtual storage either above 16 Mb or below 16 Mb), IPL/NIP ignores the PAK list.
IEASYSxx	x			New, changed, and deleted parameters: <ul style="list-style-type: none"> - CMB specifies the I/O device classes for which measurement data is to be collected. The CMB specifications are in addition to DASD and tape device classes, for which measurement data is always collected. - The ALT parameter is no longer supported. Have operators use the SYSCTL console frame to specify an alternate nucleus. If ALT is specified, MVS/XA truncates processing and asks the operator for another member. Because some system parameters might already have been processed, have operators re-IPL and request an IEASYSxx member that does not contain the ALT parameter. - The MLPA and FIX parameters have an additional option, NOPROT. NOPROT indicates that the LPA modules listed in the IEALPAxx or IEAFIXxx PARMLIB member are not to be page-protected. Unless the NOPROT option is specified, the system page-protects those modules. See "Page Protection" in Chapter 3. - A new option on the DUMP parameter, 'DASD,xx-yy', specifies which currently cataloged SYS1.DUMPnn data sets the system is to use. If none are specified, the system uses any that are cataloged. - CSA has an additional option that specifies the size of CSA above 16 Mb. - SQA has an additional option that specifies the size of SQA above 16 Mb.

Figure 2-3 (Part 4 of 6). New, Changed, or Deleted PARMLIB Members

Member	Release			Description of Change
	1.0	1.1	1.2	
IEASYSxx (continued)		x		<p>New parameters:</p> <ul style="list-style-type: none"> - LNKAUTH specifies whether all data sets in the LNKLST concatenation are to be treated as APF authorized or only those named in the APF table. The default is to treat all of the data sets as APF authorized, as do previous releases of MVS. See "Using a New Directory for LNKLST Data Sets" for more information. - LPA identifies the LPALSTxx member to be processed. See the LPALSTxx entry in this table for more information. - MSTRJCL specifies which MSTRJCLxx member in SYS1.LINKLIB the system is to use. If you omit the MSTRJCL parameter, the system uses the JCL in the IBM-supplied default, MSTRJCL00. <p>MSTRJCLxx members are new in Release 1.1. For more information, see "New MSTRJCLxx Members in the SYS1.LINKLIB Data Set," under "Defining System Data Sets."</p> <ul style="list-style-type: none"> - PAK identifies the IEAPAKxx member to be processed. See the IEAPAKxx entry in this table for more information. <p>The BLDL and BLDF parameters are obsolete. If specified, the operator receives warning message IEA240I. The LNKLST lookaside (LLA) function provides a directory of modules in the LNKLST concatenation. The new directory eliminates the need for an IEABLDxx member, and consequently, the BLDL and BLDF parameters. For more information, see "Using a New Directory for LNKLST Data Sets" in Chapter 8.</p>
				x

Figure 2-3 (Part 5 of 6). New, Changed, or Deleted PARMLIB Members

Member	Release			Description of Change
	1.0	1.1	1.2	
IECIOSxx	x			The LCH parameter is no longer supported. A new MIH control statement specifies intervals at which MIH is to scan for MIH conditions. MIH allows installations to specify different time intervals for different devices and/or types. See "Specifying MIH Intervals (IECIOSxx)."
		x		A new statement, HOTIO, allows you to change: - The threshold IOS uses to detect hot I/O conditions - The recovery actions IOS takes when it detects a hot I/O condition See "Processing Hot I/O Interrupts" in Chapter 4 for more information.
LNKLSTxx		x		The data sets in the LNKLST concatenation no longer have to be APF authorized. Also, you can concatenate up to 123 data sets. Earlier releases allow no more than 16. See "Using a New Directory for LNKLST Data Sets" in Chapter 8 for more information.
LPALSTxx		x		A new member that lists data sets to be concatenated to the SYS1.LPALIB data set. See "Concatenating Data Sets to the SYS1.LPALIB Data Set" under "Defining System Data Sets."
MPFLSTxx	x			Release 1.0 adds two new keywords: - .MSGCOLR specifies how messages are to be displayed on MCS color consoles that can use seven colors and other forms of highlighting. An MPFLSTxx member can also include the statement .MSGCOLR NOCHANGE, which maintains the color and highlighting attributes in effect when the member became active. - .MSGIDS NOCHANGE requests that the system not change the message IDs already in effect.
				x
SMFPRMxx		x		Release 1.1 ignores the BUFNUM parameter. The SMF buffers are moved to the new SMF address space. Consequently, the number of buffers SMF can obtain is limited only by the amount of virtual storage in the SMF address space. SMF initially obtains 100 buffers and requests more as needed. If the number of buffers in use exceeds 1000, SMF informs the operator via message IEE978E.

Figure 2-3 (Part 6 of 6). New, Changed, or Deleted PARMLIB Members

SYS1.PROCLIB Changes

The SYS1.PROCLIB data set in Release 1.0 contains a new DUMPSRV procedure and a new statement in the PRDMP procedure. The RMF procedure in RMF Version 3 is also changed. Release 1.1 adds two new procedures, LLA and IEESYSAS. Release 1.2 changes the PRDMP procedure. Either use the SYS1.PROCLIB shipped with the product, or copy the new and changed procedures into your version of SYS1.PROCLIB.

Following are descriptions of the new and changed procedures. If your MVS/XA system is part of a loosely-coupled JES3 configuration that includes MVS/370 systems, also read "Using SYS1.PROCLIB in a Loosely-coupled JES3 Configuration" in Chapter 9.

DUMPSRV Procedure

The new DUMPSRV procedure starts the dump service (DUMPSRV) address space:

```
//DUMPSRV EXEC PGM=IEAVTDSV
```

IEESYSAS Procedure

IEESYSAS is a new procedure in Release 1.1 that starts full-function system address spaces, which include the SMF address space.

```
//IEESYSAS PROC PROG=IEFBR14  
// EXEC PGM=εPROG
```

LLA Procedure

The LLA procedure starts the LNKLST lookaside (LLA) function:

```
//LLA EXEC PGM=CSVLLCRE
```

The IEACMD00 PARMLIB member shipped with Release 1.1 contains a START LLA command, which starts the LLA procedure. See "Using a New Directory for LNKLST Data Sets" in Chapter 8 for more information about the LLA function.

PRDMP Procedure

The PRDMP procedure in Release 1.2 is changed. Release 1.2 PRDMP runs as a command processor under TSO. As a result, the EXEC statement has been changed and three DD statements, SYSPRINT, SYSTSIN, and SYSTSPRT, are now required. You can, however, specify dummy DD statements for any of the three. The INDEX DD statement is new in Release 1.0.

The new Release 1.2 procedure is:

```
//PRDMP PROC DUMP=DUMPOO  
//DMP EXEC PGM=IKJEFT01, PARM=AMDPRDMP  
//SYSTSIN DD DUMMY, DCB=(RECFM=F, LRECL=80, BLKSIZE=80)  
//SYSTSPRT DD DUMMY  
//SYSPRINT DD SYSOUT=A  
//TAPE DD DSN=SYS1.&DUMP, DISP=SHR  
//INDEX DD SYSOUT=A  
//PRINTER DD SYSOUT=A  
//SYSUT1 DD UNIT=SYSDA, SPACE=(4104, (1027, 191))
```

The EXEC statement must invoke IKJEFT01, the TSO terminal monitor program. Specifying PARM=AMDPRDMP on EXEC causes IKJEFT01 to invoke PRDMP. You can also specify on PARM additional parameters for AMDPRDMP, as in previous PRDMP procedures. The same parameters are valid.

SYSPRINT, which previously was optional, directs system messages (except those IKJEFT01 issues) to the specified data set. If you use a dummy statement, the system does not log those messages.

SYSTSIN and SYSTSPRT are both DD statements that IKJEFT01 uses. SYSTSIN specifies a data set that contains commands or subcommands IKJEFT01 is to execute. SYSTSPRT identifies the data set in which the system is to log messages that IKJEFT01 issues. If you specify a dummy SYSTSPRT DD statement, the system does not log the messages.

The INDEX DD statement requests that PRDMP write the dump index to a sequential data set other than the PRINTER data set. If the INDEX DD statement precedes the PRINTER DD statement, the index is printed before the dump. If the INDEX DD statement is missing, PRDMP prints the index on the PRINTER data set after the dump.

RMF Procedure

RMF Version 3 contains a new procedure for starting RMF. The new procedure adds one statement:

RMF Version 3 procedure:

```
// ...  
// ...  
// ...  
//IEFPARM DD DDNAME=IEFRDER (new statement)  
//IEFRDER DD DSN=SYS1.PARMLIB,DISP=SHR
```

The IEFPARM statement must come before IEFRDER.

The new statement is required because RMF Version 3 opens IEFPARM, not IEFRDER, as does RMF Version 2. The IEFRDER statement is necessary to allow operators to override or specify additional options on IEFRDER via the START command. (See the *RMF Reference and User's Guide* for details.)

The RMF Version 3 procedure cannot start RMF Version 2 because the Version 2 procedure opens IEFRDER. When processing the Version 3 procedure, the system associates the data set information on the IEFRDER statement with the IEFPARM ddname, then deletes the IEFRDER name. Therefore, if MVS/370 executes the Version 3 procedure, RMF cannot find the IEFRDER statement when it attempts to open SYS1.PARMLIB.

You can, however, modify the IEFPARM statement in the Version 3 procedure to obtain a procedure that can start either RMF Version 2 or 3. Replace the IBM-supplied IEFPARM statement with the following one:

```
// ...  
// ...  
// ...  
//IEFRDER DD DSN=SYS1.PARMLIB,DISP=SHR  
//IEFPARM DD DSN=*.IEFRDER,DCB=*.IEFRDER,  
VOLUME=REF=*.IEFRDER,DISP=SHR
```

Notice that, unlike the original IEFPARM statement, the modified statement must come *after* the IEFRDER statement. You can specify additional data set information on IEFRDER. However, you must also specify the same keywords and values on IEFPARM or the system ignores the information.

Using Default RNLs

You can IPL a system at the Release 1.2 level without modifying the IBM-supplied resource name lists (RNLs). Previously, if GRS=JOIN or GRS=START, you had to include entries in the SYSTEMS exclusion RNL for global RESERVE requests that the system issues during IPL processing (for example, the RESERVE request for the system master catalog). If earlier releases encounter a global RESERVE request before global resource serialization is initialized and the request is not in the SYSTEMS exclusion RNL, the system stops. (Global ENQs are not mentioned here because the system issues none before resource serialization is initialized.)

Release 1.2 IPL processing treats global RESERVE requests as local requests until global resource serialization is initialized. For each global RESERVE processed, the system issues message ISG066I, which states that the resource is temporarily excluded from global processing.

In the following situations, however, you still need to add entries to the default SYSTEMS exclusion RNL:

- Other systems in the global resource serialization complex have additional entries in the RNL (for example, systems at earlier levels of MVS, which require entries in the SYSTEMS exclusion RNL.) The RNLs of all systems in the complex must be identical.
 - Some global RESERVE requests are to be treated as local requests after global resource serialization is initialized.
-

Duration of the RMF Initialization Process

The first RMF initialization following an IPL takes up to a minute longer in MVS/XA than in MVS/370. The increase represents the time required for the initialization routines to obtain configuration information from the IOCDS. The MVS/370 initialization routines do not use comparable information.

Generating Stand-Alone Dump

You can generate the stand-alone dump program (SADMP) and initialize the volumes on which it resides in one batch job instead of two, as required in MVS/370. The old two-step procedure still works. *SPL: Service Aids* describes how to perform the same functions in one step. You can use either Assembler H Version 1 or Version 2 to generate MVS/XA SADMP. Do not use Assembler XF.

MVS/SP Version 2 introduces several other improvements to stand-alone dump that might affect how you code the AMDSADMP macro instruction. The next topic describes the changes.

Stand-Alone Dump Macro Instruction Changes

The stand-alone dump macro (AMDSADMP) has several new and changed keywords that improve its usefulness. The following figure summarizes the changes:

Keyword	Status	Description of Change
CONSOLE	Changed	<p>Accepts up to 21 device addresses and device types in MVS/XA. In MVS/370, CONSOLE identifies only one device. The default device type is also changed. The MVS/XA default is 3278, the MVS/370 default is 3215. The default device address (01F) remains the same.</p> <p>In MVS/XA, CONSOLE must include the addresses of all consoles that SADMP can use. Unlike the MVS/370 version, MVS/XA SADMP does not accept interrupts from any console not listed. Also, after performing an SADMP IPL, the MVS/XA operator must press the ENTER or ATTN key to signal which console SADMP is to use.</p>
DUMP	New	<p>Allows the user to select storage areas to be dumped in an unformatted dump. The areas specified are in addition to the areas that a stand-alone dump normally includes. DUMP is valid only for high speed stand-alone dumps.</p>
LOADPT	New	<p>Specifies an absolute address where the stand-alone dump real storage dump module (AMDSARDM) is to be loaded. LOADPT allows users to avoid bad or offline storage.</p>
MSG	New	<p>Requests that SADMP display only messages that require action. If MSG=ACTION is not coded, SADMP displays both information and action messages. Suppressing information messages speeds up dump processing.</p>
PROMPT	New	<p>Requests that SADMP prompt the operator at execution time for additional storage areas to be dumped. PROMPT provides the same function as the DUMP keyword, but allows the operator to make storage requests at the time a dump is taken instead of when SADMP is generated. You can specify PROMPT on the same macro as DUMP. Like DUMP, PROMPT is valid only for high speed stand-alone dumps.</p>

Figure 2-4. Stand-Alone Dump Macro Instruction Changes

Chapter 3. Programming Considerations

This chapter describes differences that might affect user-written assembler programs, including user modifications to the system. The changes are grouped according to the type of programs each might affect.

Changes that Might Affect Unauthorized Programs

With few exceptions, programs that use only unauthorized services and published external interfaces will work unchanged in MVS/XA. The exceptions include programs that use the following macro instructions:

- IOHALT
- IOSGEN UCBLINK
- RESETPL
- SPIE (in two circumstances only)
- STATUS with the STOP,SYNCH option specified

You can modify the affected programs before installing MVS/SP Version 2. See the topics describing each macro.

Some unauthorized programs might also require modification because of changes described in “SDWA Changes” on page 3-9 and “Differences in GETMAIN Processing” on page 3-9.

Topics describing changes that apply to all programs are:

- “CHKPT Macro Instruction” on page 3-6
- “TSO TEST Command” on page 3-10
- “Deleted Instructions” on page 3-11
- “Macro Expansions in JES Modifications” on page 3-11
- “Limiting Concurrent Global Resource Serialization Requests” on page 3-11
- “Format Changes to Hardcopy Log Records” on page 3-12
- “Link Editing Allocation User Routines” on page 3-13
- “Removal of the Interval Timer” on page 3-13
- “Changed Instructions” on page 3-28
- “Summary of New and Changed Macros” on page 3-37

Changes that Might Affect Authorized Programs

Authorized programs are those that execute either in:

- Supervisor state (bit 15 in the PSW is zero)
- A system key (bits 8-11 in the PSW are in the range 0-7)
- As part of an APF-authorized job step task (bit JSCBAUTH in the JSCB is 1)

Although many authorized programs will work unchanged in MVS/XA, you might have to modify some. Those most likely to require modification are programs that:

- Use system interfaces that are not documented externally.
- Communicate directly with system modules (for example, via branch entry)
- Access system control blocks that are changed or that have been moved to virtual storage above 16 Mb

- Modify the system
- Have dependencies on the names or virtual storage locations of system modules

In addition to topics already mentioned, the following describe changes that might affect authorized programs:

- “Checklist for Determining if Authorized Programs Must be Changed” on page 3-13
- “Changes to the SVC Table” on page 3-17
- “Changes to the Locking Structure” on page 3-17
- “Determining Which Locks a Processor Holds” on page 3-17
- “Page Protection” on page 3-17
- “PSA Low Address Protection” on page 3-18
- “Fetch-Protected PSA Areas” on page 3-19
- “Patch Areas in the PSA” on page 3-19
- “Real Addressing Considerations” on page 3-19
- “Cross Memory Entry Table Entries” on page 3-22
- “Interfaces to System Services” on page 3-22

31-bit Addressing Considerations

During the migration period, most users do not need to be concerned with addressing mode. The information in this section is for installations with programs that:

- Access system control blocks that have been moved to virtual storage above 16 Mb.
- Use unpublished internal interfaces to communicate with system programs that must be entered in 31-bit addressing mode.

The following topics describe changes that support 31-bit addressing:

- “Using the EXCPVR Macro Instruction” on page 3-19
- “Interfaces to System Services” on page 3-22
- “31-bit Addressing Considerations” on page 3-24
- “Changed Instructions” on page 3-28
- “New Instructions” on page 3-30
- “Modifying Programs that Invoke Modules Above 16 Mb” on page 3-31
- “Retrieving Data from a Control Block Above 16 Mb” on page 3-34
- “Performing I/O in 31-bit Addressing Mode” on page 3-35
- “Using the EXCP Macro” on page 3-36
- “Entry Points in IEFW21SD” on page 3-37

New Function

The information in this section does not affect existing programs. The topics describe new function available to both authorized and unauthorized users:

- “Using the EXCPVR Macro Instruction” on page 3-19
- “New Instructions” on page 3-30
- “Using the EXCP Macro” on page 3-36
- “Summary of New and Changed Macros” on page 3-37
- “New Parameters on the GETMAIN Macro Instruction” on page 3-42
- “SDUMP Macro Instruction” on page 3-43

- “SETLOCK RELEASE,TYPE=(reg) | ALL Macro Instruction” on page 3-43
- “Using GTF to Trace User Events” on page 3-44
- “Unit Verification” on page 3-44

“Chapter 9: Coexistence Considerations” contains additional considerations for maintaining programs that must run in both MVS/XA and MVS/370.

The following topic, “Macro Instructions Mentioned in This Publication,” lists the executable macros that might require attention and indicates why.

Macro Instructions Mentioned in This Publication

Figure 3-1 lists the executable macros that might require attention when converting to MVS/XA. The macros are included in the list for one or more of the following reasons:

1. The MVS/XA expansion of the macro will not work in MVS/370 (that is, the macro is downward incompatible). The downward incompatibility affects only programs that must run on both MVS/370 and MVS/XA systems. If such a program uses one of these macros, ensure that the program either:
 - Includes the MVS/370 expansion of the macro
 - Has two paths, one that MVS/370 executes, the other that MVS/XA executes

For more information, see “Handling Downward Incompatible Macros” in Chapter 9.
2. The MVS/370 macro expansion does not work in MVS/XA. You must reassemble, using the MVS/XA MACLIB, all programs that use these macros. For most of the macros, you can install compatibility PTFs or compatible program products on an MVS/370 system and reassemble the affected programs before installing MVS/XA. See the references in the notes column.
3. The macro expansion passes different parameters to the associated service routine. The parameter changes affect only programs that generate, test, or alter the parameters. You must change those programs. “Appendix A. Parameter Changes in Incompatible Macros” describes the differences between the MVS/370 and MVS/XA parameter lists.
4. The MVS/XA expansion is required if used in programs that execute in 31-bit addressing mode. Thus, you must assemble such programs using Assembler H Version 2 and the MVS/XA MACLIB.
5. The macro provides new function. The list includes macros that are new or that have new parameters or new options on existing parameters. The added function does not affect existing programs. Figure 3-6 summarizes the new function that each macro provides.

You must use the MVS/XA MACLIB to assemble programs that use new functions. With two exceptions, the programs will then not work on MVS/370 systems. New GETMAIN options are one exception. The new AMODE=24 option on the SYNCH macro is the other. See “New Parameters on the

GETMAIN Macro Instruction” in this chapter and “Downward Incompatible SYNCH Macros” in Chapter 9 for more detail.

6. The macro requires attention for another reason. The notes column refers to the topic that describes the reason.

System Macros and Facilities documents authorized macros. *Supervisor Services and Macros* documents unauthorized macros. Some of the macros listed as unauthorized have parameters that only authorized users can specify. Those macros are documented in both publications.

Unauthorized Macros	1. Incompatible MVS/XA Expansion	2. Incompatible MVS/370 Expansion	3. Parameter Changes	4. 31-Bit Mode Requirement	5. New Function	6. Other	Notes
ABEND					x		
ATTACH	x		x	x			The MVS/XA expansion is required only if parameter addresses are greater than 16 Mb.
BLSABDPL					x		
BLSQMDEF					x		
BLSQMFLD					x		
BLSRESSY					x		
BTAM RESETPL		x	x				See "RESETPL (BTAM) Macro Instruction."
CALL				x			The MVS/XA expansion is required only if VL is specified when the program executes in 31-bit mode.
CHKPT	x			x		x	See "CHKPT Macro Instruction."
CPOOL					x		
CPUTIMER					x		
ENQ					x	x	See "Limiting Concurrent Global Resource Serialization Requests."
ESPIE					x		
ESTAE	x		x	x			
EVENTS	x		x	x			
EXCP					x		See "Using the EXCP Macro."
GETMAIN					x		See "New Parameters on the GETMAIN Macro Instruction."
GQSCAN						x	See "Limiting Concurrent Global Resource Serialization Requests."
GTRACE					x		
LINK				x			The MVS/XA expansion is required only if PARAM, VL=1, DE, or ERRET is specified.
IOHALT		x	x				See "IOHALT Macro Instruction (SVC 33)."
IOSGEN UCBLK						x	Deleted. See "IOSGEN UCBLK Macro Instruction."
LOAD					x		
PGSER					x		
RESERVE					x	x	See "Limiting Concurrent Global Resource Serialization Requests."
RETURN						x	See the entry for RETURN in "Summary of New and Changed Macros"
SETRP					x		
SMFEXIT	x		x	x			

Figure 3-1 (Part 1 of 2). Unauthorized Macro Instructions Mentioned in This Publication

Unauthorized Macros	1. Incompatible MVS/XA Expansion	2. Incompatible MVS/370 Expansion	3. Parameter Changes	4. 31-Bit Mode Requirement	5. New Function	6. Other	Notes
SMFIOCNT					x		
							SNAP x
SPIE						x	Works differently in some situations. See "Differences in SPIE Processing."
SPLEVEL					x		See "Handling Downward Incompatible Macros" in Chapter 9.
STATUS STOP,SYNCH		x					See "STATUS STOP,SYNCH Macro Instruction."
STAX	x		x	x			The MVS/XA expansion is required only if parameter addresses are greater than 16 Mb.
STIMER	x		x	x			
STIMERM					x		
SYNCH	x				x	x	See "Downward Incompatible SYNCH Macros" in Chapter 9.
WTL					x		
WTO					x		See the entry for WTO in "Summary of New and Changed Macros."
WTOR	x		x	x	x		
XCTL				x			The MVS/XA expansion is required only if PARAM or VL=1 is specified.

Figure 3-1 (Part 2 of 2). Unauthorized Macro Instructions Mentioned in This Publication

Authorized Macros	1. Incompatible MVS/XA Expansion	2. Incompatible MVS/370 Expansion	3. Parameter Changes	4. 31-Bit Mode Requirement	5. New Function	6. Other	Notes
CALLDISP				x			
CALLRTM					x		
CIRB					x		
DATOFF					x		
EXCPVR					x		See "Using the EXCPVR Macro Instruction."
FESTAE	x			x			
IOSLOOK					x		Replaces other methods of locating UCBs. See "IOSGEN UCBLook Macro Instruction."
INTSECT	x			x			
MODESET				x			The MVS/XA expansion is required only if EXTKEY=RBT234 is specified.
MGCR					x		
NUCLKUP					x		
PTRACE					x		
SCHEDULE	x						Downward incompatible only if SCOPE=GLOBAL is specified.
SDUMP	x		x	x	x		Downward incompatible only if new parameters are specified. See "SDUMP Macro Instruction."
							See also the entry for SDUMP in "Summary of Macros that Provide New Function."
SETFRR	x						Downward incompatible only if INLINE=YES is specified.
SETLOCK	x		x		x		Downward incompatible only if RELEASE TYPE=(reg) ALL is specified. See "Determining Which Locks a Processor Holds" for an example of new SETLOCK function.
SVCUPDTE					x		
VSMLIST					x		
VSMLOC					x		
VSMREGN					x		

Figure 3-2. Authorized Macro Instructions Mentioned in This Publication

CHKPT Macro Instruction


User programs that successfully take checkpoints in MVS/370 can take checkpoints in MVS/XA. However, a program that has taken a checkpoint must restart on the same operating system (either MVS/370 or MVS/XA).

IOHALT Macro Instruction (SVC 33)


You must recompile modules that use the IOHALT macro and modify programs that issue SVC 33 directly (without using IOHALT). The SVC 33 service routine requires different input in registers 0 and 1:

MVS/370 Input

Register 0




Register 1




Options: X'00' - Halt I/O
X'80' - Modify the CCW

MVS/XA Input

Register 0



Register 1



Options: X'01' - Halt I/O
X'81' - Modify the CCW

PTF UZ29156 provides the new version of IOHALT and the new SVC 33 interface that are compatible with MVS/XA. You can install the PTF on an MVS/370 system and reassemble and modify the affected programs before installing MVS/SP Version 2. (You do not have to reassemble programs that will not be run on an MVS/XA system.) The reassembled programs will work on both MVS/370 and MVS/XA systems.

If you use GTF to trace modules that use IOHALT or SVC 33 and you install the IOHALT compatibility PTF, you might need to install an additional PTF on MVS/370. Unless you have installed MVS/SP Version 1 Release 1 or a later release, install either:

- PTF UZ32985 for systems with SE2 installed
- PTF UZ32984 for systems with SE1 installed
- PTF UZ32983 for MVS 3.8 systems with neither SE1 nor SE2 installed

The PTFs allow GTF to trace programs that use either the old or the new SVC 33 interface. MVS/SP Version 1 Release 1 and later releases incorporate the PTF changes.

IOSGEN UCBLook Macro Instruction

You must change programs that use the IOSGEN UCBLook macro or that directly access the UCB look-up table. Neither the IOSGEN UCBLook function nor the UCB look-up table is supported in MVS/XA.

To obtain UCB addresses in MVS/XA, use either:

- The UCB scan routine (IOSVSUCB)
- The IOSLOOK macro

IOSVSUCB allows you to scan each UCB in the system or in a specified device class. Each time it is invoked, IOSVSUCB returns the address of one UCB's common segment. To scan several UCBs, invoke IOSVSUCB repeatedly. Both authorized and unauthorized programs can use IOSVSUCB.

IOSLOOK returns the address of the common segment of the UCB associated with a given device number. Unlike IOSVSUCB, IOSLOOK requires that users be in supervisor state.

Both IOSVSUCB and IOSLOOK are documented in *System Macros and Facilities*.

Both services are also available in MVS/370. MSV/SP Version 1 Release 3 and later releases include IOSVSUCB. PTF UZ28392 includes IOSLOOK. Therefore, you can change the affected programs before installing MVS/XA. The changed programs will run on both MVS/370 and MVS/XA.

RESETPL (BTAM) Macro Instruction

Programs, including CICS and IMS BTAM modules, that use the OS/VS BTAM expansion of RESETPL will not work in MVS/XA. You must reassemble them using the RESETPL macro included in BTAM/SP (5665-279). (BTAM/SP is required to run BTAM application programs and subsystems in MVS/XA.)

You can install BTAM/SP on MVS/370 and reassemble the affected programs before installing MVS/XA. The reassembled programs will work on both MVS/370 and MVS/XA systems.

Note: Because the RESETPL expansion issues an IOHALT macro, you must also install the IOHALT compatibility PTF on MVS/370 before reassembling the programs. See “IOHALT Macro Instruction (SVC 33)” for more detail.

Differences in SPIE Processing

Most programs using SPIE macros will continue to work correctly in MVS/XA. However, you need to modify programs that create a SPIE to protect a program running under a different RB.

MVS/XA terminates the SPIE when the program that created it completes, whether normally or abnormally. In MVS/370, the SPIE usually remains in effect until all programs in the step complete (task termination time). The exception occurs when the program that creates the SPIE abends. If that happens, MVS/370 terminates the SPIE also.

Following are two examples of programs that do not work the same in MVS/XA as in MVS/370:

- Program A links to Program B, which issues a SPIE macro and returns. In MVS/XA, the SPIE is deleted. In MVS/370, it remains in effect.
- Program A issues a SPIE macro, followed by an XCTL macro to invoke Program B. In MVS/XA, the SPIE is deleted. In MVS/370, the SPIE is in effect for Program B.

You can change affected programs before installing MVS/XA.

STATUS STOP,SYNCH Macro Instruction

The SYNCH operand on the STATUS STOP macro is no longer supported. Change programs that issue STATUS STOP,SYNCH to issue STATUS STOP without the SYNCH operand. You can make the changes before installing MVS/XA.

SDWA Changes

The SDWA is increased in size. The additional storage is included in previously existing or new SDWA extensions. The sizes of the FRR work area and the ESTAE save area remain the same.

Programs that use indirect pointers into the SDWA work unchanged in MVS/XA. However, you must modify programs that:

- Depend on the 72-byte save area passed to ESTAE exits being located at a given offset into the SDWA. MVS/XA uses register 13 to pass to FRRs the address of the user save area, as does MVS/370.
 - Depend on the 200-byte FRR work area that is passed to FRR routines being located at a given offset into the SDWA. MVS/XA uses Register 0 to pass to FRRs the address of that work area, as does MVS/370.
 - Use explicit length values to free the SDWA. Modify the programs to use the value in the SDWALNTH field.
 - Depend on the order of the SDWA and its extensions. Modify the programs to use indirect pointers.
 - Place data in the SDWA variable recording area (VRA) without updating the SDWAURAL field. Programs must maintain an accurate count in the SDWAURAL field to prevent data from being overlaid.
 - Assume that the unused section of the SDWA contains zeros. Programs need to ignore data in the unused area.
-

Differences in GETMAIN Processing

Two differences in GETMAIN processing might cause programs to fail in isolated instances:

- Although the MVS/XA GETMAIN service routine does not introduce any new parameter restrictions, it does enforce some restrictions that were documented but not enforced in MVS/370. With one exception, the GETMAIN routine no longer allows the parameter list, address list, or length list specified on the LC, LU, VU, EC, or EU forms of GETMAIN to overlap. If the request is for a single element, MVS/XA allows the pointer to the address list to point to itself. All other overlaps cause the program to fail with ABEND code x'504'.

Because programs seldom use the forms of GETMAIN mentioned or overlap parameters, you probably will not want to spend time looking for programs that have to be changed. Instead, keep in mind the parameter restrictions. If a program fails with ABEND code x'504', you can change it then.

- GETMAIN routines obtain storage differently in MVS/XA. Although, the rules GETMAIN uses have never been externally documented, at least one installation has written a program that now fails because of the change. Specifically, the program expects that additional virtual storage will be contiguous with currently owned virtual storage. If your installation has written similar programs, you need to modify them.
-

TSO TEST Command

To use TSO TEST on an MVS/XA system, you must install the MVS/XA version of TSO/E (5665-285). TSO TEST is not part of the MVS/XA base control program. If you issue TSO TEST and have not installed TSO/E for MVS/XA, you receive message IKJ56500I stating that the TEST command is not found. In addition, unless TSO/E for MVS/XA is installed, user programs that issue either SVC 61 or SVC 97 receive a return code of 4.

Following are TSO TEST compatibility considerations:

- You can test on MVS/370 programs created in MVS/XA as long as the programs do not use any MVS/XA instructions, new macros, new parameters or options on existing macros, downward incompatible macros, or addresses above 16 Mb.
- When executing TSO TEST on an MVS/XA system, AT subcommands and LIST subcommands that specify the instruction data type support only MVS/XA instructions.
- User-written TEST subcommands that do any of the following will not work with the MVS/XA version of TSO TEST:
 - Use IKJEGSTA as an ESTAE exit. The parameter list that IKJEGSTA requires is incompatible.
 - Use the TCOMTAB mapping macro to access fields in the TCOMTAB control block. Many labels on the TCOMTAB macro are deleted because TSO TEST no longer uses the corresponding fields.
 - Use labels and equates in TSO TEST mapping macros to determine the length of the corresponding TEST control blocks. The names of the equates used to define the lengths of control blocks are changed.
- Installations that altered the TSO TEST subcommand table (IKJEGSCD) must rebuild their changes in the new table. To update the table:
 - Copy the IKJEGSCD CSECT of assembly module IKJEGMNL in the TEST load module into a separate data set.
 - Make the required changes.
 - Assemble and again link edit IKJEGSCD into TEST.

The IKJEGSUB macro that generates IKJEGSCD in MVS/370 is deleted from TSO/E for MVS/XA.

- UNALLOC is now an alias for the TSO FREE command. AND and OR are new subcommands of TEST. In type 32 SMF records, the UNALLOC command and the AND and OR subcommands of TSO TEST are recorded as *OTHER, unless your installation adds those names to CSECT IEEMB846.
-

Deleted Instructions

The following instructions are deleted from the standard 370-XA instruction set:

ISK (Insert Storage Key)
SSK (Set Storage Key)
All 370 I/O instructions

ISK and SSK are deleted because MVS/XA replaces segment protection with page protection.

Macro Expansions in JES Modifications

The MVS/XA expansions of fourteen macros are downward incompatible. That is, their MVS/XA expansions will not work in MVS/370. The MVS/XA MACLIB contains both the MVS/370 and MVS/XA expansions of those macros. JES2 and JES3 both require the MVS/370 expansions. Therefore, JES system programs that issue downward incompatible macros use the SPLEVEL macro to ensure that the MVS/370 expansions are generated at assembly time. (SPLEVEL specifies which level is to be used.) The programs, however, do not issue SPLEVEL directly. Each includes another macro, whose expansion issues SPLEVEL. JES2 programs use \$HASPGEN in Release 1.0 and \$HASPEQU in Release 1.1 and later releases. JES3 programs use IATYMOD.

If you modify JES system programs and your modification includes a downward incompatible macro, either:

- Ensure that your modification comes after the \$HASPGEN, \$HASPEQU, or IATYMOD macro. (The macros appear near the beginning of programs.)
- Use SPLEVEL to ensure that the MVS/370 expansion of the downward incompatible macro is assembled.

“Handling Downward Incompatible Macros” in Chapter 9 lists the macros and describes SPLEVEL.

Limiting Concurrent Global Resource Serialization Requests

In Release 1.1, global resource serialization limits the number of ENQ, RESERVE, and certain types of GQSCAN requests a single job, started task, or TSO user can have outstanding at a given time. The GQSCAN requests it limits are those that specify the TOKEN parameter. The change is designed to prevent one address space from using up all of GRS virtual storage, which causes subsequent GRS requests to fail.

Generally, the new processing does not require any action on your part. However, you need to be aware of the changes. Users might receive new ABEND or return codes indicating their programs failed because of too many concurrent global resource serialization requests. Also, you might want to change the limits global

resource serialization enforces, although the default values are satisfactory for most installations.

To enforce the limit, global resource serialization uses a new threshold for each address space. The thresholds are in the GVTCREQ fields of the GVTs. If the number of outstanding ENQ, RESERVE, or specific types of GQSCAN requests reaches the threshold (the default is 4096), global resource serialization:

- Rejects subsequent ENQ and RESERVE requests from *unauthorized* callers in the address space. The system terminates unconditional requests with ABEND code x'538', and rejects conditional requests with a return code of x'014'. In earlier releases, if GRS virtual storage is depleted, users receive a return code of x'08'.
- Allows *authorized* callers in the address space to issue a limited number of additional ENQ and RESERVE requests. The number cannot exceed the tolerance value specified in the GVTCREQA field of the GVT. The tolerance value is also new. Its default value is 4111. Global resource serialization allows authorized callers the additional ENQ and RESERVE requests to enable recovery and normal termination routines to obtain the resources required to finish processing.
- Rejects with a return code of x'14' GQSCAN requests that specify the TOKEN option and request more information than can fit into the caller's buffers. Global resource serialization returns the buffers of information but does not continue the scan as it normally would. (If the threshold had not been reached, global resource serialization would have queued the request for continuation, returned the full buffers to the caller, and, after the caller cleared the buffers, resumed the scan.)

If you find the threshold and tolerance values in the GVT are too high or low, you can change them for your installation using the AMASPZAP service aid. For details, see *Service Aids*.

Format Changes to Hardcopy Log Records

The formats of all hardcopy log records except those written using JES3 are changed in Release 1.2 to provide additional machine-readable information. As a result, you need to modify most programs that scan the SYSLOG data set. Scan programs that run on a JES3 system might work unchanged. However, be aware that records logged before JES3 is initialized and all records written via the LOG command or the WTL macro have the new format. If your installation keeps the hardcopy log on a JES2 multi-access spool that systems at earlier levels can access, the scan programs must be sensitive to which system wrote the record.

The new log format includes the following additional information:

- A record ID, which identifies the type of record written (for example, a WTOR, label line, or command response). The record ID appears only in SYSLOG and not in printed output.
- The system ID.
- The date the message was issued.

- The ID of the console from which the command or command response was issued.
- User exit and message suppression flags.
- The full text of the message. If the text requires more than one line, one or more WTL entries might be interspersed among the continuation lines. If printed, however, the text appears on consecutive lines.

The text of entries written using the LOG command begin with the prefix '0'. The text of entries written using the WTL macro begin with either a user-specified prefix or an 'X'.

A new macro, IHAHCLOG, maps the new record format. When modifying your programs, use the mapping macro instead of offsets to access the data.

Link Editing Allocation User Routines

Release 1.1 removes the following routines from the device allocation load module, IEFW21SD, and places them in their own single-CSECT load modules. Therefore, you need to link edit them differently:

IEFDB401 - Dynamic allocation user exit
 IEFXVNSL - Non-standard tape label routine
 IEFAB445 - Allocation space defaults CSECT
 IKJEFD00 - Dynamic allocation interface routine

IEFDB401 and IEFXVNSL can reside above or below the 16 Mb line. Specify their RMODEs on the link edit statements for each. The sysgen link edit control statements omit the RMODE specification. IEFAB445 resides below the 16 Mb line. IKJEFD00 resides above that line, but you can invoke it via an interface routine, IKJDAIR, in either 24- or 31-bit addressing mode.

Release 1.1 changes some of the entry points in IEFW21SD. Programs to be executed in 31-bit addressing mode must use the new rather than old entry points. See "Entry Points in IEFW21SD" for more information.

Removal of the Interval Timer

The 370-XA architecture deletes the interval timer. Modify programs that use the interval timer to use the CPU timer instead. Although the CPU timer works like a stopwatch, you can use it like an interval timer. Use the STIMER macro to set the CPU timer, the CPUTIMER or TTIMER macro to obtain its current value, and the SRBTIMER macro to set a time limit for SRB processing. Although you can use either the CPUTIMER or TTIMER macro, CPUTIMER is faster and you can use it in SRB or task mode. You can use TTIMER in task mode only.

Checklist for Determining if Authorized Programs Must be Changed

The following checklist is for use in examining **authorized** assembler programs for incompatibilities and is not applicable to other programs. Programs written in high level (non-assembler) languages are compatible and require no change. Most unauthorized assembler programs also work unchanged. The few exceptions are noted in the introduction to this chapter.

Most of the following programs require modification and/or reassembly:

- Programs that issue any of the following macros:

RESETPL (a BTAM macro)
IOHALT (or SVC 33)
IOSGEN UCBLK
STATUS STOP,SYNCH

In all cases, you can change programs that use these macros before installing MVS/XA. For details, see the topics describing the macros.

- Programs that access system control blocks that are changed or that now reside in virtual storage above 16 Mb. “Appendix B. Control Block Changes” lists the control blocks requiring attention.
- Programs that directly invoke system modules that now require entry in 31-bit addressing mode, or that require parameter addresses to be 31-bit values.

Programs using SVCs or published macros to invoke service routines that now execute in 31-bit addressing mode generally work unchanged in MVS/XA. In most cases, the macro invokes a routine that changes modes, if necessary, before entering the service routine.

The MVS/XA components having a large percentage of modules that execute in 31-bit addressing mode include: contents supervision (CSV), GTF, IOS, RMF, RSM, RTM, SRM, system trace, and VSM in Release 1.0; ALLOCATE in Release 1.1; SVCDUMP in APAR OZ78216; and VSAM record management load modules and contents supervision in Release 1.2.

See “Interfaces to System Services” for more detail.

- Programs that use the high-order byte of address fields for flags. When running in 31-bit addressing mode, MVS/XA treats addresses as 31-bit values and, if applicable, uses the high-order bit to set the PSW A-mode bit.

Specific examples of programs that will fail include those that use the high-order byte of:

- Address fields they pass to IARUTRV (translate real to virtual routine). IARUTRV, which replaces IEAVTRV in MVS/XA, treats the real addresses as 31-bit values.
 - The SRBEP or SRBRMTR field in the SRB. MVS/XA treats each field as a 31-bit value, and uses the high-order bit to set the PSW A-mode bit.
 - The SVC screening table address in the TCB (the TCBSVCA2 field). MVS/XA also treats that address as a 31-bit value and uses the high-order bit to set the PSW A-mode bit.
- Programs that treat UCB addresses as 2-byte values. In MVS/XA, UCB addresses are three bytes instead of two.

- Programs that directly access the UCB look-up routine. It does not exist in MVS/XA. "IOSGEN UCBLook Macro Instruction" describes alternate ways of obtaining the same information.
- Programs that depend on the structure of the nucleus and the FLPA. In MVS/XA, the FLPA no longer resides in the nucleus buffer. Also, neither the nucleus nor the FLPA is mapped V=R, and modules in those areas might not be loaded into contiguous real frames.

Examples of programs that must be modified are:

- V=R programs that use EXCP to perform I/O into or out of the FLPA.
- Programs in the nucleus or FLPA that run DAT-off. "DAT-off Restrictions" describes how to change the programs.
- Programs that use the CVTNUCB field to determine if they have been loaded into the FLPA. Change the programs to test the CVTFLPAS and CVTFLPAE fields to determine residency in the FLPA below 16 Mb, and the CVTEFLPS and CVTEFLPE fields for residency above 16 Mb. Each pair of fields indicates the beginning and ending addresses of the FLPA areas.
- Programs sensitive to virtual storage location changes; for example, programs that treat parameter addresses below 64 K as invalid.
- Programs sensitive to changes in the locking structure. Programs requiring modification are those that:
 - Use the IOSCAT or the IOSLCH lock. Those locks are deleted in MVS/XA.
 - Obtain the DISP lock in order to hold the highest lock in the system.
 - Request the DISP lock after obtaining the ASM lock.
 - Use the PSAHLHI field to determine locking heirarchy.

For more information, see "Changes to the Locking Structure" and "Determining Which Locks a Processor Holds."

- Programs that use the following system-created data:
 - GTF, system trace, or LOGREC records. The record formats are different. Also, the structure of the system trace table is changed.
 - Programs that scan the SYSLOG data set. Release 1.2 changes the format of the hardcopy log records. See "Format Changes to Hardcopy Log Records."
 - Dump data. Dump contents and formats have changed. See Chapter 6 for more information.
 - SMF data that is changed. Chapter 7 identifies which SMF records are changed and briefly describes the differences.

- Programs that examine the PSW field in MVS/XA control blocks (for example, programs that use trace data or print reports). The PSW format is changed. Among other differences, the instruction addresses are contained in 4-byte, instead of 3-byte, fields.
- Programs that use the LRA instruction. LRA always returns a 31-bit address in MVS/XA, even when executed in 24-bit addressing mode.
- Programs that call IEFSCAN or that directly access MVS/370 device allocation tables (DEVNAMET, IEFDEVPT, and DEVMASKT). IEFSCAN and the tables are deleted in MVS/XA. “Unit Verification” describes how both authorized and unauthorized programs can perform unit verification in MVS/XA.
- Programs that use extended ECBs for POST exits. The programs must be authorized to fetch from the ECB extension, as well as to fetch and store the extended ECB.
- Programs that specify an ACON length other than 4 (for example, AL3(location)), if the location in parentheses is above 16 Mb.
- Programs that examine the SVTDACTV or SVTPWAIT fields in the SVT (usually programs that code their own expansions of SCHEDULE or INTSECT, respectively). The offsets of these fields in the SVT have changed. Their previous locations are initialized to x'FFFFFFFF'.
- Programs that depend on CPU (processor) addresses being 0, 1, or 2. A CPU address can be any number from 0-F.
- I/O drivers that call IEASMFEX to record EXCP counts. Change the drivers to use a new SMF macro, SMFIOCNT.
- If your installation includes data sets in the LNKLST concatenation that are not APF authorized, programs that depend on the data sets being APF authorized.

The DEBAPFIN bit in the LNKLST DEB indicates whether or not all data sets in the LNKLST concatenation are APF authorized. The LLTAPFIN field in a data set's LLTAPFTB entry indicates whether the data set is APF authorized. The LLTAPFTB is a new extension to the LLT that contains one entry for each data set in the LNKLST concatenation. See “Using a New Directory for LNKLST Data Sets” for more information.

- Programs that access SMF BQEs (buffer queue elements). Release 1.1 moves the BQEs from common storage to the new SMF address space.
- Programs that read SMF data sets directly instead of via SMF dump programs. Release 1.1 initializes SMF data sets with dummy records that are shorter than valid SMF records. They contain the characters 'SMFEOFMARK.'
- Programs that obtain storage for data extent blocks (DEBs) from fetch-protected areas (subpools 0-172). If called to add a DEB table entry for a DEB that is in fetch-protected storage, the MVS/XA DEBCHK service routine issues ABEND x'16E' with reason code x'1C'. MVS/370 does not impose the same restriction.

Changes to the SVC Table

The following changes have been made to the SVC table:

SVC	Description of Change
SVC 16 (PURGE)	Changed from type 3 to type 2
SVC 46 (TTIMER)	Changed from type 3 to type 2
SVC 47 (STIMER)	Changed from type 3 to type 2
SVC 82 (DASDR)	Deleted
SVC 88 (MOD88)	Deleted
SVC 109 (Extended SVC Router)	A new entry has been added: 28 - ESPIE, a type 3 SVC
SVC 138 (PGSER)	A new type 2 SVC

Changes to the Locking Structure

The locking structure has changed in MVS/XA:

- MVS/XA uses nine new locks instead of the SALLOC lock for storage management serialization.
- A new TRACE lock serializes the system trace buffer structure.
- A new CPU lock causes the requestor to be physically disabled for I/O and external interrupts. It provides system-recognized disablement.
- The IOSCAT and IOSLCH locks have been deleted.
- The hierarchy of the ASM and DISP locks is reversed. In MVS/XA, the ASM lock's position is above the DISP lock's position in the locking hierarchy.

You must change programs that:

- Use the IOSCAT or IOSLCH locks.
 - Use the SALLOC lock to serialize storage management.
 - Obtain the DISP lock in order to hold the highest lock in the system.
 - Request the DISP lock after obtaining the ASM lock.
-

Determining Which Locks a Processor Holds

You must change programs that use the PSAHLHI (highest lock held) field to determine locking hierarchy. The PSAHLHI field is now referred to as PSACLHS (current locks held string), although the old name is retained for compatibility. The bit positions in the PSACLHS field indicate which locks the processor owns. They no longer represent the hierarchy of locks.

MVS/XA provides a new SETLOCK service that indicates whether the processor owns any locks at a higher position in the hierarchy than the one specified as input. For example, the following SETLOCK macro tests whether the processor owns any locks at a higher position than the dispatcher lock:

```
SETLOCK TEST,TYPE=HIER,LOCK=DISP,REGS=(11,12)
```

Page Protection

MVS/XA uses a new page-protection facility to enforce read-only access to the:

- Read-only nucleus (above and below 16 Mb)
- Resident BLDL list in Release 1.0
- PLPA (above and below 16 Mb)
- MLPA (above and below 16 Mb)
- FLPA (above and below 16 Mb)
- NUCMAP (an area in the non-page-protected nucleus that maps the nucleus)

Page protection is optional only for the MLPA or FLPA. Installations can turn off page protection for those areas by specifying a new subparameter, NOPROT, on the MLPA and FIX system parameters, respectively, in the IEASYSxx PARMLIB member. When NOPROT is specified, none of the MLPA or FLPA is page-protected. The system default is to page protect those areas.

You cannot include in page-protected areas any module that stores into itself. If any program running DAT-on attempts to store into a page that is page-protected, the processor generates a program interrupt, regardless of the program's state or protect key.

Modules that modify themselves might include:

- Modules that use macros which create parameter lists, but do not use the LIST/EXECUTE forms of the macros to eliminate stores into the module
- Modules marked reentrant that page-fix, serialize, and modify themselves

If you have any such modules in the PLPA, either:

1. Modify the module so that it stores the data somewhere else; for example, in dynamically-acquired storage.
2. Include the module in another library; for example, in SYS1.LINKLIB.
3. Include the module in the MLPA or FLPA and specify NOPROT for that area.

Unless you turn off page protection in the MLPA or FLPA, handle self-modifying modules in those areas as described in 1 or 2.

The page protect facility replaces MVS/370 segment protection, which enforces read-only access to segments (64 K blocks) of storage fully occupied by PLPA pages. MVS/370 installations can override the segment protection using the AMASPZAP service aid program. Installations do not have that capability in MVS/XA.

PSA Low Address Protection

PSA low address protection prevents user programs from storing into PSA locations 0 through 511. The only way to turn off PSA low address protection in MVS/XA is by using the PROTPSA macro. The CVTPRON bit in the CVT is deleted.

To disable low address protection in MVS/370, programs can either use the PROTPSA macro or change the CVTPRON bit.

Fetch-Protected PSA Areas

You might have to change some programs that fetch data from the PSA. In MVS/XA, some PSA locations are key 0 fetch-protected. That is, only programs in key 0 can fetch data from those areas:

- The last 2 K of the PSA (2 K through 4 K minus 1) are always key 0 fetch-protected.
- The first 2 K are key 0 fetch-protected from programs running on a different central processor. However, programs require no authorization to fetch data from the first 2 K of the PSA of the central processor on which they are running. The programs must still be in key 0 to store data into the first 2 K, however.

The PSA work/save areas have been moved into the fetch-protected area to improve integrity. Therefore, you must also modify programs that access the moved data while not in key 0 and programs (usually FRRs) that fetch data from the FRR six-word parameter area while not in key 0.

Fetch-protection of PSA locations is new in MVS/XA. In MVS/370, programs have to be in key 0 to store into the PSA, but no authorization is required to reference data there.

Patch Areas in the PSA

MVS/XA uses some PSA locations that are available for system patches in MVS/370. If you use areas of the PSA for system patches, ensure that your patch applications use only areas that are not system-defined. Otherwise, normal system processing might overlay the patch.

To determine which areas are safe to use, patch applications can check whether the storage contains zeros. When initializing the PSA, MVS/XA puts non-zero values in the system-defined areas within the range most commonly used as a patch area (x'600' to x'C00'). Available areas contain zeros.

Real Addressing Considerations

You might need to change programs that:


- Use the EXCPVR macro instruction
- Depend on RSM backing virtual pages with real storage below 16 Mb
- Execute DAT-off code

Using the EXCPVR Macro Instruction

EXCPVR users need to be aware of two changes:

- The list of data areas that the page fix (PGFX) appendage passes to the I/O supervisor must contain 31-bit addresses. The high-order bit of each address must be zero.
- Using EXCPVR, CCWs, and IDAWs, programmers can perform I/O to any location in real storage (above or below 16 Mb). The channel programs must use IDAWs to specify the address of buffers in real storage above 16 Mb:

CCW (format 0, the CCW format used in MVS/370)



Because the EXCP service routine (which processes both EXCP and EXCPVR macros) supports only Format 0 CCWs, CCWs and IDAWs used with EXCPVR must reside in virtual storage below 16 Mb.

As in all cases where IDAWs are used with real addresses, an IDAW is required for each 2 K real storage boundary that the data transfer operation will cross.

Unless a program uses IDAWs, EXCPVR users must ensure that all buffers are backed by real storage below 16 Mb. EXCPVR users must assume that buffers obtained via data management access methods have real addresses above 16 Mb. When requesting buffer storage, data management access methods specify LOC=(BELOW,ANY) on the GETMAIN request. LOC=(BELOW,ANY) indicates that virtual storage can be backed with real storage above 16 Mb. (See "New Parameters on the GETMAIN Macro Instruction.") The following situation can happen:

1. An program uses an access method to open a data set. The access method obtains buffer storage that might be backed by real storage above 16 Mb.
2. The program uses the access method to read data into the buffer.
3. The program attempts to write data from the buffer using EXCPVR. If the buffer is in real storage above 16 Mb, the program does not work unless it uses IDAWs to specify the real addresses above 16 Mb.

Change the program to either obtain its own buffer or use IDAWs.

SPL: Data Management describes how to use EXCPVR.

Changes in the Way RSM Backs Virtual Storage

RSM uses different algorithms to determine whether to back a virtual page with real storage above or below 16 Mb. Generally, only users who have programs with real address dependencies need to be aware of the changes. RSM:

- Attempts to back all virtual storage **above** 16 Mb with real storage **above** 16 Mb.

- Attempts to back the following virtual storage areas **below** 16 Mb with real storage **above** 16 Mb:

SQA (except subpool 226)
 LSQA
 Nucleus
 Pageable private areas
 Pageable CSA
 PLPA
 MLPA
 Resident BLDL (in Release 1.0 only)

- Always backs the following virtual storage areas **below** 16 Mb with real storage **below** 16 Mb:

V=R regions
 FLPA
 Subpool 226 (a new subpool in SQA)

- Backs subpools 227 and 228 (fixed CSA) in virtual storage **below** 16 Mb with real storage **below** 16 Mb, except when GETMAIN requests specify LOC=(BELOW,ANY).
- When satisfying a page-fix request, RSM generally backs pageable virtual pages that reside **below** 16 Mb with real storage **below** 16 Mb. (Pageable virtual pages are pages in CSA, PLPA, MLPA, or the pageable private area.) However, in the following situations, RSM attempts to use real storage **above** 16 Mb:
 - The GETMAIN request to obtain the storage specified either LOC=(BELOW,ANY), LOC=(RES,ANY), or LOC=(ANY,ANY).
 - The PGSER macro specified ANYWHERE

Note: EXCPVR users need to be aware that DFP access methods use LOC=(BELOW,ANY) on GETMAIN requests for buffer storage.

Impact of Programmers:

RSM's page backing rules are, for the most part, compatible with the way real storage is backed in MVS/370. Because programs that have real address dependencies work with fixed storage, it is expected that most existing programs will continue to receive real addresses that are less than 16 Mb. However, you must change programs that run in 24-bit addressing mode and have real address dependencies on the nucleus, SQA, or LSQA. RSM ignores requests to fix storage in the nucleus, SQA, or LSQA because those areas are already fixed. Therefore, real addresses in those areas might be greater than 16 Mb. Modify the programs to correctly handle 31-bit addresses.

DAT-off Restrictions

You must modify programs in the nucleus or FLPA that run with dynamic address translation (DAT) turned off. In MVS/370, programs can turn DAT on or off by manipulating the system mask (using the STNSM and STOSM instructions). However, because the nucleus and FLPA are not mapped V=R in MVS/XA, modules in those areas can no longer use the STNSM and STOSM instructions to

control DAT. (Programs executing as V=R jobs can use STNSM and STOSM instructions in MVS/XA and do not have to be modified unless they refer to data outside the V=R region.)

To modify modules containing DAT-off code:

1. Move the DAT-off code to a separate module. Give the module AMODE=31 and RMODE=ANY attributes. Use as its entry point, IEAVEURn, where n is a number from 1 to 4. (MVS/XA reserves four entry points in the DAT-off nucleus for users.) Use BSM 0,14 as the return instruction.
2. In the original module (which executes DAT-on), code a DATOFF macro to invoke the DAT-off module created in the previous step. DATOFF is new in MVS/XA:

```
DATOFF INDEX=INDUSRn
```

The suffix of INDUSRn must be the same as the suffix of the DAT-off module's entry point, IEAVEURn. See *System Macros and Facilities* for more detail on coding DATOFF macros.

3. Link edit the DAT-off module (IEAVEURn) into the IEAVEDAT member of SYS1.NUCLEUS (the DAT-off nucleus).

When the DATOFF macro executes, it branches to a routine in the PSA. The routine turns DAT off and branches to entry point IEAVEURn in IEAVEDAT. The DAT-off module returns via a PSA routine that turns DAT back on.

Cross Memory Entry Table Entries

You might have to change entry table entries that your installation created. MVS/XA uses a previously reserved bit in cross memory entry table entries to determine the addressing mode in which to enter the program. Entries that require modification are those that specify program addresses and either:

- Use bits in the entry description that are reserved in MVS/370
 - Specify programs to be entered in 31-bit addressing mode.
-

Interfaces to System Services

Some system services are changed to execute in 31-bit addressing mode. Some can now accept callers in either mode, but have restrictions on the length or value of parameter addresses. Others are restricted to using MVS/370-supported interfaces. During the migration phase, most programmers do not have to be concerned about changes to system service interfaces. Most programs that execute in 24-bit addressing mode and invoke system services via an SVC or a macro instruction continue to work unchanged in MVS/XA. (Exceptions are noted in the introduction to this chapter.) Interface changes might, however, affect existing programs that invoke service routines directly instead of via an SVC or macro instruction.

When modifying or developing programs that invoke system services directly or that execute in 31-bit addressing mode, programmers must now consider:

- The mode of the caller.
- The desired mode of the routine being called.
- The location of data areas passed to the service routines. Some data areas, such as the DCB, cannot reside in virtual storage above 16 Mb.
- The location of routines whose addresses are passed as parameters.
- The length of the address parameter fields. Some services expect parameter address fields to be 31 bits long even though the addresses contained in the fields might point to locations in virtual storage below 16 Mb. Other services use parameter fields that must be 24 bits long (for example the DCB address in an OPEN parameter list).

Programmers need to refer to the publications documenting the macros and SVC interfaces when using system services in 31-bit addressing mode or when invoking them directly.

System services can be categorized according to their interface requirements. Following are descriptions of the categories and examples of service routines in each. The list is not comprehensive.

Services Independent of Addressing Mode

Service routines in this category:

- Accept callers in either 24- or 31-bit addressing mode.
- Use 31-bit parameter address fields and, for callers in 31-bit mode, allow the addresses contained in those fields to point to any location.

EXAMPLES:

ABEND	EXIT	RESTORE
ATTACH*	FESTAE	SCHEDULE
CALLRTM	FREEMAIN (SVC 120)	SDUMP
CHAP	GET (VSAM)	SETFRR
CIRB	GETMAIN (SVC 120)	SETLOCK
CMSET	GETSRB	SETRP
CPOOL	GTRACE	SNAP*
DATOFF	GQSCAN	STATUS
DELETE	HOOK	STIMER
DEQ	IDENTIFY	SYNCH
DETACH	LINK*	SYSEVENT
DOM	LOAD*	TESTAUTH
DYNALLOC	PGSER	TTIMER
ENQ	POST	WAIT
ESPIE	PTRACE	WTO
ESTAE	PUT (VSAM)	WTOR
EVENTS	RESERVE	XCTL*

*When a DCB parameter is specified, the DCB must reside in 24-bit addressable storage.

Services with Some Restrictions on the Address Parameter Values

Services in this category:

- Accept SVC callers in either 24- or 31-bit addressing mode.
- Might require that branch entry callers be in 24-bit addressing mode.
- Require that one or more parameter addresses point to locations below 16 Mb. In some cases, the length of an address field must be 24 bits. In other cases, the length of an address field must be 31 bits long, but the address contained in the field must be a 24-bit value.

EXAMPLES:

BLDCPOOL	GETLINE	PUTLINE
BLDVRP	GETCELL	QEDIT
CLOSE	GETMAIN (SVC 4 and 10)	RACDEF
DLVRP	MGCR	RACLIST
EVENTS	MODCB	RACHECK
EXCP	OPEN	RACINIT
EXCPVR	PGFIX	RACROUTE
EXTRACT	PGFREE	SHOWCM
FRACHECK	PGLOAD	SMFWTM
FREECELL	PGOUT	SMFEWTM
FREEMAIN (SVC 5 and 10)	PGRLSE	STACK
GENCB	PURGE	TESTCB

Services that Do Not Support 31-bit Addressing

Services in this category:

- Accept callers in 24-bit addressing mode only.
- Require that all parameter addresses point to storage below 16 Mb. Parameter lists (both in-line and remote), control blocks, buffers, and user exit routines must reside in virtual storage below 16 Mb.

EXAMPLES:

SPIE
STAE
SEGLD
SEGWT

Data management macro instructions for all DFP access methods except VSAM (specifically, SAM, PAM, DAM, and ISAM)

31-bit Addressing Considerations

A 370-XA system can treat instruction and data addresses as 24- or 31-bit values. A new concept, addressing mode, describes the size of addresses being used. The value of a bit in the PSW (the PSW A-mode bit) determines the addressing mode. If the bit is 0, the system treats addresses (except those returned from the LRA instruction) as 24-bit values. If the bit is 1, the system treats them as 31-bit values. Programs executing while the system is in 24-bit addressing mode can address up to 16 Mb of virtual storage. Programs executing in 31-bit mode can address up to two gigabytes (approximately 2 billion bytes) of virtual storage.

Impact of 31-bit Addressing on Programmers

During the migration phase, most programmers do not have to be concerned with addressing mode. Most existing user-written programs that use standard system interfaces run unchanged on an MVS/XA system in 24-bit mode.

Programmers need to be concerned about addressing mode only if they:

- Have existing user-written programs that access system control blocks that have been moved to virtual storage above 16 Mb. (Appendix B, “Control Block Changes” lists those control blocks.) Programs that run in 24-bit addressing mode must switch modes to access data above 16 Mb. The next topic describes ways of changing the addressing mode. “Retrieving Data from a Control Block Above 16 Mb” illustrates how 24-bit mode programs can be changed to reference virtual storage above 16 Mb.
- Have existing user-written programs that use non-standard interfaces to invoke system programs (for example, programs that branch enter system programs rather than use macros, SVCs, or documented entry points). Some system programs must now be entered in 31-bit addressing mode or using a BASSM instruction.

Also, some system programs now expect input addresses to be 31-bit values. Modules that run in 24-bit mode must ensure that the addresses they pass to programs in 31-bit mode do not contain flags or other data in the high-order byte, unless the 31-bit mode program ignores the first byte or sets it to zero.

See “Modifying Programs that Invoke Modules Above 16 Mb” for examples of how you can make affected programs work in MVS/XA.

- Develop application programs, exit routines, or system modifications that execute in 31-bit addressing mode. Developing new programs to execute in 31-bit addressing mode is not described in this publication. See *SPL: 31-Bit Addressing*.

The following address mode related topics give programmers an introduction to how mode switching is performed so they can assess the work required to modify existing programs:

- “Changing Addressing Mode”
- “Establishing a Program’s Addressing Mode”
- “BSM (Branch and Set Mode) Instruction”
- “BASSM (Branch and Save and Set Mode) Instruction”
- “Modifying Programs that Invoke Modules Above 16 Mb”
- “Retrieving Data from a Control Block Above 16 Mb”
- “Performing I/O in 31-bit Addressing Mode”
- “Using the EXCP Macro”

See *SPL: 31-bit Addressing* for more detail.

Changing Addressing Mode

The only way to change the addressing mode is to change the value of the PSW A-mode bit. Following are ways of changing the A-mode bit:

- New 370-XA instructions:

BSM (branch and set mode)

BASSM (branch and save and set mode)

BSM and BASSM both save the current addressing mode, set a new addressing mode, and branch to an address. BASSM also saves a return address. The instructions allow problem programs in different addressing modes to communicate. See “New Instructions” for more detail.

- Supervisor assisted linkages (XCTL, LINK, and ATTACH). When a module uses XCTL, LINK, or ATTACH to invoke another routine, MVS/XA ensures that the called routine receives control in the correct addressing mode. (The way programs establish an addressing mode is described in the next topic.) Programs issuing XCTL, LINK, or ATTACH macros do not have to be aware of the addressing mode of the called routines except to ensure that the parameter requirements are met. When the routine called using LINK or ATTACH returns, the supervisor restores the addressing mode of the caller.
- Supervisor calls (SVCs). The supervisor saves and restores the issuer’s addressing mode and ensures that the service routine receives control in the correct mode.

Programs that reside below 16 Mb and pass parameters located in virtual storage below 16 Mb can issue SVCs without being aware of the service routine’s addressing mode or input requirements. However, before using SVCs in programs that will execute in 31-bit mode and/or use parameters located above 16 Mb, consult documentation on the SVC interface. Some SVCs require that input parameters be located below 16 Mb. See “Interfaces to System Services” for more detail.

- SYNCH macro. A new parameter, AMODE, allows programs to specify the addressing mode in which the called routine is to get control.
- SRB dispatch. When the SRB is dispatched, MVS/XA replaces the PSW A-mode bit with the high-order bit of the SRBEP or SRBRMTR field.
- PC and PT instructions, which establish the identified addressing mode.
- LPSW instruction.

Establishing a Program’s Addressing Mode

Every program that executes in MVS/XA is assigned two new attributes, an AMODE (addressing mode) and an RMODE (residency mode). (Existing programs are assigned default AMODE/RMODE attributes, which are described below.) AMODE specifies the addressing mode in which the program is designed to receive control. Generally, the program is also designed to execute in that mode, although a program can switch modes and can have different AMODE attributes for different entry points within a load module. The RMODE indicates where in virtual storage the program can reside.

Valid AMODE and RMODE specifications are:

AMODE=24	Specifies 24-bit addressing mode
AMODE=31	Specifies 31-bit addressing mode
AMODE=ANY	Specifies either 24- or 31-bit addressing mode
RMODE=24	Indicates that the module must reside in virtual storage below 16 Mb. You can use the RMODE=24 specification for 31-bit programs that have 24-bit dependencies.
RMODE=ANY	Indicates that the module can reside anywhere in virtual storage

You do not have to specify AMODE and RMODE attributes for a program. When none are specified, the system assigns the following defaults: AMODE=24, RMODE=24. To override the defaults, specify AMODE and/or RMODE on one or more of the following:

- AMODE and RMODE statements within the assembler source code for a program. Only Assembler H Version 2 recognizes AMODE and RMODE statements.

```
XYZ    CSECT
XYZ    AMODE xxx
XYZ    RMODE xxx
```

- The EXEC statement of a link edit step:

```
//LKED EXEC PGM=HEWLHO96, PARM=' AMODE=xxx, RMODE=xxx,
```

- The linkage editor MODE control statement (one per load module):

```
MODE AMODE (xxx) , RMODE (xxx)
```

- The LINK or LOADGO TSO commands:

```
LINK AMODE (xxx) , RMODE (xxx)
LOADGO AMODE (xxx) , RMODE (xxx)
```

Each of the AMODE/RMODE specifications listed overrides the previous ones in the list. For example, the AMODE/RMODE specifications on the MODE control statement override the specifications on the linkage editor EXEC statement.

MVS/XA uses a program's AMODE attribute to determine whether a program invoked using ATTACH, LINK, or XCTL is to receive control in 24- or 31-bit addressing mode. MVS/XA uses the RMODE attribute to determine whether a program must be loaded into virtual storage below 16 Mb or can reside anywhere in virtual storage (above or below 16 Mb).

Assembler H Version 2 establishes flags in the external symbol dictionary (ESD) to indicate the specified (or default) AMODE and RMODE of each CSECT. The MVS/XA linkage editor retains these flags in the composite external symbol dictionary (CESD). The linkage editor also inserts AMODE and RMODE flags in the partitioned data set (PDS) directory entry for each load module. The linkage editor by default uses the AMODE and RMODE indicators from the ESD entries. As noted earlier, the linkage editor also accepts AMODE and RMODE specifications in the EXEC statement and in the MODE control statement. If either of these are used to specify AMODE or RMODE, they are reflected in the PDS directory entry, but do not affect the information in the CESD.

You can use the MVS/XA version of AMBLIST to print the directory entry and the CESD to determine a program's AMODE and RMODE. You can also use the LOAD macro to determine the addressing mode in which a module expects to receive control. The high order bit of the entry point address that LOAD returns indicates the addressing mode.

Note: Do not confuse AMODE with the current addressing mode. Specifying an AMODE attribute guarantees that the module will receive control in the specified mode **only** when invoked using one of the methods defined in this topic. Specifying an AMODE does not, for example, prevent a program in 24-bit mode from issuing a BALR to a program with an AMODE of 31, although the program will not execute as expected. Also, there is nothing to prevent a programmer from specifying an AMODE of 31 on a program designed to execute in 24-bit mode, although doing so is incorrect.

Restrictions on Using a Linkage Editor Overlay Structure

Programs executing in 31-bit addressing mode cannot use a linkage editor overlay structure.


Changed Instructions

The following instructions work differently either when executed in 31-bit addressing mode or when executed on a 370-XA processor: BAL, BALR, BAS, BASR, CLCL, EDMK, LA, LRA, MVCL, and TRT. The following topics describe the differences.

Also remember that when executing in 31-bit mode, 370-XA processors treat all virtual addresses as 31-bit values. When executing in 24-bit mode, they treat virtual addresses as 24-bit values.


BAL and BALR (Branch and Link) Instructions

The way BAL and BALR work depends on the addressing mode. In 24-bit addressing mode, BAL and BALR work the same way as they do when executed on a 370 processor. BAL and BALR put information from the PSW into the high-order byte of the first operand register and put the return address into the remaining 3 bytes before branching:



ILC - instruction length code
 CC - condition code
 PGM MASK - program mask

In 31-bit addressing mode, BAL and BALR put the return address into bits 1 through 31 of the first operand and save the current addressing mode in the high-order bit. Because the addressing mode is 31-bit, the high-order bit is always a 1.



Note that when executed in 31-bit addressing mode, BAL and BALR do not save the instruction length code, the condition code, or the program mask. A new 370-XA instruction, IPM (Insert Program Mask), saves the program mask and the condition code.

BAS and BASR (Branch and Save) Instructions

BAS and BASR instructions execute on 308x processors in either 370 or 370-XA mode. They:

- Save the return address and the current addressing mode in the first operand.
- Replace the PSW instruction address with the branch address.

The high-order bit of the return register indicates the addressing mode.

Note that BAS and BASR perform the same function that BAL and BALR perform when BAL and BALR execute in 31-bit addressing mode.


CLCL, EDMK, MVCL, and TRT Instructions

When executed in 31-bit addressing mode, the following four instructions load 31-bit values into return registers and leave bit 0 unchanged. When executed in 24-bit addressing mode, they load 24-bit addresses and leave bits 0-7 unchanged:


- CLCL (Compare Logical Long)
- EDMK (Edit And Mark)
- MVCL (Move Long)
- TRT (Translate And Test)

Most programs using these instructions will run unchanged in 31-bit addressing mode. You need to change only programs that depend on bits 1-7 remaining unchanged.

Return Registers in 24-bit Addressing Mode



Return Registers in 31-bit Addressing Mode



LA (Load Address) Instruction

The LA instruction works differently when executed in 31-bit addressing mode. It loads a 31-bit value and clears the high-order bit instead of the entire high-order byte.

LRA (Load Real Address) Instruction

The LRA instruction performs the same functions as in 370 mode. However, it always returns a 31-bit address regardless of the issuing program's addressing mode. Also, the meaning of condition code 1 (CC=1) from an LRA instruction might be different. Because some page tables for the user region above 16 Mb are themselves pageable in MVS/XA, a condition code of 1 can mean either that:

- The page table does not exist because the virtual space has not been obtained.
- The page table is paged out or has not yet been built.

In MVS/XA, neither situation is necessarily an error. Users receiving a condition code of 1 can access the page in question to determine which condition exists. Accessing the page causes either an x'0C4' program check or segment fault/page fault resolution.

New Instructions

Standard instructions that are new in 370-XA mode include:

BASSM (Branch and Save and Set Mode)
BSM (Branch and Set Mode)
DXR (Divide Extended)
IPM (Insert Program Mask)
TRACE
All I/O instructions

The following topics describe BSM and BASSM in more detail. For more information on the other instructions, see *Principles of Operation*.

BSM (Branch and Set Mode) Instruction

BSM is a branch instruction that also sets the addressing mode.

The BSM instruction:

- Saves the current addressing mode. BSM puts into the high-order bit of the first operand the value of the current PSW A-mode bit. The rest of the operand is unchanged.
- Sets a new addressing mode. BSM replaces the PSW A-mode bit with the high-order bit of the second operand.
- Replaces the PSW instruction address with the branch address. Note that BSM sets the new addressing mode before computing the branch address. Thus, a program executing in 24-bit mode can use BSM to branch to a program in virtual storage above 16 Mb.

Uses for BSM

Programs called via a BASSM instruction (described in the next topic) can use BSM to return to the caller in the caller's addressing mode.

When the first operand is 0 (for example, BSM 0,14), BSM:

- Does NOT save the current addressing mode
- Sets the PSW AMODE bit
- Executes a branch

When the second operand is 0 (for example, BSM 14,0), BSM:

- Saves the current addressing mode
- Does NOT change the PSW AMODE bit
- Does NOT execute a branch

BASSM (Branch and Save and Set Mode) Instruction

The BASSM instruction works like BSM, except that it saves the return address as well as the current addressing mode in the first operand.

BASSM:

- Saves the next instruction address in bits 1 through 31 of the first operand.
- Saves the current addressing mode in the high-order bit of the first operand.
- Replaces the PSW A-mode bit with the high-order bit of the second operand.
- Replaces the PSW instruction address with a branch address. Like BSM, BASSM sets the new addressing mode before computing the branch address. Thus, programs executing in 24-bit mode can use BASSM to call programs in virtual storage above 16 Mb.

Uses of BASSM

Programs can use BASSM when branching to modules that might have different addressing modes. In addition, a program called via BASSM can return to its caller in the caller's addressing mode using either BSM or BASSM, provided the called program saves the full contents of the linkage register.

When the second operand is 0 (for example, BASSM 14,0), BASSM:

- Saves the current addressing mode
- Saves the next instruction address
- Does NOT change the PSW AMODE bit
- Does NOT execute a branch

When the first operand is 0 (for example, BASSM 0,14), BASSM does not suppress the saving operation.

Modifying Programs that Invoke Modules Above 16 Mb

You must change existing user-written programs that branch to MVS/XA system programs that must be entered in 31-bit addressing mode or via a BASSM instruction. Following are two examples of how you can adapt such programs work in MVS/XA. One method requires that you change user-written programs to use BASSM and BSM instructions. The other method uses a linkage assist routine, which does not require that you change your program.

To understand the following examples, you need to know:

- How BASSM and BSM work. Both are described in "New Branch Instructions."
- That the LOAD macro returns in the high order bit of Register 0 the addressing mode in which the module expects to receive control.

The modified program in the first example uses LOAD to determine not only the entry point address of the SYSRTN module, but also its addressing mode. The linkage assist routine in the second example uses LOAD to determine the same information for XYZNEW.

Using BASSM and BSM Instructions

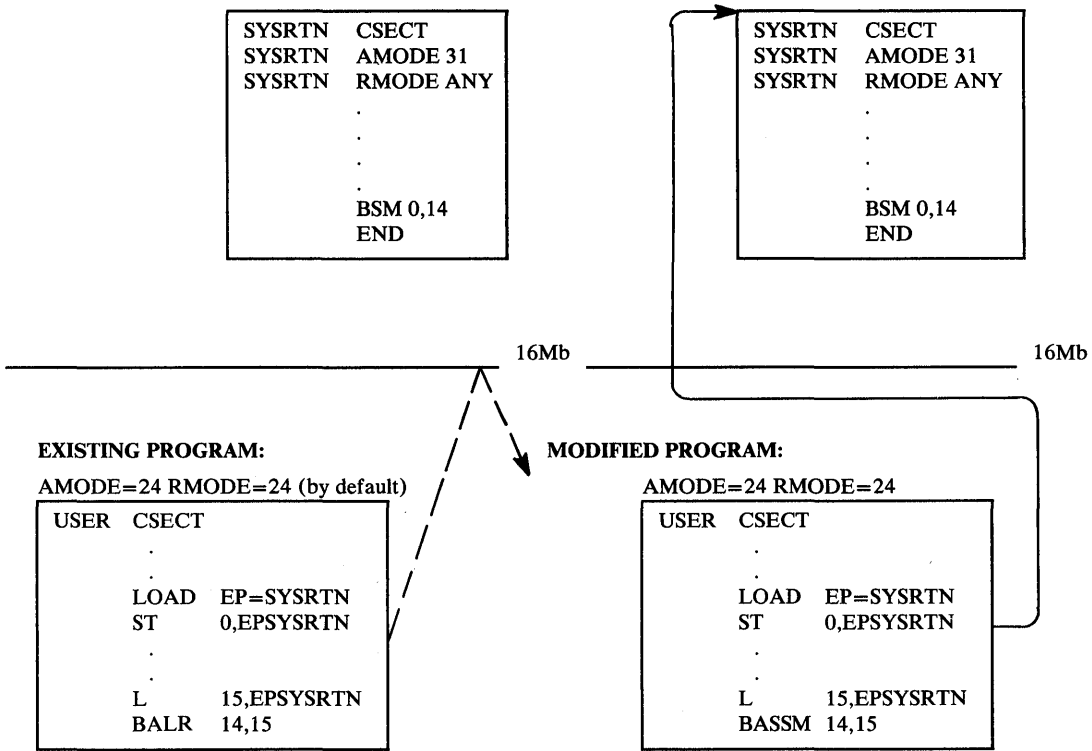
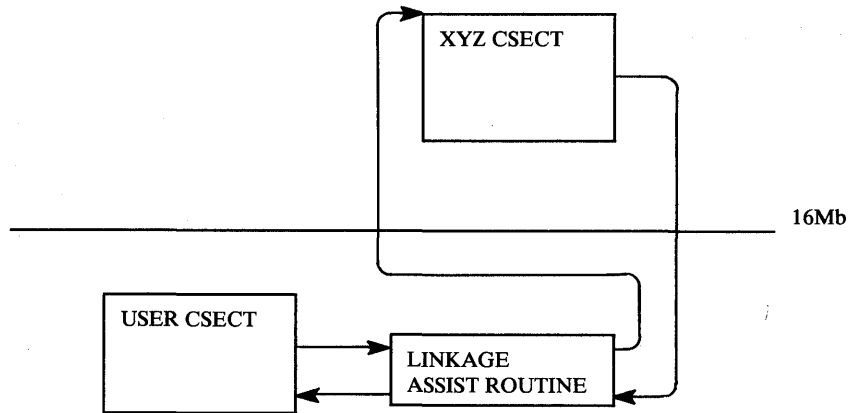


Figure 3-3. Example of Using BSM and BASSM

Using Linkage Assist Routines

Linkage assist routines are programs that handle the mode switching required to pass control between programs in 24- and 31-bit addressing mode:



The linkage assist routine should:

- Save registers.
- Ensure that all parameter addresses and linkage registers contain 31-bit address values.
- Obtain a new save area.
- Branch to the target module using a BASSM instruction.
- Receive control from the target module after it executes.
- Free the new save area.
- Restore registers.
- Return to the caller.

SPL: 31-bit Addressing provides guidelines for using linkage assist routines.

Linkage assist routines are mentioned in this publication for two reasons:

- You can use linkage assist routines to make user-written programs work in MVS/XA without change. Using a linkage assist routine is practical when:
 - A 24-bit mode program repeatedly calls a program that now executes in 31-bit mode.
 - Several 24-bit mode programs call the same program that now executes in 31-bit mode.

If a user-written program invokes a 31-bit mode program only once, other methods of mode switching might be preferable to using a linkage assist routine.

- Some system programs that have been moved above 16 Mb are invoked via linkage assist routines. This change is transparent to most users. Programmers might, however, notice linkage assist routines when tracing control flow during problem determination. Some addresses in the CVT now point to linkage assist routines instead of to the target module itself.

Following is one example of a linkage assist routine that passes control between a 24-bit mode user program and a 31-bit mode system program. Using the routine requires renaming the target routine and giving the linkage assist routine the original name of the target module.

OUXB



EXISTING PROGRAM

16Mb

USER	CSECT
.	.
.	.
L	15,ASCBOUXB
L	2,0,(15)
.	.

MODIFIED PROGRAM

USER	CSECT
.	.
L	15,ASCBOUXB
L	1,LABEL1
	BSM 0,1
LABEL1	DC A(LABEL2 + X'80000000')
LABEL2	DS 0H
	L 2,0,(15)
	LA 1,LABEL3
	BSM 0,1
LABEL3	DS 0H

(Put a 1 into the high-order bit of Register 1.)
(Switch to 31-bit addressing mode.)

(Retrieve the data from above 16Mb.)
(Put a 0 into the high order bit of Register 1.)
(Switch to 24-bit addressing mode.)

Figure 3-5. Retrieving Data from Above 16 Mb

Performing I/O in 31-bit Addressing Mode

To perform I/O, a program executing in 31-bit addressing mode must either:

- Use VSAM services, which accept callers in either 24- or 31-bit addressing mode. (See "Services with Some Restrictions on Address Parameter Values.") Programs using VSAM can access buffers that reside above 16 Mb.
- Use the EXCP macro. All parameter lists, control blocks, and EXCP appendage routines must reside in virtual storage below 16 Mb. See "Using the EXCP Macro."
- Use the EXCPVR macro. All parameter lists, control blocks, and appendage routines must reside in virtual storage below 16 Mb. See "Using the EXCPVR Macro Instruction."

- Use an intermediate routine that executes in 24-bit addressing mode as an interface to non-VSAM access methods, which accept callers executing in 24-bit addressing mode only.

To perform I/O to buffers located in virtual storage above 16 Mb, programs must use either:


- VSAM. Specify on the access method control block (ACB) at OPEN time that I/O buffers are to reside above 16 Mb. The ACB must be below 16 Mb, but the request parameter list (RPL) can be above 16 Mb.
- The EXCP macro and new virtual IDAW support, which “Using the EXCP Macro” describes.
- The EXCPVR macro. IDAWs can contain real addresses above 16 Mb, as described in “Using the EXCPVR Macro Instruction.”

Using the EXCP Macro

EXCP users can now:

- Back all I/O buffers with real storage above 16 Mb. To back I/O buffers below 16 Mb with real storage above 16 Mb, callers must specify LOC=(BELOW,ANY) on the GETMAIN request. (See “New Parameters on the GETMAIN Macro Instruction.”)
- Perform I/O to virtual storage areas above 16 Mb. CCWs in the channel program that EXCP initiates can point to a virtual IDAW. The virtual IDAW contains the 31-bit virtual address of an I/O buffer, which can be anywhere in virtual storage. The EXCP service routine supports only Format 0 CCWs, the CCW format used in MVS/370.

CCW (Format 0)



Any CCW that causes data to be transferred can point to a virtual IDAW. Virtual IDAW support is limited to non-VIO data sets. Programmers must be aware of this fact when coding the JCL to execute a program that uses virtual IDAWs.

Although the I/O buffer can be in virtual storage above 16 Mb, the virtual IDAW that contains the pointer to the I/O buffer and all other areas related to the I/O operation (CCWs, IOBs, DEBs, DCBs, and appendages) must have 24-bit virtual addresses.

See *SPL: Data Management* for information on using the EXCP macro.

Entry Points in IEFW21SD

The following entry points in the device allocation load module, IEFW21SD, are changed in Release 1.1. When writing programs to be executed in 31-bit addressing mode, use the new entry points. Programs that run in 24-bit addressing mode can continue to use the old entry points.

<i>Old Entry Point</i>	<i>New Entry Point</i>
IEFAB4DC	IEFGB4DC
IEFAB445	None
IEFAB4UV	IEFGB4UV

Summary of New and Changed Macros

Figure 3-6 lists macros that are new or that have new or changed options. When assembling programs that use any of the new function, use the MVS/XA MACLIB. With the following exceptions, the object code generated will be downward incompatible:

- The new LOC, VRC, and VRU options on the GETMAIN macro are downward compatible, as described in “New Parameters on the GETMAIN Macro Instruction.”
- The MVS/XA MACLIB expansions of SYNCH macros that specify AMODE=24 are downward compatible. However, if the AMODE parameter is omitted or if it specifies any option other than 24, the MVS/XA expansion of SYNCH will not run on an MVS/370 system. See “Downward Incompatible SYNCH Macros” in Chapter 9 for more information.

Macro	Release			Description of Change
	1.0	1.1	1.2	
ABEND	x			A new keyword, REASON, specifies a reason code that supplements the completion code for abnormal termination. The list form of the SNAP macro to which the DUMPOPT parameter points can specify all of the new SNAP dump options.
BLSABDPL			x	A new macro that maps the BLSABDPL exit parameter list. Dump analysis and formatting exit routines that run under IPCS, PRDMP, or SNAP can use BLSABDPL when invoking the new services described in "New Services for Dump Processing Exits" in Chapter 5.
BLSQMDEF and BLSQMFLD			x	A pair of macros that when used together define the structure of a control block. BLSQMDEF marks the beginning of the control block and defines the header. Several BLSQMFLD macros follow, each of which defines a single field in the control block. A second BLSQMDEF macro denotes the end of the control block. The definition is called a control block model. Dump analysis and formatting routines can use the models instead of format patterns for formatting control blocks. See "Format Model Processor Service" in Chapter 5 for more information.
BLSRESSY			x	A new macro that maps one IPCS symbol table record. Dump analysis and formatting routines that run under IPCS and use either the new GET or EQUATE service can use BLSRESSY to describe the record to be retrieved or stored in the symbol table. See "New Services for Dump Processing Exits" in Chapter 5 for more information about the GET and EQUATE services.
CALLRTM	x			A new keyword, REASON, specifies a reason code that supplements the completion code for abnormal termination. The list form of the SNAP macro to which the DUMPOPT parameter points can specify all of the new SNAP dump options.
CHKPT	x			New keywords: - IDADDR specifies the address of a checkpoint ID - IDLNG specifies the length of the checkpoint ID field - DDNADDR specifies the address of a DDNAME You can use IDADDR and IDLNG instead of the ID and LNG parameters.
CIRB	x			A new parameter, AMODE, specifies the addressing mode in which the specified asynchronous exit routine is to get control. If the AMODE parameter is not specified, the exit routine gets control in the issuer's addressing mode.
CPOOL	x			A new macro that: - Creates a cell pool as described by the requestor - Obtains or returns a cell to a previously constructed cell pool - Deletes a previously constructed cell pool Requestors require authorization only when the storage to be obtained is in an authorized GETMAIN subpool.
CPUTIMER	x			A new macro that obtains the current CPU (processor) timer clock value. CPUTIMER allows users to determine the amount of time remaining in a time interval established by an SRBTIMER or STIMER macro. CPUTIMER uses a PC instruction instead of an SVC. Therefore, it is faster than using TTIMER. Also, users can issue CPUTIMER in task or SRB mode. Issuers of TTIMER must be in task mode.
DATOFF	x			A new macro that provides linkage to DAT-off routines. See "DAT-off Restrictions."

Figure 3-6 (Part 1 of 5). Summary of New and Changed Macros

Macro	Release			Description of Change
	1.0	1.1	1.2	
ENQ		x		Users might receive a new return or abend code. In Release 1.1, global resource serialization limits the number of ENQ, RESERVE, and certain GQSCAN requests a single job, started task, or TSO user can have outstanding at a given time. If an address space reaches the limit, the system terminates unconditional ENQ requests with abend code x'538' and rejects conditional requests with a return code of x'014'. See "Limiting Concurrent Global Resource Serialization Requests" for more detail.
			x	ENQ has two new keywords, MASID and MTCB, which are designed for internal use. They are mentioned here only because you might see them in dumps or source code. MASID and MTCB specify the ASID and TCB address of a task, respectively. They are useful in situations where one task is performing work for another and might require resources that task owns. If an ENQ specifying MASID and MTCB fails because another task owns the resource, the task issuing the ENQ can determine whether the identified task is the owner. New return codes provide that information. If the specified task owns the resource, the issuing task can choose to use the resource anyway. Both tasks, however, must have previously established an alternate method of serialization.
ESPIE	x			A new macro that provides services similar to the SPIE macro for callers in either 24- or 31-bit addressing mode. SPIE users must be in 24-bit addressing mode.
GETMAIN	x			Three new parameters: VRC, VRU, and LOC. See "New Parameters on the GETMAIN Macro Instruction."
GTRACE	x			A new parameter, TEST, determines whether data gathering and tracing are to be done. See "Using GTF to Trace User Events." <i>SPL: Service Aids</i> describes GTRACE.
GQSCAN		x		Users might receive a new return code. In Release 1.1, global resource serialization limits the number of ENQ, RESERVE, and certain GQSCAN requests a single job, started task, or TSO user can have outstanding at a given time. If an address space reaches the limit, global resource serialization rejects subsequent GQSCAN requests that specify the TOKEN option and request more information than can fit into the caller's buffers. The issuer receives a return code of x'14.' Global resource serialization returns the buffers of information, but does not continue the scan. For more information, see "Limiting Concurrent Global Resource Serialization Requests."
			x	Release 1.2 allows you to specify generic (partial) rnames and qnames on the RESNAME keyword. GQSCAN attempts to obtain information about any resource that matches the specified part. The keywords and parameters that provide the new support are GENERIC, SPECIFIC, rname length, and qname length. GENERIC indicates that the qname and rname on GQSCAN are generic names. The qname and rname lengths specify the number of characters in the qname or rname that must match the qname and rname specified on GQSCAN. SPECIFIC indicates that the complete qname and rname must match. It is the default.
LOAD	x			New keywords: - EOM=NO YES specifies whether a module in global storage is to be deleted at address space termination (YES) or end-of-task time (NO). The default is NO. Previously, the system deleted the modules at end-of-task time. - LOADPT requests that the virtual address of the first byte of the load module be returned at the location indicated on the LOADPT keyword. The high-order bit of the entry point address that LOAD returns indicates the addressing mode in which the routine expects to receive control.
MGCR			x	With Release 1.2 installed, you can use MGCR to issue internal REPLY commands, as well as internal START commands. This new function enables you to write WTO/WTOR exits that respond to particular WTOR macros. <i>User Exits</i> shows an example of such an exit. (WTO/WTOR exits are also new in Release 1.2. See "New WTO/WTOR User Exits" for more information.)

Figure 3-6 (Part 2 of 5). Summary of New and Changed Macros

Macro	Release			Description of Change
	1.0	1.1	1.2	
NUCLKUP	x			A new macro that retrieves either: (a) the address and addressing mode of a CSECT or entry point in the DAT-on nucleus, or (b) the name and address of a CSECT that resides at a given address in the DAT-on nucleus.
PGSER	x			A new macro that performs the same services as PGANY, PGFIX, PGFIXA, PGFREE, PGFREEA, PGLoad, PGOUT, and PGRlse. PGSER can use 31-bit addresses, the other services listed cannot.
PTRACE	x			A new macro that creates system trace table entries.
RACROUTE	x			A new macro that requests the RACF services that FRACHECK, RACDEF, RACLST, RACHECK, and RACINIT invoke. The REQUEST parameter on RACROUTE indicates which service is to be performed. Any of the parameters that are valid on the other RACF macros are also valid on RACROUTE. Thus, using RACROUTE is very similar to using the other macros. Note one difference, however. RACROUTE requires a 512-byte work area, while the other RACF macros do not.
RESERVE		x		Users might receive a new return or abend code. In Release 1.1, global resource serialization limits the number of ENQ, RESERVE, and certain GQSCAN requests a single job, started task, or TSO user can have outstanding at a given time. If an address space reaches the limit, the system terminates unconditional RESERVE requests with abend code x'538', and rejects conditional requests with a return code of x'014'. See "Limiting Concurrent Global Resource Serialization Requests" for more detail.
			x	Release 1.2 adds two new keywords, MASID and MTCB. See the ENQ entry for more information.
RETURN		x		Release 1.1 changes the flag that the T parameter requests the system to set. The new flag is the low-order bit of the fourth word in the called program's save area. The system sets that bit to one after the called program has returned to its caller. In previous releases, the same flag is the entire fourth word in the save area.
SDUMP	x			New parameters: <ul style="list-style-type: none"> - ALLNUC, a new option on the SDATA keyword, requests that the DAT-off nucleus and the entire DAT-on nucleus be dumped. - SUBPLST and KEYLIST specify the subpool and key of storage to be dumped. - TYPE=NOLOCAL indicates that SDUMP is not to obtain a local lock. Changed parameters: <ul style="list-style-type: none"> - NUC requests that the DAT-on, non-page-protected section of the nucleus be dumped. In MVS/370, NUC causes the entire nucleus to be dumped. The ALLNUC option requests the entire MVS/XA nucleus. - TRT requests trace data from the active trace facilities, as in MVS/370. However, in MVS/XA, if the caller is unauthorized, the dump includes system trace data for the caller's address space only. MVS/XA dumps for authorized requestors and all MVS/370 dumps include the entire system trace table. Also, unlike MVS/370, MVS/XA dumps can include both system trace and GTF data, because both trace facilities can be active at the same time. If SDUMP specifies any new parameters, the macro must be assembled using the MVS/XA macro expansion. If it is not, the system ignores the new parameters and flags the others as errors.
			x	
SETLOCK	x			New parameters support the new lock types. Programs that issue SETLOCK RELEASE,TYPE=(reg) ALL must use the MVS/XA expansion of SETLOCK in some situations. See "SETLOCK RELEASE,TYPE=(reg) ALL Macro Instruction."

Figure 3-6 (Part 3 of 5). Summary of New and Changed Macros

Macro	Release			Description of Change
	1.0	1.1	1.2	
SETRP	x			<p>A new keyword, REASON, specifies a reason code that supplements the completion code for abnormal termination.</p> <p>The list form of the SNAP macro to which DUMPOPT points can specify all of the new SNAP dump options.</p>
SMFIOCNT	x			<p>A new macro that supplies to SMF either the EXCP count, the device connect time, or both.</p>
SNAP	x			<p>New parameters:</p> <ul style="list-style-type: none"> - SUM, a new option on the SDATA keyword, requests a new summary dump. See "New User Summary Dumps" in Chapter 6. - SUBPLST requests individual subpools. When SNAP specifies the SUBPLST option, the length of the list and standard forms of the macro expansion increase. The object code generated is downward incompatible. - ALLVNUC, a new option on the SDATA keyword, requests that the entire DAT-on nucleus be dumped. In SYSMDUMPS, ALLVNUC causes the entire nucleus to be dumped. - SUBTASKS, a new option on the PDATA keyword, requests that program data for all subtasks of a designated task be dumped. <p>Changed parameters:</p> <ul style="list-style-type: none"> - NUC requests that the DAT-on, non-page-protected section of the nucleus be dumped. SYSABEND dumps also include the PSA and CVT. In MVS/370, NUC causes the entire nucleus to be dumped. - TRT requests trace data from the active trace facilities, as in MVS/370. In MVS/XA, the dump can include system trace data and GTF data. In MVS/370, system trace and GTF cannot be active at the same time. Therefore, MVS/370 dumps never include both system trace and GTF data. - SDATA=ALL requests all of the SDATA options except ALLVNUC. In MVS/370, it includes all SDATA options.
SPLEVEL	x			<p>A new macro that controls which expansion of a macro the assembler generates. The MVS/XA MACLIB contains both MVS/XA and MVS/370 expansions of macros whose MVS/XA expansions do not work in MVS/370. See "Handling Downward Incompatible Macros" in Chapter 9.</p>
STIMERM			x	<p>A new macro that sets a timer for a real time interval, as does the existing STIMER macro. STIMERM is different from STIMER, however, in that:</p> <ul style="list-style-type: none"> - Each task can have up to 16 STIMERM intervals in effect at the same time. Only one STIMER interval is allowed. - STIMERM sets only real time intervals; STIMER sets both task and real time intervals. - STIMERM can also test how much time remains in the interval and cancel the interval. TTIMER provides those functions for STIMER intervals. - STIMERM can pass a 4-byte parameter to the exit routine that receives control when the interval expires. STIMER cannot. <p>STIMERM is provided so that a task can easily have a task interval and one or more real time intervals in effect at the same time. A task can, for example, set an STIMER interval to measure task time and an STIMERM interval to simultaneously measure real time.</p>

Figure 3-6 (Part 4 of 5). Summary of New and Changed Macros

Macro	Release			Description of Change
	1.0	1.1	1.2	
SVCUPDTE	x			A new macro that allows authorized programs to dynamically update the SVC table. In MVS/370, defining an SVC after sysgen requires an IPL of the system.
SYNCH	x			A new parameter, AMODE, indicates the addressing mode in which the specified program is to get control. If the AMODE parameter is not specified, the program gets control in the issuer's addressing mode. Unless SYNCH specifies the AMODE=24 parameter, programs that use SYNCH and are assembled using the MVS/XA MACLIB will not run on MVS/370 systems. See "Downward Incompatible SYNCH Macros" in Chapter 9 for more information.
VSMLIST	x			A new macro that returns information about virtual storage allocation within an address space.
VSMLOC	x			A new macro that verifies a given virtual storage area has been allocated to satisfy a GETMAIN request.
VSMREGN	x			A new macro that returns the starting addresses and sizes of the private area regions associated with a given TCB in the current address space.
WTL			x	A new keyword, OPTION=PREFIX or OPTION=NOPREFIX, indicates whether the WTL text contains a prefix to identify the log record. If the text already contains a prefix, specify the PREFIX option. If you specify NOPREFIX or omit the OPTION keyword altogether, the system inserts a two-character prefix. ('X' is the default prefix.)
WTO		x		You can specify routing code 11 (programmer information) on multiline WTO messages, an option earlier releases do not allow. If, because of that restriction, you have programs that issue a series of single line WTOs with the same message ID, you can improve performance by combining the messages into one multiline message. You can also use the HRDCOPY option when writing multiline WTOs. In earlier releases you cannot write multiline messages to hardcopy only.
			x	A new CMD option on the MCSFLAG keyword enables you to record system commands in the system log. MCSFLAG=CMD indicates that the text is a system command and requests that it be entered in the system log.
WTOR			x	Release 1.2 adds the CMD option on the MCSFLAG keyword. See the WTO entry.

Figure 3-6 (Part 5 of 5). Summary of New and Changed Macros

New Parameters on the GETMAIN Macro Instruction

MVS/XA provides three new parameters on GETMAIN: VRC, VRU, and LOC.

VRC and VRU Parameters

VRC (variable request conditional) and VRU (variable request unconditional) are two new forms of GETMAIN. Both issue SVC 120.

```
GETMAIN VRC,LV=(maximum length,minimum length)
GETMAIN VRU,LV=(maximum length,minimum length)
```

VRC and VRU request a single area of virtual storage having a length between the maximum and minimum lengths specified. MVS/XA returns the address of the allocated virtual storage in Register 1 and the length of the storage in Register 0.

Callers in 24- or 31-bit addressing mode can use VRC or VRU. However, MVS/XA treats all parameter lengths and addresses as 31-bit values.

VRU and VRC are exceptions to the general rule that programs using new MVS/XA function are not downward compatible. Both generate object code that runs on MVS/370 systems. MVS/370 treats VRC and VRU parameters as RC

and RU, respectively, and obtains the maximum length of storage specified on the LV operand. MVS/370, of course, also uses 24-bit parameter values.

LOC Parameter

The new LOC parameter has two subparameters for specifying whether virtual storage is to be obtained above or below 16 Mb and how it is to be backed if fixed (below 16 Mb or anywhere). (RSM always allocates real storage anywhere until the storage is fixed.) Possible LOC specifications are:

LOC=(BELOW) VSM must allocate virtual storage below 16 Mb.

LOC=(ANY) VSM can allocate virtual storage anywhere.

LOC=(RES) VSM allocates storage according to the requestor's residence. If the requestor resides in virtual storage below 16 Mb, VSM allocates storage below 16 Mb. If the requestor resides above 16 Mb, VSM allocates storage anywhere.

LOC=(parm1,ANY) RSM attempts to back the page with real storage above 16 Mb. If unsuccessful, RSM backs the page with real storage below 16 Mb. Note that, regardless of the LOC specification, RSM backs virtual storage with real storage anywhere until the storage is fixed (either by definition or by a PGFIX or PGSER macro).

The first subparameter (parm1) can be BELOW, ANY, or RES.

LOC is especially useful in programs with 24-bit dependencies. Programs that reside above 16 Mb must specify LOC=(BELOW) on requests for storage that has 24-bit dependencies.

You can specify LOC only on the RU, RC, VRU, and VRC forms of GETMAIN (SVC 120). VSM satisfies all other forms with virtual storage below 16 Mb, which RSM backs with real storage below 16 Mb.

Like VRU and VRC, LOC is downward compatible. Regardless of the LOC specification, however, MVS/370 always obtains storage below 16 Mb.

SDUMP Macro Instruction

If you use any new SDUMP keywords, be aware that SDUMP generates a longer parameter list. In addition, once the longer list is assembled in a module, the assembler generates the long form of all subsequent SDUMP parameter lists in the module, regardless of which keywords the SDUMP macros specify. The long list is similar to the short list, except that it has additional bytes appended to the end.

If you do not want the long form to be generated on subsequent macros, use the SPLEVEL macro to request the MVS/370 expansion. See "Handling Downward Incompatible Macros" in Chapter 9.

SETLOCK RELEASE,TYPE=(reg) | ALL Macro Instruction

If any new locks are held when the SETLOCK RELEASE,TYPE=ALL macro is issued, you must use the MVS/XA expansion. The MVS/370 expansion does not recognize the new locks and, therefore, does not release them. Likewise, if the SETLOCK RELEASE,TYPE=(reg) macro is issued and the bit string in the register specifies a new lock, you must use the MVS/XA expansion. If you use the MVS/370 expansion, the system does not release the new locks. You can use an SPLEVEL macro to request either expansion. See "Handling Downward Incompatible Macros" in Chapter 9.

Using GTF to Trace User Events

MVS/XA provides an alternate way of using GTF to trace USR events. Using the new method:

- Applications no longer need to supply and support an external interface that requests tracing. Starting GTF with the appropriate USRP options specified is sufficient to allow applications to trace their own USR events. (The USRP option is new in MVS/XA. See the TRACE entry in Figure 4-2 in Chapter 4.)

In MVS/370, applications have to provide their own interface for requesting USR event tracing. One example of an interface is the DIAGNS=TRACE subparameter of the DCB parameter on a DD statement, which requests module flow tracing through OPEN, CLOSE, and EOV. Also, any program can support and request tracing of their own USR events by specifying a trace keyword in the PARM field of the EXEC statement.

- Applications can use the new TEST keyword on the GTRACE macro to determine whether or not GTF tracing is active for their USR events. Depending on the return code from GTRACE, applications can either gather the trace data and have it written or bypass tracing.

In MVS/370, applications have to make two tests before issuing GTRACE, one to determine if the application has requested tracing, and another to determine if GTF is active for USR tracing.

The MVS/370 method of establishing tracing capability continues to work in MVS/XA. GTRACE TEST offers an alternate method for MVS/XA users.

The way that applications build the data records to be traced and issue the GTRACE macro to write them to the SYS1.TRACE data set has not changed.

See *SPL: Service Aids* for more information on using GTRACE.

Unit Verification

Two modules in MVS/XA provide unit verification: IEFAB4UV for programs in scheduler key (key 1) and IEFEB4UV for programs in user key (key 8-15) and task mode. The MVS/370 device allocation tables (DEVNAMET, IEFDEVPT, and DEVMASKT) and module IEFSCAN have been deleted. You must change programs that call IEFSCAN or that directly access the device allocation tables to use either IEFAB4UV or IEFEB4UV.

IEFAB4UV

The MVS/XA version of IEFAB4UV provides all of the services that the MVS/370 version provides, all of the IEFSCAN services, and some new function. Specifically, IEFAB4UV can:

1. Check whether the device numbers supplied as input are all associated with the same group.
2. Check whether the device numbers supplied as input are associated with the unit name specified in the eligible device table (EDT).

3. Return the unit name associated with an input value such as a device type. The unit name is the EBCDIC representation of the IBM generic device name (for example, 2305) or the user-defined esoteric name (for example, TAPE).
4. Return the UCB addresses associated with a specified unit name.
5. Return group identification for each input UCB.
6. Indicate whether a specified unit name is an internal representation of the unit name (that is, whether the unit name is an index into the EDT). This service is used with 2 and 4.
7. Return the internal representation of a specified unit name, which can then be used as an index into the EDT.
8. Convert a four-byte UCB device type to an internal representation of a unit name, which can then be used as an index into the EDT.
9. Return general information about a specified unit name, including:
 - Whether the unit name is esoteric, VIO eligible, contains 3330V units, or contains teleprocessing class devices
 - The number of device classes in the unit name
 - The number of generic device types in the unit name
10. Indicate that the parameter list should not be altered, thereby allowing the parameter list to be in storage that is not protected by key 1. This service is used with 2.

IEFEB4UV

IEFEB4UV performs the functions described in 3, 4, 6, 7, 8, and 9 of the IEFAB4UV description for programs in user key and task mode.

See *SPL: System Modifications* for information on using IEFAB4UV or IEFEB4UV.

Chapter 4. Operating Considerations

This chapter contains information that pertains to operators and operational procedures. System programmers might also be interested in some of the command changes it describes (for example, SLIP, DISPLAY DUMP, and DUMP). The topics included in this chapter are:

- “Loading 370-XA Microcode at Power-on Reset Time”
 - “SYSCTL (SCP Manual CNTL) Console Frame”
 - “Storing Status Before Taking a Stand-alone Dump” on page 4-3
 - “Using Labeled Tapes for Stand-alone Dumps” on page 4-3
 - “JCL Changes to Jobs that Allocate SYS1.DUMP Data Sets” on page 4-3
 - “Processing Hot I/O Interrupts” on page 4-3
 - “Extended Color Support on 3279 MCS Consoles” on page 4-4
 - “Controlling Message Traffic on Operator Consoles” on page 4-5
 - “New Response to Message IOS201E”
 - “Summary of New, Changed, or Deleted Commands”
-

Loading 370-XA Microcode at Power-on Reset Time

To initialize a 308x processor to run in 370-XA mode, the operator must use the CONFIG frame to:

1. Release the current configuration (option A=1)
2. Once the current configuration is released, select 370-XA mode (option M=1)
3. Select the processors and storage that are to be online at power-on-reset time (P and S options)
4. Perform the power-on-reset function (option A=2)

The configuration frame is unchanged, except that it has a new option, M=1, for selecting 370-XA mode. Previously, however, the frame appeared with some options already selected. Releasing the current configuration might be a new step for some operators.

SYSCTL (SCP Manual CNTL) Console Frame

An SYSCTL console frame comes with 308x processors. The SYSCTL frame is very similar to the SC frame available on 3036 consoles attached to 3033 processors. The frame allows the operator to request some of the functions offered

on the existing OPRCTL (operator control) frame, as well as some additional functions. The SYSCTL frame allows the operator to:

- Specify an alternate nucleus. Operators should use the SYSCTL frame instead of previous methods for specifying an alternate nucleus. Previous methods include:
 - Using the OPRCTL frame or the SC frame.
 - Storing an alternate nucleus identifier in absolute location 8 after the hardware IPL completes and while the processor is in instruction step mode. VM users must still use this method.
 - Specifying 'ALT=xx' when asked to specify system parameters at IPL time. MVS/XA does not support the ALT parameter.
- Specify the device number from which to IPL the operating system.
- Specify the device number from which to IPL the stand-alone dump program. Having the stand-alone dump IPL option separate from the operating system IPL option prevents the operator from inadvertently loading the wrong program. In addition, issuing the stand-alone dump IPL command from an SYSCTL frame causes the hardware to automatically store status.
- Specify how MVS/XA is to perform restart processing. The operator can request that MVS/XA:
 - Option 0 - Display information about the work in progress. The operator can then choose to either terminate the interrupted work and invoke the necessary recovery routines, or return to the point of interruption.
 - Option 1 - Attempt to diagnose and repair the problem. In response, the system takes several corrective actions.

With an OPRCTL frame, the operator does not have an option. MVS performs restart processing as described in option 0 for the SYSCTL frame.

When performing restart processing, operators should use the SYSCTL frame. If the system is in a restartable wait state, operators should either:

- Select option 0 on the SYSCTL frame. Operators must not select any other option.
- Use the restart button. Using the restart button is a valid option only if the operator knows which processor is the target of the restart. The current frame might not identify the target processor, in which case the hardware uses a previously-established target.
- Use a bottom line command. The bottom line command allows the user to specify a target processor and can be used with any frame.
- Use the OPRCTL frame.

- Request instruction recording. The operator can have the hardware record a total of 470 instruction addresses on a disk in the service processor. The operator can obtain a hardcopy using stand-alone dump or SVC dump formatted by print dump. Previously, when a program was looping, the operator had to record by hand the instruction counter addresses before taking a dump.
- Allow instruction stepping on both processors.

See the *Operator's Guide* for the 308x processor for more information on using the SYSCCTL frame.

Storing Status Before Taking a Stand-alone Dump

If the operator IPLs stand-alone dump using an SYSCCTL (SCP manual CNTL) frame, the hardware automatically stores status. The operator is not required to store status manually.

If the operator uses an OPRCTL (operator control) frame to IPL stand-alone dump, the operator must store status manually. If the operator does not store status, the stand-alone dump might be missing critical information.

Using Labeled Tapes for Stand-alone Dumps

Stand-alone dump (SADMP) can use labeled tapes that are not password-protected. If the operator mounts such a tape, SADMP displays the VOLSER and asks the operator if the tape is to be used. Note, however, that **SADMP writes over the label**. If used again as a labeled tape, the tape has to be re-labeled.

In MVS/370, SADMP rejects all labeled tapes.

JCL Changes to Jobs that Allocate SYS1.DUMP Data Sets

Jobs that allocate SYS1.DUMPnn data sets (for example, AMDPRDMP or IEBCGENER to unload dump data sets) must specify DISP=SHR on the JCL. You must change DISP=OLD to DISP=SHR on any DD statements that allocate SYS1.DUMPnn data sets.

MVS/XA now allocates dump data sets to the DUMPSRV address space to improve integrity. The data sets are allocated with DISP=SHR and DUMPSRV does not free them after taking a dump. Thus, jobs that request SYS1.DUMPnn data sets with DISP=OLD cannot access them.

Processing Hot I/O Interrupts

Hot I/O interrupt processing is changed. (Hot I/O interrupts are consecutive, unsolicited interrupts on a subchannel. They are caused by hardware malfunctions.) In MVS/XA:

- IOS uses a single criterion for detecting hot I/O conditions. Because of the new I/O architecture, no other thresholds are necessary. In MVS/370, IOS uses separate thresholds to detect excessive time-outs and hot I/O conditions on channels, devices, and control units.

- Installations can specify hot I/O recovery actions that IOS performs automatically. Unlike MVS/370, recovery does not have to involve the operator.

You can specify recovery actions for the following device categories:

- Reserved DASD
- Non-reserved DASD
- All other devices

Valid recovery actions are:

- Asking the operator to direct recovery (as in MVS/370)
- Boxing the hot device (forcing the device offline in such a way that future I/O requests for the device are returned to the I/O driver as permanent errors)
- Performing channel path recovery
- Forcing the channel path offline

The following figure shows the IBM-supplied recovery actions:

Device Category	Non-recursive Condition	Recursive Condition
RESERVED DASD	Request direction from the operator	Request direction from the operator
NON-RESERVED DASD	Perform channel path recovery	Request direction from the operator
ALL OTHER	Request direction from the operator	Request direction from the operator

Figure 4-1. Default Hot I/O Recovery Actions

The hot I/O threshold and recovery actions are contained in a new module, IOSRHIDT, which replaces IECVHIDT. If Release 1.1 or a later release is installed, you can change the defaults using the new HOTIO keyword in the IECIOSxx PARMLIB member. If the system is at the Release 1.0 level, use the AMASPZAP service aid instead. *SPL: System Modifications* describes how to change the defaults in more detail.

Extended Color Support on 3279 MCS Consoles

MVS/XA provides additional ways of highlighting messages on 3279 MCS color consoles, Models 2B and 3B. You can:

- Display message types or console fields in up to seven different colors. (MVS/SP Version 1 Release 3 provides four-color support for 3279 MCS color consoles. Four-color support is still available on Models 2A, 2C, and 3A.)
- Highlight messages with underscoring, blinking, or reverse video display (black characters on a colored background).

Color-coding and other highlighting techniques help operators distinguish the importance of messages. As such, highlighting can be an effective way of controlling message traffic.

You can specify highlighting attributes for your installation in an MPFLSTxx PARMLIB member. If none are specified, the system uses default values. The same highlighting attributes are in effect for all 3279 consoles, Models 2B and 3B. A SET MPF=xx command, where xx identifies an MPFLSTxx PARMLIB member, causes the system to use the attributes in the specified MPFLSTxx member. When highlighting attributes are changed, the system puts the name of the old MPFLSTxx member in the SET MPF command section of a type 90 SMF record. The DISPLAY MPF command displays the current specifications. *SPL: Initialization and Tuning* describes how to use MPFLSTxx.

Controlling Message Traffic on Operator Consoles

In general, as the system workload increases, messages appear on the operator console at a faster rate. To keep the message rate manageable, your installation needs to evaluate its current methods of tailoring message output. Methods of controlling message traffic include using:

- The WTO/WTOR user exits, which are new in Release 1.2, or the existing WTO user exit (IEECVXIT). The new WTO/WTOR exits can modify processing for any message. They can also alter processing in more ways than IEECVXIT can. See “New WTO/WTOR User Exits” in Chapter 5 for more information.
- Additional operator consoles with multiple console support (MCS). *System Commands* describes how to use MCS consoles.
- The message processing facility, which suppresses nonessential messages from the operator console. *Initialization and Tuning* describes how to use the message processing facility.
- Console clusters, which reduce message traffic on a single console. *System Commands* describes how to use console clusters.
- The TRACK command to display system status, instead of having JOB STARTED/ENDED messages displayed on the console. *System Commands* describes the TRACK command.
- Message routing codes to direct application messages to the appropriate console. *System Commands* describes how to assign routing codes.
- Color displays to help operators distinguish important messages. Four-color message coding is available on 3279 consoles, Models 2A, 2C, and 3A; seven-color message coding is available on 3279 consoles, Models 2B and 3B. *System Commands* and *SPL: Initialization and Tuning* describe how to assign color attributes.
- Precise inquiries. By requesting only pertinent data, operators can reduce message volume. For example, by using D U,,,130,1 instead of D U to display the status of device 130, the operator generates 3 lines of output instead of 52 (a default maximum).

- A console other than the master console to start tasks, such as VTAM, that communicate with the console on which they are started. VTAM retains the ID of the console from which the START VTAM command is issued and directs all messages to that console. Therefore, if possible, start VTAM on a console defined to receive TP messages to reduce traffic on the master console.
 - A terminal dedicated to RMF. Use RMF Monitor II reports to determine what the system is doing. See *RMF Reference and User's Guide* for more information.
-

New Response to Message IOS201E

MVS issues message IOS201E after it recovers from an error condition that required it to stop processors that shared a resource. The message indicates whether or not the resource was lost (that is, whether a task was in the process of updating the resource and lost its reserve before the update was finished). After Release 1.1 is installed, operators must reply U to message IOS201E before the system restarts the processors that have been stopped. If the system cannot communicate with the operator console, it puts itself in restartable wait state x'114'. In earlier releases, the system displays the message for five seconds then automatically restarts all processors that were stopped.

Requiring a response improves system integrity. It ensures that the operator is aware of the lost resource and allows the operator options for recovery. The operator might, for example, want to re-IPL instead of restart processors that depended on the update being completed.

Summary of New, Changed, or Deleted Commands

Figure 4-2 summarizes the commands that are new, changed, or deleted in MVS/XA. Most of the changes are compatible (that is, commands that specify existing parameters, keywords, or options can be entered the same way in MVS/370 and MVS/XA). Exceptions are:

- TRACE (except TRACE STATUS)
- MODE E1, E2, E3, and E4
- VARY CHP, CPU, STOR, and PATH
- REPLY id,ASID in response to a DUMP command

These commands must be specified differently in MVS/XA.

The syntax of the following commands is the same. However, they produce different output in MVS/XA:

- DISPLAY M=DEV
- DISPLAY MPF
- REPLY SDATA=(NUC) or (TRT) in response to a DUMP command

System Commands describes the syntax of the commands and how to use them.

Command	Release			Description of Change
	1.0	1.1	1.2	
CANCEL	x			<p>Using the new A keyword, operators can terminate specified jobs, time-sharing users, or started processes that do not have unique names or have not yet been assigned job names. The A keyword identifies the task to be cancelled.</p> <p>If two tasks with the same name are both active when the operator issues a CANCEL command, the MVS/XA system rejects the command. The operator can then use A to specify the ASID of the task to be cancelled and reissue the command. In MVS/370, the system terminates the first task found.</p> <p>Operators can also use the A keyword to cancel partially-initiated tasks that have not yet been assigned names. The DISPLAY A,L command identifies those tasks as 'STARTING' or '*LOGON*.' The operator can cancel them by issuing a CANCEL command that specifies the A keyword and a job name of 'STARTING' or a user ID of '*LOGON*.'</p>
CHNGDUMP	x			<p>Following are new or changed options for SYSABEND, SYSMDUMP, SYSUDUMP, and SVC dumps: ALLVNUC, ALLNUC, NOSYM, NUC, SUM, TRT, and SUBTASKS. Figure 6-1 in Chapter 6 describes each of the options.</p>
CONFIG	x			<p>A new reconfiguration command that:</p> <ul style="list-style-type: none"> - Physically and logically reconfigures channel paths, processors, and storage. The CONFIG CPU, CONFIG CHP, and CONFIG STOR commands, respectively, replace the MVS/370 VARY CPU, VARY CH, and VARY STOR commands. - Allows the operator to reconfigure the system as specified in a CONFIGxx PARMLIB member. The CONFIG MEMBER command requests this function. Note that the CONFIG MEMBER command does not reconfigure devices or DASD volumes. - Allows the operator to select elements to be reconfigured from a display of elements that are in the current configuration or that can be brought online. The operator issues either CONFIG ONLINE or CONFIG OFFLINE to request this function. - Allows the operator to determine the online or offline status of all available processors, channel paths, and storage in the configuration.
CONTROL M			x	<p>A new keyword, UEXIT, activates and deactivates the IEAVMXIT user exit. If IEAVMXIT is in the LNKST concatenation at IPL time, the system automatically activates it. Thereafter, you can use the UEXIT keyword to control IEAVMXIT's status.</p> <p>IEAVMXIT is new in Release 1.2. See "New WTO/WTOR User Exits" in Chapter 5 for more information.</p>
CONTROL S		x		<p>Two new keywords:</p> <ul style="list-style-type: none"> - L=cc specifies which console the CONTROL S command is to affect. In earlier releases, you can change the console specifications for only the console on which the CONTROL command is issued. - MFORM specifies the form in which messages are to be displayed. The choices are: (1) the message text alone, (2) the message text with the issuer's job name or the job ID, or (3) the message text with either the issuer's job name or the job ID and a timestamp.

Figure 4-2 (Part 1 of 8). Summary of New, Changed, or Deleted Commands

Command	Release			Description of Change
	1.0	1.1	1.2	
CONTROL V		x		When processing a CONTROL V command that switches a console between full-capability and message stream modes, the system reestablishes the console specifications that were in effect the last time the console was in that mode during the IPL. Earlier releases reestablish the console specifications set during system generation.
			x	<p>A new LEVEL keyword specifies which type of messages a particular console will accept. You can use LEVEL in addition to routing codes to further control message traffic.</p> <p>The options on LEVEL are:</p> <p>ALL - All messages routed to the console. ALL is the default.</p> <p>R - WTORs.</p> <p>I - Immediate action messages (descriptor codes 1 and 2).</p> <p>CE - Critical eventual action messages (descriptor code 11).</p> <p>E - Eventual action messages (descriptor code 3).</p> <p>IN - Informational messages, excluding broadcast and action messages.</p> <p>NB - No broadcast messages, regardless of the descriptor code.</p> <p>UNCOND is also an option on LEVEL. It indicates the system is to perform the LEVEL request unconditionally, even if some broadcast or informational messages are sent only to the hardcopy log as a result. (The system sends WTORs and action messages that are suppressed from all consoles to the master console as well as to the hardcopy log.)</p> <p>If you specify LEVEL=UNCOND and messages are sent to hardcopy only, you receive a warning message. To identify which messages are sent only to hardcopy, use a DISPLAY CONSOLES command with the HCONLY keyword specified.</p> <p>If you omit the UNCOND option on a command that would cause some messages to go to hardcopy only, the system rejects the command.</p>

Figure 4-2 (Part 2 of 8). Summary of New, Changed, or Deleted Commands

Command	Release			Description of Change
	1.0	1.1	1.2	
DISPLAY CONSOLES			x	<p>Several new keywords request status information about specific consoles. They are useful for limiting the display to information that is pertinent. Previously, DISPLAY CONSOLES always produced the status of all consoles in the system.</p> <p>The new keywords and the consoles they specify are:</p> <p>ACTIVE - Active consoles, including the master console and the hardcopy console. ACTIVE is the default.</p> <p>NACTIVE - Consoles that are not active.</p> <p>SS - Consoles that subsystems can use.</p> <p>CN(xx) - Consoles whose IDs are listed on CN.</p> <p>U(aaa) - Consoles identified by the device numbers on U.</p> <p>ROUT(rr) - Consoles that receive messages with the routing codes specified on ROUT.</p> <p>BACKLOG - Consoles with a message backlog.</p> <p>MASTER - The master console and any pseudo-master consoles.</p> <p>* - The console from which the DISPLAY command is issued.</p> <p>LIST - All consoles defined to the system. Thus, the output from LIST is equivalent to the output from previous DISPLAY CONSOLES commands.</p> <p>You can now route the DISPLAY CONSOLES command. A new L keyword specifies the display area, console, or both where the system is to display the output.</p> <p>The HONLY keyword is also new. It requests information about messages that are recorded only on the hardcopy log and not sent to any console. The display lists routing codes not assigned to any console. If any broadcast messages are being sent to hardcopy only, the display also includes the word BROADCAST.</p> <p>The HONLY keyword is introduced in Release 1.2 because of changes to the CONTROL V command. You can now use CONTROL V to specify which type of messages a console does or does not accept. Thus, it is possible that some broadcast and informational messages are routed to hardcopy only. For more detail, see the entry for the CONTROL V command.</p>
DISPLAY DUMP	x			<p>Two new keywords, TITLE and ERRDATA, display information stored in the header record of DASD SYS1.DUMP data sets that are full. The TITLE keyword lists the dump titles; the ERRDATA keyword displays error data from the dump header. TITLE and ERRDATA are valid options for displaying DASD dump data set information only. You can specify the DSN operand with the TITLE and ERRDATA keywords to restrict the display to a specific group of DASD dump data sets.</p> <p>The OPTIONS parameter supports the new SYSABEND, SYSUDUMP, SYSMDUMP, and SVC dump keywords. Figure 6-1 describes the new keywords.</p> <p>The information displayed when the STATUS operand is specified has changed. The display lists the full and available dump data sets grouped by device (DASD and tape). It does not provide any titles. Users must specify the TITLE operand to display dump titles. Previously, the display gave the status (full or available) of each dump data set on a separate line and included the titles of full dump data sets.</p>

Figure 4-2 (Part 3 of 8). Summary of New, Changed, or Deleted Commands

Command	Release			Description of Change
	1.0	1.1	1.2	
DISPLAY GRS			x	<p>Release 1.2 provides new keywords for displaying additional global resource serialization information:</p> <ul style="list-style-type: none"> - CONTENTION displays information about all resources for which at least one task is waiting. For a description of the information shown, see the RES keyword entry. - RES displays information about every resource that is currently allocated. The display includes the resource's rname, qname, and scope, and the following information about each task that owns or is waiting for the resource: the sysname, jobname, ASID, TCB address, type of request (shared or exclusive), and the task's status (owner or waiting). - RNL displays one or more RNLs, depending on which options RNL specifies. - HEX displays CONTENTION, RES, and RNL information in hexadecimal as well as regular format. - ALL displays the information that CONTENTION, RNL, SYSTEM, and LINK request. (SYSTEM and LINK are old keywords.) <p>If you specify DISPLAY GRS with no keywords, you see SYSTEM and LINK information, which is consistent with earlier releases.</p> <p>Note: You can display CONTENTION, RES, and ALL information even if the system is not part of an active global resource serialization complex.</p>
DISPLAY M	x			<p>New and changed parameters:</p> <ul style="list-style-type: none"> - ddd requests the online or offline status of all channel paths to device ddd. - nn requests the status of each device on channel path nn. - CHP requests the status of all channel paths in the system. - DEV displays the number of online channel paths to each device defined during sysgen. In MVS/370, DEV displays the online or offline status of all devices and the channel sets to which the online devices are connected.
DISPLAY MPF	x			<p>The display now also includes the color and highlighting defaults for 3279 MCS consoles.</p>
			x	<p>DISPLAY MPF displays additional information in Release 1.2. It has two new keywords for limiting the output, MSG and COLOR.</p> <p>MSG displays:</p> <ul style="list-style-type: none"> - The IDs of messages MPF is to suppress - The IDs of action messages that the action message retention facility does not retain - Which WTO/WTOR user exits are associated with which messages - Whether or not the general WTO/WTOR exit (IEAVMXIT) is active <p>COLOR displays the color, intensity, and highlighting options in effect for 3279 consoles.</p> <p>If you specify no keywords, you see all of the information.</p>
DUMP	x			<p>The syntax of the DUMP command is the same; the content of the reply to the command is different.</p> <p>The NUC option on the SDATA keyword requests only the non-page-protected DAT-on section of the nucleus. Previously, it requested the entire nucleus.</p> <p>ALLNUC, a new option on the SDATA parameter, requests a dump of the DAT-off nucleus and the entire DAT-on nucleus, including the page-protected section.</p> <p>Users must specify the ASID list in hexadecimal rather than decimal values. Thus, the ASID specifications in the dump request are consistent with the message describing its completion and the internal description of the dump contents.</p>

Figure 4-2 (Part 4 of 8). Summary of New, Changed, or Deleted Commands

Command	Release			Description of Change
	1.0	1.1	1.2	
DUMPDS	x			<p>Using the new DUMPDS command, operators can:</p> <ul style="list-style-type: none"> - Add or delete SYS1.DUMP data sets after IPL/NIP time without having to re-IPL. Before adding a data set, it must be allocated and cataloged. In MVS/370, installations can only add and delete dump data sets at IPL/NIP time. - Be aware that if the DUMPSRV address space terminates, you must again add dump data sets that were added using the DUMPDS command. When DUMPSRV restarts, the system adds the dump data sets that were added at IPL/NIP time. - Clear SYS1.DUMP data sets on DASD or tape. The DUMPDS CLEAR command avoids having to either run a utility job (AMDPRDMP or IEBGENER) or reset, load, and ready a tape drive to clear a DASD dump data set. <p>Operators can issue the DUMPDS command only from consoles that have system authority.</p>
FORCE	x			<p>Using new A and ARM options on the FORCE command, operators can terminate specified jobs, time-sharing users, or started processes that:</p> <ul style="list-style-type: none"> - Do not have unique names - Have not yet been assigned job names. - Are not eligible for cancellation via the CANCEL command. <p>If two tasks with the same name are both active when the operator issues a FORCE command, the MVS/XA system rejects the command. The operator can then use A to specify the ASID of the task to be cancelled and reissue the command. In MVS/370, the system terminates the first task found.</p> <p>Operators can also use the A keyword to cancel partially-initiated tasks that have not yet been assigned names. The DISPLAY A,L command identifies those tasks as 'STARTING' or '*LOGON*.' The operator can cancel them by issuing a FORCE command that specifies the A keyword and a job name of 'STARTING' or a user ID of '*LOGON*.'</p> <p>The ARM keyword requests that the specified job, time-sharing user, or started process be terminated, even if it is not eligible for cancellation via the CANCEL command. Operators can use ARM to cancel any address space except one that is not eligible for termination (ASCBNOMT=1).</p>
MODE	x			ENABLE, E1, E2, E3, and E4 are no longer valid parameters and cause the MODE command to be rejected.
MODIFY		x		New keywords, LLA, REFRESH, cause the system to rebuild the LNKLST lookaside (LLA) directory. The LLA directory is new in Release 1.1. See "Using a New Directory for LNKLST Data Sets" in Chapter 8 for more information.
MONITOR			x	<p>You can now route the MONITOR command. In earlier releases, the output is always displayed at the console from which MONITOR is issued. To route MONITOR, either:</p> <ul style="list-style-type: none"> - Include the L=cc keyword on MONITOR. L=cc specifies the ID of the console on which the output is to be displayed. - Issue an MSGRT command that specifies the new MN keyword.
MSGRT	x			A new keyword, CF, specifies to which MCS console the system is to route CONFIG commands.
		x		You can now issue the MSGRT K command from display, as well as non-display, consoles. Earlier releases restrict its use to non-display consoles only.
			x	A new MN keyword routes MONITOR and STOPMN commands. MN specifies the console to which the commands are to be routed. Neither of those commands can be routed in earlier releases.

Figure 4-2 (Part 5 of 8). Summary of New, Changed, or Deleted Commands

Command	Release			Description of Change
	1.0	1.1	1.2	
SET		x		<p>DAE=xx is a new keyword that specifies which ADYSETxx PARMLIB member the system is to process. ADYSETxx contains options for controlling dump analysis and elimination (DAE). Issuing SET DAE=xx causes the system to begin using the options in the specified member. See "Dump Analysis and Elimination (DAE)" in Chapter 6 for more information.</p> <p>You can now use SET SMF to restart SMF after it is terminated. Because SMF runs in its own address space, you no longer need to perform an IPL to restart it.</p>
SLIP MOD			x	<p>You can now enable or disable several SLIP traps with one command. Release 1.2 allows asterisks in place of any or all of the four characters of the ID keyword. The system enables or disables all traps having identifiers that match the characters you do specify. For example, SLIP MOD,ENABLE,ID=0*** causes the system to enable all SLIP traps that have identifiers beginning with 0.</p>
SLIP SET	x			<p>New options and keywords in Release 1.0:</p> <ul style="list-style-type: none"> - NOSYSA, NOSYSM, NOSYSU, and NOSVCD are new options on the ACTION keyword. They suppress respectively SYSABEND, SYSMDUMP, SYSUDUMP, and SVC dumps for specified abend conditions. See "Suppressing Dumps" in Chapter 6 for more information. - ALLNUC on the SDATA keyword requests the entire nucleus (both the DAT-on and DAT-off sections). <p>Indirect addresses on SLIP command keywords (DATA, LIST, SUMLIST, and TRDATA) can be 31-bit values. To indicate that an indirect address is to be treated as a 31-bit value, use a '?' instead of a '%' as the indirect address indicator. When a keyword specifies '%', the system treats the indirect address as a 24-bit value.</p>
			x	<p>New options and keywords in Release 1.1:</p> <ul style="list-style-type: none"> - NUCMOD specifies a nucleus module or an address range within a nucleus module. When specified on an error event trap, NUCMOD indicates the range within which the error must occur. On IF or SB PER traps, it establishes the range of addresses to be monitored. On SA PER traps, NUCMOD specifies boundaries for the instruction causing the storage alteration. - NOSUP, a new parameter on the ACTION keyword of error event traps, indicates that the system is to take dumps for the trapped event, regardless of any attempts to suppress the dumps. Thus, NOSUP overrides dump suppression via dump analysis and elimination (DAE) or the ABDUMP pre-dump exit routine. - AND (&) or OR () on the DATA keyword logically compare triplets of target locations, operators, and values. You can specify AND and OR together with any number of triplets. You can also group triplets using parentheses. If you specify no logical operator, AND is the default, which is consistent with earlier releases. In earlier releases, AND is the implied logical operator. - REASON, a new keyword on error event traps, specifies a reason code for filtering errors. It is valid only with the COMP (completion code) keyword.

Figure 4-2 (Part 6 of 8). Summary of New, Changed, or Deleted Commands

Command	Release			Description of Change
	1.0	1.1	1.2	
SLIP SET (continued)			x	<p>New keywords and options in Release 1.2 are:</p> <ul style="list-style-type: none"> - RECORD, a new option on the ACTION keyword of error event traps, causes recovery routines to record the trapped error in SYS1.LOGREC. - STRACE, a new option on the ACTION keyword of PER traps, causes the system to write a system trace table entry for each trapped event. Thus, you can use PER to isolate problems without stopping the system when the program event occurs. - PVTMOD is now valid on all PER and non-PER traps. Previously, it was an option only on non-PER and storage alteration PER traps. Thus, for the first time you can monitor instruction fetch and successful branch events in the private area. - PVTEP, LPAEP, and NUCEP specify entry points in modules that reside in the private area, LPA, and nucleus, respectively. They cause monitoring to begin at the address associated with the entry point or at the specified offset from that address. PVTEP and LPAEP are particularly useful for monitoring the section of a module that begins at an alternate entry point. NUCEP is equivalent to NUCMOD. <p>You can now monitor events in modules whose names end in x'CO' (SVC load modules). Because x'CO' is not alphanumeric, the system does not accept in on the LPAMOD, LPAEP, NUCMOD, NUCEP, PVTMOD, or PVTEP keywords. However, you can now use an asterisk in place of x'CO'. The system interprets the asterisk as x'CO'.</p>
START	x			<p>A new keyword, SUB, directs the JCL for a started task to the internal reader of a secondary JES or to the master subsystem. By routing started tasks to a secondary JES or to the master subsystem, operators can:</p> <ul style="list-style-type: none"> - Run started tasks independently of the primary JES. - Start certain tasks before the primary JES initialization is completed. <p>See <i>SPL: System Modifications</i> for more information.</p>
		x		<p>A new keyword, LLA, starts the LLA procedure, which in turn starts the LNKLST lookaside (LLA) function. The LLA function builds and maintains a new directory of modules in the LNKLST concatenation. BLDL then searches that directory instead of the PDS directories for LNKLST modules. For more information about the LLA function, see "Using a New Directory for LNKLST Data Sets" in Chapter 8.</p>
STOP		x		<p>A new keyword, LLA, stops the LNKLST lookaside (LLA) function. The system then searches PDS directories instead of the LLA directory to locate modules in the LNKLST concatenation. For more information, see "Using a New Directory for LNKLST Data Sets" in Chapter 8.</p>
STOPMN			x	<p>You can now use STOPMN to suppress MONITOR command output from any console except the one from which STOPMN is issued. A new keyword, L=cc, identifies the console on which the output is to be suppressed.</p>

Figure 4-2 (Part 7 of 8). Summary of New, Changed, or Deleted Commands

Command	Release			Description of Change
	1.0	1.1	1.2	
TRACE	x			<p>The syntax of the TRACE command has changed. Except for TRACE STATUS, MVS/370 TRACE commands do not work in MVS/XA. In addition, the TRACE command controls system tracing differently:</p> <ul style="list-style-type: none"> - MVS/XA continues system tracing after system initialization, unless an installation requests that tracing be stopped. Thus, no TRACE command is required in the COMMNDxx PARMLIB member to request system tracing. In MVS/370, to continue system tracing after system initialization, you have to issue a TRACE ON command. - The MVS/XA trace facility performs branch tracing, address space tracing, or explicit tracing. ("Changes to the System Trace Facility" in Chapter 5 lists the system events in each category.) A new TRACE operand, BR, allows installations to start and stop branch tracing independently of address space and explicit tracing. Options are: <ul style="list-style-type: none"> --Explicit and ASID tracing on, branch tracing off --All tracing on --All tracing off <p>The system can perform branch tracing only when the other trace options are active.</p> <ul style="list-style-type: none"> - TRACE can start or stop system tracing or change the TRACE options at any time after system initialization is completed. To start system tracing in MVS/370, you have to issue a TRACE ON command before the system starts JES2 or JES3. Also, in MVS/370 you cannot use the TRACE command to stop system tracing after subsystem initialization is completed. - TRACE allows you to dynamically change the size of the system trace table using a new option on the ST keyword. The default size is 16 K per processor. In MVS/370, the trace table size is fixed at IPL time. <p>The master trace (MT) and display trace status (STATUS) functions of the TRACE command are not changed. However, you must use the new syntax of the TRACE command to start or stop master trace, as well as system trace. The order in which you specify the parameters is changed. <i>System Commands</i> describes the new syntax.</p>
VARY CPU, VARY CH, VARY STOR	x			<p>The VARY CPU, VARY CH, and VARY STOR commands are deleted. The new CONFIG CPU, CONFIG CHP, and CONFIG STOR commands perform equivalent functions in MVS/XA. See the CONFIG entry in this table.</p>
VARY PATH	x			<p>The syntax is changed to support the new I/O architecture. MVS/370 VARY PATH commands do not work in MVS/XA. Also, the VARY PATH,OFFLINE command might take longer to process. The system will not vary a device path offline until I/O activity to the target device has completed. If I/O is not completed after 150 seconds, the system issues message IEE717D, which gives the operator a chance to cancel the command.</p>

Figure 4-2 (Part 8 of 8). Summary of New, Changed, or Deleted Commands

Chapter 5. System Modifications

This chapter contains information related to modifying the system. The topics it includes are:

- “Print Dump Exit Control Table (ECT) Modifications”
- “Updating SYSTEMS Exclusion RNLs”
- “Serializing VSAM Data Sets”
- “Limiting User Region Size using IEFUSI Instead of IEALIMIT” on page 5-2
- “Obtaining an Extended Region Size of More Than 32 Mb” on page 5-3
- “Bypassing the Storage Availability Check Before a Job Executes” on page 5-3
- “Changing the Hot I/O Threshold and Recovery Actions” on page 5-4
- “Pre-dump Exits” on page 5-4
- “Post-dump Exits” on page 5-4
- “RMF Exits” on page 5-4
- “JES2 User Exits” on page 5-4
- “JES2 Interfaces” on page 5-5
- “JES3 Dynamic Support Programs (DSPs) and User Exits” on page 5-5
- “PRDMP Exits” on page 5-6
- “PRDMP Header Exits” on page 5-6
- “SMF Exits” on page 5-6
- “New WTO/WTOR User Exits” on page 5-7
- “New Services for Dump Processing Exits” on page 5-8

Print Dump Exit Control Table (ECT) Modifications

The system uses some previously-empty ECT entries for new print dump exits. If your installation added ECT entries, you must add them to the new ECT at different offsets. The ECT is in module AMDPRECT.

Updating SYSTEMS Exclusion RNLs

If a system with Release 1.1 installed is part of a global resource serialization complex, you might have to add the resource name for the SYS1.DAE data set (SYSDSN, SYS1.DAE) to the SYSTEMS exclusion RNLs of other systems in the complex. Because SYS1.DAE cannot be shared among systems, the SYSTEMS exclusion RNL shipped in Release 1.1 includes an entry for SYS1.DAE. Global resource serialization requires that all systems in a complex have identical RNLs. Therefore, you need to add the same entry to the SYSTEMS exclusion RNLs of any system in the complex that is not at Release 1.1 or a later level.

Serializing VSAM Data Sets

If your installation has a global resource serialization complex that includes both systems with MVS/XA DFP 1.1 or MVS/370 DFP installed and systems with neither, you need to take certain actions to ensure that VSAM data sets are serialized correctly. MVS/XA DFP 1.1 and MVS/370 DFP include a new level of VSAM. The ENQs that the new VSAM OPEN processing issues to enforce VSAM cross-region share options 1 and 2 are different. The scopes of the ENQs are changed to SYSTEMS, and their rnames (minor names) include catalog names. Only the qnames (major names) remain the same.

A catalog name is system-independent information. Therefore, if all systems in the complex include the new level of VSAM and all systems accessing the data set belong to the complex, you can use the VSAM ENQ to serialize access to the data set as a global resource. The scope of SYSTEMS causes global resource serialization to treat the data set as a global resource by default.

The old level of VSAM uses ENQs with scopes of SYSTEM and rnames that include catalog ACB addresses. Because catalog ACB addresses vary from system to system, you cannot use these ENQs to serialize access to data sets as global resources. The scope of SYSTEM causes global resource serialization to treat the data set as a local resource by default.

The following figure summarizes the differences in the ENQs:

	Old VSAM ENQs	New VSAM ENQs
QNAME	SYSVSAM	SYSVSAM
RNAME	system-dependent	system-independent
SCOPE	SYSTEM (local resource)	SYSTEMS (global resource)

Because of the differences, systems in complexes that include both levels of VSAM cannot share the VSAM data sets globally. Therefore, place a generic entry for SYSVSAM in the SYSTEMS exclusion RNL of each system. Also ensure that the SYSTEM inclusion RNLs do NOT include an entry for SYSVSAM.

After all systems in the complex have the new level of VSAM installed, remove the SYSVSAM entry from the SYSTEMS exclusion RNLs. The VSAM OPEN processing then enforces share options 1 and 2 as follows:

- If a data set assigned share option 1 is opened for output, no other user in the complex can open it for output or input.
- If a data set assigned share option 2 is opened for output, other users in the complex can open it for input but not output.

Limiting User Region Size using IEFUSI Instead of IEALIMIT

Installations can now use the SMF step initiation exit (IEFUSI) to limit the sizes of user regions above and below 16 Mb. The methods available in MVS/370 continue to work in MVS/XA:

- Specifying the REGION parameter on JCL statements
- Assigning default values through JES2 or JES3
- Using the IEALIMIT installation exit

However, using IEFUSI has the following advantages:

- IEFUSI is a separate load module in the LPA. IEALIMIT must reside in the nucleus. Thus, you must link edit the nucleus every time you replace IEALIMIT with a new version.
- IEFUSI users can readily obtain information required to set a region size and region limit. IEALIMIT must scan system control blocks to gather that information. Thus, IEFUSI is easier to write and less susceptible to system changes.

- IEALIMIT requires that the local lock be held and, therefore, cannot issue SVCs. IEFUSI has neither of those restrictions.
- IEFUSI can control the region size and region limit of both the area above and the area below 16 Mb. IEALIMIT can set values for the area below 16 Mb only; VSM uses defaults defined in the code for the area above 16 Mb.

To indicate that VSM is to use IEFUSI instead of IEALIMIT for controlling the user region area, you must set a flag in the IEFUSI parameter list. VSM then bypasses the IEALIMIT exit and limits the user area as IEFUSI requests. IEFUSI does not have to set new limits. It can, for example, request that VSM set limits identical to the IEALIMIT defaults.

SPL: User Exits describes how to use IEALIMIT. *SPL: SMF* describes how to use IEFUSI. *SPL: System Modifications* provides general information on limiting the user region.

Obtaining an Extended Region Size of More Than 32 Mb

VSM changes in Release 1.2 make it easier for a job to obtain an extended region size greater than 32 Mb:

- You can now specify values greater than 16 Mb on the JCL REGION parameter. The values can be expressed in K or Mb units and can be as high as 2096128 K or 2047 Mb.
- VSM now uses the REGION parameter (if nonzero) to calculate the extended region size and the limit for GETMAIN requests above 16 Mb. VSM sets both to the smaller of: (1) the size of the extended private area, or (2) the REGION parameter value or 32 Mb, whichever is greater. In earlier releases, VSM sets both to 32 Mb.

If the REGION parameter is greater than 16 Mb, the only limits on GETMAIN requests below 16 Mb and the region size below 16 Mb are the limits that IEALIMIT or IEFUSI set, or the size of the private area. *System Modifications* contains more information about limiting the user region size.

Bypassing the Storage Availability Check Before a Job Executes

If Release 1.2 is installed, you can control whether or not VSM checks that the amount of storage requested on the REGION parameter is available before permitting a job to execute. The change is intended to prevent jobs from failing simply because programmers specify REGION values that are unnecessarily large.

Earlier levels of VSM always check for availability of storage below 16 Mb but never check for availability above 16 Mb. The Release 1.2 level of VSM does the same by default.

To change the default, use IEFUSI. Bits 1 and 2 in the first word of the VSM parameter list passed to IEFUSI control checking. Bit 1 indicates whether VSM is to check for available storage below 16 Mb. Bit 2 controls checking above 16 Mb. See *System Modifications* and *SPL: SMF* for more detail.

Changing the Hot I/O Threshold and Recovery Actions

A new HOTIO statement in the Release 1.2 IECIOSxx PARMLIB member makes it easier to change: (1) the threshold IOS uses for detecting hot I/O conditions, and (2) the recovery actions IOS performs when it detects a hot I/O condition. Module IOSRHIDT contains the threshold and recovery actions. It replaces the MVS/370 IECVHIDT module. Previously, you had to use the AMASPZAP service aid to update IOSRHIDT (or IECVHIDT).

The first release of MVS/SP Version 2 changes hot I/O processing. "Processing Hot I/O Interrupts" in Chapter 4 briefly describes the changes. *System Modifications* contains more detail. *Initialization and Tuning* describes how to write the HOTIO statement.

Pre-dump Exits

Installations can now provide exit routines that get control before the ABDUMP routine takes a dump. The exits can analyze the requested dump and either:

- Continue with the dump as requested
- Modify the dump options and continue with the dump
- Terminate the dump request.

IBM supplies a load module, IEAVTABX, that contains blank entries for the pre-dump exit routine names. *SPL: System Modifications* describes pre-dump exits. *SPL: User Exits* provides coding information.

Post-dump Exits

Installations can now provide exit routines that get control after each SYSMDUMP and SVC dump is taken. The post-dump exit routines can examine dumps in dump data sets, evaluate the dump and the dump symptoms, and take appropriate action (for example, tell the operator to clear the dump data set using the new DUMPDS command or to start a PROC to offload the dump data set). IBM supplies a new load module, IEAVTSEL, that contains an SDUMP exit list with blank entries.

SPL: System Modifications describes post-dump exits in more detail. *SPL: User Exits* provides coding information.

RMF Exits

RMF exits require modification. RMF invokes all user exits in 31-bit addressing mode and expects return in that mode. In addition, most RMF user exits need to access control blocks that have been moved to virtual storage above 16 Mb.

JES2 User Exits

JES2 invokes user exits in 24-bit addressing mode and expects return in that mode. Thus, JES2 user-exit routines must reside in virtual storage below 16 Mb. They can, however, switch modes after entry as long as they do not use JES2 interfaces while in 31-bit addressing mode. Note that these restrictions do not apply to the SMF exits that JES2 takes in the JES2 or user address space.

If a JES2 exit contains any of the fourteen downward incompatible macros listed in Chapter 9, you also need to ensure that the exit (the part of it that uses JES2 interfaces) contains their MVS/370 expansions. JES2 requires the MVS/370 expansions.

The MVS/XA MACLIB contains both the MVS/370 and MVS/XA expansions of the downward incompatible macros. To ensure that exit routines include the MVS/370 expansions, use either a \$HASPGEN or a \$HASPEQU macro before any other macro. (If your system is at the Release 1.0 level, use \$HASPGEN; if it is at a later level, use \$HASPEQU.) Both \$HASPGEN and \$HASPEQU issue SPLEVEL macros that request the MVS/370 expansions (SPLEVEL SET=1).

If an exit routine requires the MVS/XA expansion of a downward incompatible macro, use an SPLEVEL macro before and after the incompatible macro. On the SPLEVEL macro that precedes the incompatible macro, request the MVS/XA expansion (SPLEVEL SET=2). On the SPLEVEL macro that follows, request MVS/370 expansions of subsequent downward incompatible macros.

For more information about downward incompatible macros or SPLEVEL, see “Handling Downward Incompatible Macros” in Chapter 9.

JES2 Interfaces

Programs executing in 31-bit addressing mode cannot invoke JES2's subsystem interface (SSI) routines or its spool access method (HAM) routines. Spool data set requests (for example, OPEN, CLOSE, GET, PUT) must be made while executing in 24-bit addressing mode. All control blocks used as input to the JES2 subsystem via SSI (for example, SSOBs) must reside in virtual storage below 16 Mb.

JES3 Dynamic Support Programs (DSPs) and User Exits

JES3 invokes DSPs and user exits in 24-bit addressing mode and expects return in that mode. Thus, JES3 user-written programs must reside in virtual storage below 16 Mb. They can, however, switch modes after entry as long as they do not use JES3 interfaces while in 31-bit addressing mode.

If a JES3 exit contains any of the fourteen downward incompatible macros listed in Chapter 9, you also need to ensure that the exit (the part of it that uses JES2 interfaces) contains their MVS/370 expansions. JES3 requires the MVS/370 expansions.

The MVS/XA MACLIB contains both the MVS/370 and MVS/XA expansions of the downward incompatible macros. To ensure that exit routines include the MVS/370 expansions, use an IATYMOD macro before any other macro. IATYMOD issues an SPLEVEL macro that request the MVS/370 expansions. (SPLEVEL SET=1).

If an exit routine requires the MVS/XA expansion of a downward incompatible macro, use an SPLEVEL macro before and after the incompatible macro. On the SPLEVEL macro that precedes the incompatible macro, request the MVS/XA expansion (SPLEVEL SET=2). On the SPLEVEL macro that follows, request MVS/370 expansions of subsequent downward incompatible macros.

For more information about downward incompatible macros or SPLEVEL, see “Handling Downward Incompatible Macros” in Chapter 9.

PRDMP Exits

You need to change PRDMP exit routines that use the print service and supply their own output buffers. In MVS/XA, the print service routine expects a 132-byte output buffer. It prints 132 bytes, beginning at the output buffer address specified in the ADPLBUF field of the exit parameter list (BLSABDPL). In MVS/370, the output buffer is 121 bytes, but only 120 bytes are printed.

PRDMP exits that use the PRDMP-supplied output buffer continue to work unchanged in MVS/XA. The buffer is set to blanks after each use. However, if you modify the ADPLBUF field to point to a 121-byte output buffer, either:

- Each printed line will contain 12 bytes of unexpected data.
 - The PRDMP exit will fail with an x'0C4' ABEND.
-

PRDMP Header Exits

The Release 1.1 level of PRDMP allows a new type of user exit, header exits. Using header exits you can add information to the title pages of dumps. PRDMP calls header exits when processing title pages.

Dump analysis and elimination (DAE) supplies one header exit, ADYHDFMT, which formats and prints DAE symptom data. You can supply additional header exits, but do not change the ECT (exit control table) entry for ADYHDFMT (entry 21). PRDMP also calls that entry when processing the new DAEDATA PRDMP verb.

SMF Exits

You might have to modify two SMF exits, IEFU29 and IEFU84. In Release 1.1, IEFU29 and sometimes IEFU84 run in the new SMF address space instead of the master scheduler address space. (IEFU84 runs in the SMF address space when SMF calls it during system initialization to write record types 0, 8, 19, 22, and 90.) If either exit requires data located in the private area of the master scheduler address space, you need to change the exit. Use cross memory instructions (SSAR, MVCP, and MVCS) to move data between the two address spaces.

Other SMF exits continue to run unchanged in the same address spaces as in previous releases of MVS. You can, however, write SMF exits that run in 31-bit addressing mode, reside above 16 Mb, or address data above 16 Mb. Exits that run in 31-bit addressing mode must return control to SMF using a BSM instruction. In addition, SMF records that IEFU83 or IEFU84 passes to SMF must reside below 16 Mb. See *SPL: 31-bit Addressing* for help in writing programs that use 31-bit addresses.

New WTO/WTOR User Exits

Installations with Release 1.2 installed can use new WTO/WTOR exits to modify message processing. The new exits are in addition to IEECVXIT. Unlike IEECVXIT, they can modify processing for any message. They can also alter processing in more ways than IEECVXIT can. The new exits can:

- Alter routing and descriptor codes.
- Change the message text.
- Change the console on which the message is displayed.
- Queue messages to a particular active console; queue them unconditionally to any console, regardless of whether it is active; or queue messages by routing codes only.
- Direct messages to the hardcopy log only, to consoles only, or to both the hardcopy log and consoles.
- Delete messages, except WTORs.
- Control whether or not messages are broadcast to active consoles.
- Override MPF suppression.
- Issue SVCs (for example, SVC 34 and SVC 35).
- Reply to or suppress WTORs.
- Control whether or not the action message retention facility retains a message.

The system invokes WTO/WTOR exits after IEECVXIT and MPF processing is completed and before calling the subsystem interface (SSI). It can invoke only one WTO/WTOR exit for each message processed. However, you can provide several exits and specify which the system is to call for particular messages. You can also write one general exit that the system invokes for all messages not associated with a specific exit. You must name the general exit IEAVMXIT. Both the general and specific WTO/WTOR exits must execute in 31-bit addressing mode and reside in an authorized data set that is included in the LNKLIST concatenation.

To associate a specific exit with a message or group of messages, include statements like the following in the MPFLSTxx PARMLIB member:

```
message ID  USEREXIT(name of exit routine)
```

The message ID can identify either a particular message or a class of messages, for example, IEF404I or IEF*. The system then calls the specified exit when processing those messages. (You can also put the new RETAIN and SUP keywords on the same statement as USEREXIT. See “New, Changed, or Deleted PARMLIB Members” in Chapter 2 for more information.) Use multiple MPFLSTxx members and the SET MPF=xx or SET MPF=NO command to control which specific exits are active.

If the system is to call IEAVMXIT when processing a message, you need not do anything except make that exit available and possibly activate it. If IEAVMXIT is in the LNKLIST concatenation at IPL time, the system automatically activates it. Thereafter, you can activate and deactivate it using a new UEXIT=Y|N keyword in the CONTROL V command.

New Services for Dump Processing Exits

Release 1.2 provides several new services for dump processing exits that are invoked from IPCS, PRDMP, and SNAP:

- Format model processor service
- Control block formatter service
- ECT service
- GET symbol service
- EQUATE symbol service
- Select ASID service

In addition, Release 1.2 includes a new exit services router, which serves as an interface between dump exits and new and existing exit services. However, Release 1.2 continues to support existing interfaces between dump exits and the services available to them. Thus, existing dump exits will run unchanged in Release 1.2.

Exit Services Router

Dump exits can use the exit services router to invoke the following services:

- Storage access service
- Print service
- Format model processor service
- Control block formatter service
- Index service
- ECT service
- GET symbol service
- EQUATE symbol service
- Select ASID service

The storage access, print, and index services are not new. The service router simply provides a new method of invoking them.

Dump exits can invoke any of the services listed by calling the exit service router and passing it three parameters:

- The address of the user exit parameter list, ABDPL, which is mapped by BLSABDPL.
- A service code indicating which service the router is to invoke.
- The address of the parameter list the requested service expects (except when invoking the print or index service, which uses no parameter list). ABDPL includes each service's parameter list.

Most services also require some additional information in fields of the ABDPL. *IPCS User's Guide and Reference* describes those requirements. It also describes in more detail how to invoke services using the exit services router.

Format Model Processor Service

The format model processor service formats and prints a control block or other data area using a format model. The exit routine supplies the model and the virtual address of the storage to be formatted. The format model processor returns the formatted control block.

In addition to formatting control blocks, you can use format models and the format model processor to:

- Decode flag bytes
- Summarize dump data
- Present data in hexadecimal representation

The *IPCS User's Guide and Reference* describes each of these uses in more detail.

Using format models is an alternative to using format patterns. A format model consists of several BLSQMDEF and BLSQMFLD macros. The BLSQMDEF macro begins the model and describes the header. A series of BLSQMFLD macros follow, one for each field to be formatted. A second BLSQMDEF macro denotes the end of the model. The models can be part of a program load module or separate load modules. *System Macros and Facilities* describes how to write BLSQMDEF and BLSQMFLD macros. The *IPCS User's Guide and Reference* shows how to write format models.

Format models are flexible. Programs that use models can call a user-supplied exit routine to examine or modify a formatted line before printing it. They can also vary which fields of the model are printed. Thus, programs can use the same model to format different levels of the control block. Models can format arrays within a control block. They are also independent of the line length, starting column, and column spacing. PRDMP, IPCS, or SNAP modules determine that information.

Control Block Formatter Service

The control block formatter service formats and prints a control block, given a control block acronym and a bit string indicating which type of fields are to be included. You can use this service to format any of the following control blocks:

ASCB	LCCA	SRB
ASXB	LLE	STKE
CDE	PCCA	SUPVT
CSD	PSA	SVT
CVT	RB	TCB
EED	RTCT	TIOT
FRRS	RTM2WA	XSB
IHSA	SCB	XTLST

Using the control block formatter has several advantages over using format patterns:

- The user needs to know nothing about the control block structure.
- You can invoke the service once and format the entire control block.
- If the control block format changes, only the control block model that the control block formatter service uses needs to be modified.
- You can request multiple levels of detail using only one control block model.

ECT Service

The ECT service searches the ECT (exit control table) for the requested exit, then links to it. This service allows one dump exit to invoke another.

GET Symbol Service

The GET symbol service returns a specified symbol from the IPCS symbol table. Thus, it performs the same function as the GET subcommand of IPCS and is meaningful only during IPCS processing. PRDMP and SNAP modules ignore requests for this service.

EQUATE Symbol Service

The EQUATE symbol service adds a symbol to the IPCS symbol table, as does the EQUATE subcommand of IPCS. Like the GET symbol service, it is meaningful only during IPCS processing. PRDMP and SNAP modules ignore requests for this service.

Select ASID Service

The select ASID service returns a list of pointers to ASCBs in a dump. The exit routine specifies which ASCB pointers are to be returned using one or more of the following keywords:

ALL	All ASCBs in the dump
CURRENT	ASCBs for address spaces that were active at the time the dump was taken
ERROR	ASCBs for all address spaces that terminated abnormally (ASCBCC≠0) or that contain TCBs representing tasks that completed abnormally (TCBCMP≠0 or TCBRTWA≠0)
TCBERROR	ASCBs for all address spaces that contain TCBs representing tasks that completed abnormally (TCBCMP≠0 or TCBRTWA≠0). TCBERROR specifies a subset of the ASCBs that ERROR selects.
ASIDLIST	ASCBs corresponding to the ASIDs listed on ASIDLIST.
JOBLIST	ASCBs corresponding to the jobs listed on JOBLIST.

Chapter 6. Problem Determination

This chapter includes information related to dumping services, trace facilities, and debugging.

The following topics describe differences in dump content and format:

- “New and Changed Dump Options” on page 6-2
- “New Symptom Dumps for Task-Mode Abends” on page 6-4
- “New User Summary Dumps” on page 6-5
- “Dump Format Changes” on page 6-6

Topics describing new ways of suppressing dumps include:

- “New Operands on the SLIP Command for Suppressing Dumps” on page 6-7
- “MVS/XA’s Use of SLIP Commands” on page 6-7
- “Dump Analysis and Elimination (DAE)” on page 6-8

Print dump (PRDMP) changes are described in:

- “New and Changed PRDMP Control Statements” on page 6-10
- “Print Dump Index” on page 6-12
- “Print Dump Requirements for Printers” on page 6-13

The topics below describe IPCS changes introduced in Release 1.2. Note that to use Release 1.2 IPCS dialogs, you must also have ISPF Version 2 installed.

- “New and Changed IPCS Subcommands” on page 6-13
- “Accessing Additional Sources of Dump Data Using IPCS” on page 6-16
- “New IPCS Panels” on page 6-16
- “Changes to the IPCS BROWSE Panels” on page 6-17
- “Changes to the Titles of IPCS Print Files” on page 6-18

“Using the MVS/XA Versions of IPCS and PRDMP on Other Systems” on page 6-18 describes how to obtain MVS/XA versions of IPCS and PRDMP on earlier systems.

Debugging considerations include:

- “Changes to the System Trace Facility” on page 6-20
 - “SDWA Changes” on page 6-22
 - “Addressing Mode Reflected in Dumps” on page 6-22
 - “Specifying Reason Codes” on page 6-23
 - “System Termination Facility Wait State Codes” on page 6-23
 - “Exceeding the Region Limit” on page 6-23
 - “Diagnosing Checkpoint/Restart Errors” on page 6-24
-

New and Changed Dump Options

Figure 6-1 summarizes the new and changed dump options for user and system dumps. It indicates whether the option is valid on the SNAP macro, the SDUMP macro, and/or the DUMP command. None of the changes are incompatible. However, the following MVS/370 options produce different dump data in MVS/XA: NUC, TRT, CB, SPLS, and SQA. Some options (for example, NUC and SUBPLST) are designed or changed to eliminate unnecessary dump data. Other options (such as SUM) improve the usefulness of dump data.

Figure 2-3 describes changes to dump options in the IEAABD00, IEADMP00, and IEADMR00 PARMLIB members.

Dump Option	SNAP	SDUMP	DUMP	Data Included in the Dump
ALLNUC (new)		x	x	The entire nucleus, both the DAT-on nucleus and the DAT-off nucleus. (NUC no longer requests the entire nucleus.) Although ALLNUC is an SDATA keyword, specifying SDATA=ALL does not cause the ALLNUC information to be dumped. You must explicitly specify the ALLNUC option to obtain the entire nucleus. Your installation might want to keep one copy of a dump of the page-protected nucleus to use with other dumps. To obtain a dump of only the nucleus, use a DUMP command with the ALLNUC, NOSQA, and NOSUM SDATA options specified.
ALLVNUC (new)	x			The entire DAT-on nucleus. Users cannot obtain the DAT-off nucleus in a formatted dump. However, ALLVNUC causes the system to dump both the DAT-on nucleus and the DAT-off nucleus in SYSMDUMPs. Although ALLVNUC is an SDATA option, specifying SDATA=ALL does not cause the ALLVNUC information to be dumped. When ALLVNUC is specified, dumps also include the PSA and the CVT. Unauthorized users receive only the section of the PSA that is not fetch-protected (locations 0 through 2 K minus 1). Authorized (key 0) users receive the entire PSA.
CB (changed)	x			In SYSABEND, SYSUDUMP, and SNAP dumps, a storage map that contains: <ul style="list-style-type: none"> - The starting storage address - The subpool number - The length of the storage allocated to the task - The storage key - Flag information pertaining to the storage and the owning TCB, if the storage is shared In MVS/370, the CB option causes formatted VSM control blocks to be dumped.
CSA and LSQA (changed)	x	x	x	The specified area below and above 16 Mb. No additional SDATA options are required to include extended storage areas in a dump.
KEYLIST (new)		x		Areas in the subpools listed on the SUBPLIST keyword that have the specified key(s). KEYLIST is only valid when specified with SUBPLST. It allows users to obtain a subset of the specified subpool storage.
NUC (changed)	x	x	x	The DAT-on, non-page-protected section of the nucleus, the PSA, and the CVT. Authorized (key 0) users receive the entire PSA. Unauthorized users receive only the section of the PSA that is not fetch-protected (locations 0 through 2 K minus 1). Users that need the page-protected section must specify either the ALLVNUC or ALLNUC option. In MVS/370, NUC requests the entire nucleus. The change in NUC output is designed to eliminate from dumps those areas of the nucleus that are not expected to change (the page-protected areas). Those areas are normally not required to debug problems. The change was made in such a way that most installations can suppress the page-protected areas without having to respecify dump options.

Figure 6-1 (Part 1 of 2). New, Changed, or Deleted Dump Options

Dump Option	SNAP	SDUMP	DUMP	Data Included in the Dump
SPLS (changed)	x			Subpool storage printed in ascending address order, instead of by ascending subpool number as in MVS/370. To obtain storage by subpool ID, specify the new SUBPLST parameter.
SQA (changed)	x	x	x	No system trace entries are included because the trace buffers have been moved to the TRACE address space. SUM and TRT are the only dump options for including the system trace table in a SNAP/ABDUMP dump. SUNDUMP and TRT are the only options for including it in a system dump. The SQA option requests both the SQA area above and below 16 Mb.
SUBPLST (new)	x	x		Specified subpools, which appear in ascending order, as does the storage contained in each subpool. Another parameter for requesting subpool storage already exists, SPLS. However, SPLS causes all user subpools to be dumped. Also, the storage is printed in ascending address order instead of by subpool number.
SUBTASKS (new)	x			Program data (PDATA) information for subtasks.
SUM (new)	x		x	Summary dumps for abending tasks. See "New User Summary Dumps."
TRT (changed)	x	x	x	The trace data included in dumps is changed. See "Changes to the System Trace Facility." The trace output is formatted the same in SNAP dumps, stand-alone dumps, and dumps printed via print dump. It is different from MVS/370 trace output. As in MVS/370, TRT causes trace data from the active trace facilities to be dumped. In MVS/XA, the dump can include master trace data, system trace data, and GTF data. In MVS/370, system trace and GTF cannot be active at the same time. Therefore, MVS/370 dumps never include both system and GTF trace data.

Figure 6-1 (Part 2 of 2). New, Changed, or Deleted Dump Options

New Symptom Dumps for Task-Mode Abends

When an ABEND or program check occurs and the ABDUMP module gets control, the system automatically issues a ten-line symptom dump. The symptom dump, which appears in the job listing, includes:

- The ABEND code and error ID
- The PSW at the time of the error, the instruction length code, and the interrupt code
- The name and address of the active load module, if the PSW points to an active load module
- The offset into the module where the error occurred, if the PSW points to an active load module

- Six bytes of storage before and six bytes after the PSW address at the time of the error
- The register contents at the time of the error

Unless the NOSYM parameter is specified in the appropriate PARMLIB member (IEAABD00, IEADMP00, or IEADMR00), all users get the symptom dump regardless of whether or not they supply a dump DD statement. However, TSO users must specify the WTPMSG parameter on the PROFILE command to see the symptom dump output. Installations that do not want symptom dumps can suppress them by specifying the NOSYM option in the appropriate PARMLIB member or by using the CHNGDUMP command. Specifying NOSYM in the IEADMP00 PARMLIB member also suppresses symptom dumps for users who do not specify dump DD statements.

The symptom dump is designed to help users identify duplicate problems and, in some cases, solve the problems without a full dump. Thus, symptom dumps can reduce the time required for problem determination. Note that symptom dumps are only issued for task-mode ABENDs.

New User Summary Dumps

Users can now obtain a summary dump for abending tasks. A new SDATA keyword, SUM, requests a summary dump. Summary dumps show information that can help identify duplicate problems, followed by control blocks and storage areas. In many cases, the summary dump is sufficient to debug user program checks and ABEND dumps.

If SUM is the only dump option specified, the summary data for SYSUDUMP and SYSABEND dumps has the following format:

1. The dump title.
2. The ABEND code and PSW at the time of error.
3. The name and address of the load module in error.
4. The offset into the load module where the error occurred.
5. Control blocks (the same as if the CB option were specified).
6. Error control blocks (RTM2WAs and SCBs).
7. Save areas.
8. The general purpose registers at the time of error (from the RTM2WA).
9. The contents of the active load module and the load module associated with the last PRB.
10. One K of storage before and after the addresses pointed to by the PSW and registers at the time of error. The summary dump includes only storage areas for which the caller is authorized.

11. System trace table entries for the dumped ASID. (GTF records are not dumped.)

If SUM is not the only SDATA option, the summary data might be dispersed throughout the dump, depending on the other options specified.

The summary output produced for SYSMDUMP dumps is different from the summary output produced for other user dumps. You must use PRDMP and specify the SUNDUMP verb to get summary output. The output contains the same information and is in the same format as the summary information in SVC dumps.

SUM is a default option in the IEAABD00, IEADMP00, and IEADMR00 PARMLIB members. Unless installations delete the SUM options from those PARMLIB members or suppress them using the CHNGDUMP command, MVS/XA automatically produces a summary dump. In MVS/XA, SUM is the only default option in IEADMP00. Therefore, unless users that specify SYSUDUMP DD statements request additional data, they receive only summary data.

Dump Format Changes

The dump headers in user, SVC, and SYSMDUMP dumps contain additional information to aid in problem determination. Information in user dump indexes is presented differently. The SYSMDUMP and SVC printed summary dumps are restructured.

Changes to User Dump Headers

In MVS/XA, the dump headers in SNAP dumps caused by errors, SYSUDUMP, and SYSABEND dumps contain:

- The name of the load module that was executing at the time of the error.
- The offset into the load module, indicated by the PSW. The offset points the user to the failing instruction or to the next sequential instruction at the time of the error.

The SYSMDUMP dump header contains a new symptom buffer to help users identify duplicate problems or solve problems without a full dump. "New Symptom Dumps for Task-Mode Abends" describes the information contained in the symptom buffer.

User Dump Indexes

The indexes of SYSUDUMP, SYSABEND, and SNAP dumps list alphabetically the name of each active load module and the page number in the dump where it starts.

Changes to SYSMDUMP and SVC Dump Formats

The printed summary dump is restructured to make it easier to use. Following is a summary of the changes:

- Individual control blocks are formatted and dumped as separate logical records with unique IDs.
- The general, unformatted storage areas are printed in ascending address order within an address space.

- The dump index gives the starting page number of all formatted storage areas. It also has entries for each storage area requested.
- Trace data is no longer included in the formatted summary output. The trace table appears in the main body of the dump and can be formatted using IPCS or the TRACE verb of PRDMP.

The dump header records of SVC dumps and SYSMDUMP dumps contain the following additional information to aid in problem determination:

- The name of the active load module at the time of the error, if that information is available in the SDWA.
- The offset into the active load module of the instruction that caused the ABEND (in SYSMDUMP headers only).
- The PSW at the time of the error (in SYSMDUMP headers only).
- Six bytes of storage preceding the PSW address at the time of the error and six bytes following the address (in SYSMDUMP headers only). The failing instruction will be in one of those six-byte areas.
- The current SDWA control block, if available.
- Flags indicating whether SYSMDUMP or SVC dump produced the dump.
- The ID of the processor on which the dump was initiated.
- The name of the dump data set.

Note that the SYSMDUMP dump header record contains all of the information available in user symptom dumps, but in a different format.

Suppressing Dumps

The SLIP command has new operands for selectively suppressing dumps. MVS/XA uses SLIP commands to suppress user and system dumps that are normally not required for problem determination. Release 1.1 introduces dump analysis and elimination (DAE), a new SDUMP function that also suppresses dumps.

New Operands on the SLIP Command for Suppressing Dumps

The SLIP ACTION keyword has new operands, NOSYSA, NOSYSU, NOSYSM, and NOSVCD, that separately suppress SYSABEND, SYSUDUMP, SYSMDUMP, and SVC dumps, respectively. The new operands make the command more versatile. Installations can suppress specific types of dumps for ABENDs without suppressing all types.

MVS/XA's Use of SLIP Commands

MVS/XA uses SLIP commands to automatically suppress user and system dumps for ABEND codes that normally do not require a dump for problem determination. Examples are ABEND codes x'B37', x'D37', x'E37' and x'80A' (out-of-space

ABENDs). The SLIP commands that suppress those dumps are in a new PARMLIB member, IEACMD00.

When a system dump is suppressed, the system puts the SVC dump status code in the appropriate LOGREC entry, as it does in MVS/370. When a user dump is suppressed because of a SLIP command, the system issues a write-to-programmer (WTP) message to inform the user.

If your installation already uses SLIP commands to suppress dumps, compare your list with the SLIP commands in IEACMD00. Delete all unnecessary commands to conserve SQA space. SLIP commands in your PARMLIB member (COMMNDxx) override commands in IEACMD00. Therefore, if any commands in IEACMD00 are undesirable, delete them or add SLIP commands to override them in COMMNDxx. Keep in mind that IEACMD00 might be refreshed with subsequent system updates.

When the system processes IEACMD00 at IPL time, it allocates fixed storage for the SLIP action processors and the control blocks they use. "Fixed Storage for SLIP Command Processors (IEACMD00)" describes the fixed storage requirements.

Dump Analysis and Elimination (DAE)

Dump analysis and elimination (DAE) is a new SDUMP function in Release 1.1 that collects symptom data before taking SYSMDUMP or SVC dumps. DAE uses that data to identify and suppress duplicate dumps. If DAE does not suppress a dump, you see the symptom data in the dump header record.

For each dump type (SYSMDUMP and SVC), you can request that DAE do one or more of the following:

- **UPDATE** - Record the symptom data in the SYS1.DAE data set. If the data already appears in SYS1.DAE (that is, a problem with the same symptoms has already occurred), instead of creating an identical entry, DAE adds one to the incidence count in the existing entry. Incidence counts thus indicate the number of times particular problems have occurred. They appear in dump header records along with the symptom data.

If you do not request updating, the system keeps the symptom data in storage elsewhere, but only until DAE processing is stopped. Consequently, DAE cannot use that data for comparison the next time it is active.

- **MATCH** - Determine if the symptom data matches data already collected for the same dump type. Depending on other criteria described later, the system either suppresses duplicate dumps or reports matches in dump header records.
- **SUPPRESS** - Suppress dumps having symptom data that matches data already collected for the same dump type. Instead of duplicate SYSMDUMPs, users receive message IEA838I, which contains the symptom data. MVS does not notify users when DAE suppresses SVC dumps. If you request dump suppression, the dump header records of dumps that are NOT suppressed indicate why.

Note: Requesting the SUPPRESS option does not guarantee that DAE suppresses all duplicate dumps.

- The following parameters on the SLIP command prevent DAE from suppressing dumps: SVCD (SVC dump), TRDUMP (trace dump), and NOSUP (no suppression). NOSUP is new in Release 1.1. See the SLIP entry in Figure 4-2 for more detail.
- DAE does not suppress a dump unless the recovery routine that calls RTM to take the dump: (1) provides in the ABDUMP symptom area, SDWA, SDWAVRA, or SDWA extensions the symptom information that DAE requires, and (2) sets to 1 the VRADAE key in the SDWAVRA.

The following components' recovery routines are changed in Release 1.1 to allow DAE to suppress dumps they produce: DAE, allocation, converter/interpreter, display dump command processor, DUMPDS command processor, and scheduler JCL facility in Release 1.1; contents supervision, global resource serialization, SRM, TRACE, and VSM in Release 1.2.

Controlling DAE Processing

You control DAE processing via records in the new ADYSETxx PARMLIB members. Each record can specify:

- Whether DAE is to be started or stopped
- The type of processing DAE is to perform for each dump type (UPDATE, MATCH, SUPPRESS, or a combination)
- How many dump entries DAE can store in the SYS1.DAE data set

IBM supplies three ADYSETxx members in Release 1.1: ADYSET00, ADYSET01, and ADYSET02. ADYSET00 and ADYSET02 are the same. Both:

- Start DAE processing.
- Request UPDATE, MATCH, and SUPPRESS processing for SYSMDUMPs and UPDATE and MATCH processing for SVC dumps.
- Allow DAE to store up to 400 entries in SYS1.DAE.

ADYSET01 stops DAE processing. If you require different options, create additional ADYSETxx members or modify ADYSET02.

A new SET DAE=xx command specifies which ADYSETxx member MVS/XA is to use. Only one member can be in effect at a time. However, you can issue a SET DAE command at any time to change DAE processing.

The IEACMD00 PARMLIB member shipped with Release 1.1 includes the command SET DAE=00. Thus, unless you change the defaults, during initialization, the system automatically starts DAE processing with the options specified in ADYSET00.

You also have some control over:

- The symptoms DAE collects for each dump type
- The minimum number of symptoms that must match before DAE considers the dump a duplicate
- The minimum number of bytes of meaningful data each symptom must contain before DAE can use it

A new non-executable load module, ADYDFLT, contains default symptoms and specifies which are required and which are optional. It also establishes the minimum number of bytes the symptoms must contain. *System Modifications* describes how to change those defaults using the VRADATA macro.

Actions to be Taken

Before performing an IPL:

- Create a SYS1.DAE data set. If SYS1.DAE is not cataloged at IPL time, the operator receives a message stating that attempts to start DAE processing failed.

System Modifications describes how to create SYS1.DAE using JCL in the DAEALLOC member of SYS1.SAMPLIB. Consider allocating SYS1.DAE with DISP=SHR so you can browse, add, or delete records in the data set. You might, for example, want to delete information related to a problem after applying service to fix it. Note also that you cannot share SYS1.DAE among systems.

- Ensure that an ADYSETxx member specifies the desired DAE options.
 - Either ensure that the IEACMDxx member MVS/XA uses contains the appropriate SET DAE=xx command, or instruct the operator to issue that command.
-

New and Changed PRDMP Control Statements

The following figure describes the print dump (PRDMP) control statements that MVS/XA adds, changes, or deletes. For more information, see *SPL: Service Aids*.

Verb	Release			Description of Change
	1.0	1.1	1.2	
ASMDATA	x			Requests auxiliary storage management (ASM) control blocks. ASMDATA no longer formats RSM control blocks, as it did in MVS/370.
DAEDATA		x		Formats and prints the error data (symptoms) that dump analysis and elimination (DAE) collects. A new DAE header exit, ADYHDFMT, puts the same information on the title page. DAEDATA causes PRDMP to repeat the DAE symptoms later in the dump. It is designed mainly for use at the terminal.
DISPLAY	x			DISPLAY is deleted in MVS/XA. Interactive Problem Control System (IPCS) provides equivalent interactive dump scanning functions. See the IPCS information in this chapter and <i>MVS Interactive Problem Control System for MVS/System Product User's Guide and Reference</i> for more information.
FORMAT	x			Has new operands to limit formatting and printing to selected address spaces. The new operands are ASID, JOBNAME, CURRENT, ERROR, and ALL. The CURRENT address spaces are the primary, secondary, home, and CML lock holder's address spaces. When FORMAT is specified without operands, PRDMP formats and prints control blocks from the CURRENT and ERROR address spaces only. In MVS/370, the FORMAT statement has no operands. PRDMP formats and prints control blocks from all address spaces contained in the dump. To get that same output in MVS/XA, you must specify the ALL operand.
			x	Release 1.2 treats the FORMAT control statement as though it were a SUMMARY FORMAT statement. It formats major system control blocks and data associated with each address space in the dumped system.
IOSDATA	x			Requests IOS control blocks.
LPAMAP	x			Requests the names of all modules in the link pack area of the dumped system or on the LPA active queue at the time of the dump. In MVS/370, PRDMP lists only the names of modules on the LPA active queue at the time of the dump. LPAMAP also has new operands, EPA and MODNAME, to indicate how lists are to be ordered. The EPA operand requests sorting by entry point address. MODNAME requests alphabetical sorting by module name. If neither is specified, the printed output includes lists sorted both ways.
MTRACE		x		Formats and prints the master trace table for the dumped system. The console messages appear in the order in which they were issued.
NUCMAP	x			Requests the names of system modules in the nucleus when the dump was taken. The list includes the name, entry point, entry point attributes, and length of each module. Note that the NUCMAP statement does not produce a map of storage in the nucleus. To get a map of storage, use either the PRINT STORAGE= control statement or the AMBLIST utility.
RSMDATA	x			Requests real storage management (RSM) control blocks.
SUMMARY			x	Release 1.2 adds keywords that specify: (1) for which address spaces the system is to collect information, and (2) the information to be collected. The new keywords are the same as the keywords on the IPCS SUMMARY subcommand, which is also changed. See "New and Changed IPCS Subcommands."

Figure 6-2 (Part 1 of 2). New, Changed, or Deleted Print Dump Verbs

Verb	Release			Description of Change
	1.0	1.1	1.2	
TRACE	x			Requests trace data. Programmers must specify the new TRACE verb to obtain the system trace table in PRDMP output. TRACE also allows users to limit the trace data printed. See <i>SPL: Service Aids</i> for a description of trace operands.
VSMDATA	x		x	<p>A new verb that requests virtual storage management (VSM) control blocks.</p> <p>Release 1.2 provides keywords for limiting the output.</p> <p>The following keywords request the VSM control blocks in the private areas of specific address spaces:</p> <p>ASIDLIST - Address spaces whose ASIDs are listed on the ASIDLIST keyword.</p> <p>JOBLIST - Jobs listed on the JOBLIST keyword. You can use JOBNAME as an alias for JOBLIST.</p> <p>CURRENT - Address spaces that were active when the dump was taken.</p> <p>ERROR - Address spaces that terminated abnormally (ACSBCC≠0) or that contain TCBs representing tasks that completed abnormally (TCBCMP≠0 or TCBRTWA≠0).</p> <p>TCBERROR - Address spaces that include TCBs representing tasks that terminated abnormally (TCBCMP≠0 or TCBRTWA≠0). TCBERROR specifies a subset of the information requested using the ERROR keyword.</p> <p>ALL- All address spaces in the system.</p> <p>NOASIDS - Prevents the system from including VSM control blocks from the private areas of any address spaces.</p> <p>You can also limit output using:</p> <p>GLOBAL - Requests the VSM control blocks in the SQA and CSA.</p> <p>NOGLOBAL - Prevents the system from including VSM control blocks in the SQA and CSA.</p> <p>In previous releases, VSMDATA has no keywords. It requests VSM control blocks from the private areas of all address spaces, the SQA, and the CSA. To get that same information in Release 1.2, you must specify the ALL and GLOBAL keywords. If VSMDATA is specified alone, you get only the VSM control blocks that the ERROR, CURRENT, and GLOBAL keywords request.</p>

Figure 6-2 (Part 2 of 2). New, Changed, or Deleted Print Dump Verbs

Print Dump Index

The MVS/XA version of PRDMP writes a dump index and allows user exits to insert their own entries into the index. See *SPL: Service Aids*.

The MVS/XA PRDMP procedure in SYS1.PROCLIB includes a new DD statement requesting that PRDMP write the dump index to a sequential data set other than the PRINTER data set. Using a separate index data set allows you to print the index before the dump.

To have the index printed at the beginning of the dump, either use the PRDMP procedure in the MVS/XA PROCLIB, or add a DD statement in your own PRDMP procedure in PROCLIB. For example:

```
//INDEX DD SYSOUT=A
```

To obtain the index in front of the dump, the INDEX DD statement must precede the PRINTER DD statement. Unless the INDEX DD statement is specified, PRDMP prints the index on the PRINTER data set after the dump.

Print Dump Requirements for Printers

Print dump (PRDMP) output lines are 132 characters long. If your installation uses a printer with a line length of less than 132 characters, you might lose information.

New and Changed IPCS Subcommands

The Release 1.2 level of IPCS adds or changes the IPCS subcommands listed in the following figure. Only the change to EVALDUMP is incompatible.

Subcommand	Description of Change
EVALDEF	The variable list on the CLIST and DIALOG keywords can include SOURCE(variable name). The new term is an alternative to using DATASET(variable name) and DSNAME(variable name).
EVALDUMP	<p>You must use one of the following keywords to specify the dump source: ACTIVE MAIN STORAGE DSNAME(dsname) DATASET(dsname) FILE(ddname) DDNAME(ddname)</p> <p>Previous IPCS releases expect the second positional parameter to be the data set name. Therefore, you need to change programs that use the EVALDUMP subcommand.</p> <p>The formatted output from EVALDUMP now includes the keyword ACTIVE, DSNAME, or FILE before the dump source. Previous IPCS releases formatted the data set name without the keyword and surrounding parentheses. You might have to modify programs because of this change as well.</p> <p>The variable list on the CLIST and DIALOG keywords can include SOURCE(variable name). The new term is an alternative to using DATASET(variable name) and DSNAME(variable name).</p>
LISTDUMP	The format of the output is changed to display the dump source.
OPEN	You can now designate on the TITLE keyword the timestamp IPCS is to include on the top of every print file page. The default timestamp is the time problem analysis started. You might want change it to the time the dump was taken.
RENUM	A new subcommand that renumbers the symbols in the symbol stack so that every numeric suffix between the first and the last symbol is used.
RUNCHAIN	<p>A new EXEC keyword specifies a CLIST statement or an IPCS subcommand that IPCS executes each time it processes a control block in the chain. That is, IPCS locates a control block, processes it as requested, then executes the statement or subcommand on the EXEC keyword before continuing to the next control block. Thus, you can use the EXEC keyword on RUNCHAIN for iterative processing that previously required several statements.</p> <p>You can, for example, use RUNCHAIN to look at the RBs queued from a chain of TCBS. On the RUNCHAIN command that searches the TCB chain, use an EXEC keyword that specifies another RUNCHAIN command, one that searches the RB chain. The system then displays the first TCB, all of the RBs chained to it, the second TCB and its RBs, and so on.</p>
SETDEF	The format of the output is changed to display the dump source.
STACK	A new subcommand that creates a symbol in the symbol stack. It is similar to EQUATE, except that you need not specify the symbol to be created. IPCS always creates the symbol using the next suffix after the largest one used.
Subcommands that specify the dump source	<p>Except for ADDDSN, DELDSN, LISTDSN, and MODDSN, all IPCS subcommands that accept dump data set names are changed to accept these additional dump source keywords:</p> <p>ACTIVE MAIN STORAGE FILE(ddname) DDNAME(ddname)</p> <p>The existing keywords, DSNAME(dsname) and DATASET(dsname), are still valid. For more information, see "Accessing Additional Sources of Dump Data Using IPCS."</p>

Figure 6-3 (Part 1 of 2). New and Changed IPCS Subcommands

Subcommand	Description of Change
SUMMARY	<p>Several new keywords provide additional options that specify: (1) for which address spaces the system is to collect information, and (2) the information to be extracted. Also, some old keywords have been replaced by similar ones that request the same information. These changes are to make the IPCS SUMMARY subcommand like the PRDMP FORMAT and SUMMARY FORMAT verbs, which are also changed. Following are the keyword changes:</p> <p>ALL - Requests information from all address spaces in the system. In previous IPCS releases, if you do not specify from which address spaces IPCS is to take information, the default is all address spaces. In this release, the default is the ERROR and CURRENT address spaces.</p> <p>ASIDLIST - Replaces the ASID keyword. Like ASID, it specifies a list of ASIDs to be processed. The list can include a range of ASIDs, which is not allowed in earlier releases. The system accepts ASID as an abbreviation of ASIDLIST, so the change is compatible.</p> <p>JOBLIST - Replaces the JOB keyword. Like JOB, it specifies a list of job names whose associated address spaces are to be processed. To distinguish JOBLIST from the new JOBSUMMARY keyword, the minimum abbreviation for JOBLIST is JOBL. JOBNAME is also an acceptable substitute. You need to change any programs that specify JOB or J to request JOBLIST information.</p> <p>CURRENT - Requests information from all address spaces that were active when the dump was taken.</p> <p>ERROR - Requests information from address spaces that terminated abnormally (ASCBCC≠0) or that contain TCBS representing tasks that completed abnormally (TCBCMP≠0 or TCBRTWA≠0.)</p> <p>TCBERROR - Requests information from all address spaces that contain TCBS representing tasks that completed abnormally (TCBCMP≠0 or TCBRTWA≠0). TCBERROR specifies a subset of the address spaces that ERROR selects.</p> <p>ANOMALY - Requests different information in Release 1.2 than in previous releases of IPCS. In Release 1.2, ANOMALY requests the same information as TCBERROR. In previous IPCS releases, it requests a subset of the TCBSUMMARY information.</p> <p>FORMAT - A new keyword that requests major system control blocks and data associated with each address space in the dumped system. FORMAT produces the same output as the PRDMP FORMAT and SUMMARY FORMAT verbs.</p> <p>KEYFIELD - Requests key fields in the ASCBs, TCBS, and RBs of the specified address spaces. The output matches the output from the PRDMP SUMMARY verb. If you do not request specific information on the SUMMARY subcommand, IPCS gives you the information that KEYFIELD requests.</p> <p>JOBSUMMARY - Requests the following summary information: <ul style="list-style-type: none"> - A list of active CPUs - Scheduled services - For each address space specified: the jobname, ACSB location, ASID, status of the address space, local service manager queue, local service priority queue, TCB locations, completion codes, and whether or not the TCBS were active at the time of the dump </p> <p>TCBSUMMARY - Produces the same output, but in a different format.</p>

Figure 6-3 (Part 2 of 2). New and Changed IPCS Subcommands

Accessing Additional Sources of Dump Data Using IPCS

The IPCS component of Release 1.2 can access: (1) data sets stored on tape as well as direct access devices, and (2) main storage of the MVS/XA address space in which IPCS is executing. Also, multi-volume data sets no longer need to contain fixed-length records. Earlier releases of IPCS can access only cataloged data sets on direct access devices. Multi-volume data sets previously had to contain fixed-length records.

To access the additional sources of dump data, you can now specify the keywords listed below on most IPCS subcommands that specify the dump source. The only subcommands that do not accept the keywords are ADDDSN, DELDSN, LISTDSN, and MODDSN:

ACTIVE, MAIN, or STORAGE

FILE(ddname) or DDNAME(ddname)

ACTIVE, MAIN and STORAGE all request that information be taken from main storage.

FILE and DDNAME both specify ddnames currently associated with a dump data set. Three subcommands, CLOSE, DROPDUMP, and OPEN, allow more than one ddname in parenthesis.

Four other subcommands, LISTDUMP, SETDEF, EVALDEF, and EVALDUMP, are also changed to support new dump sources. Only the change to EVALDUMP is incompatible. See "New and Changed IPCS Subcommands" for descriptions of the changes.

Support for additional dump sources calls for new rules regarding which dump data sets you can move without invalidating the data set's dump directory entry. You can safely move:

- Data sets containing fixed length records to direct access devices
 - Any data set to a single reel of tape
-

New IPCS Panels

Release 1.2 includes two new IPCS panels, BLSPDISE and BLSPDSLE. BLSPDISE is a top selection panel that, when hooked into the ISPF primary option menu, provides a convenient way of initiating IPCS dialogs. The two selections on BLSPDISE are:

- BROWSE, which invokes BSLDISP, the full-screen dump viewing dialog program.
- IPCS, which displays the other new panel, BLSPDSLE. BLSPDSLE allows you to enter IPCS subcommands.

Your installation might have created similar panels in the past using instructions in the *IPCS User's Guide and Reference*. BLSPDISE is similar to the IPSELCT panel described in the guide. BLSPDSLE is identical to IPCMD, which is also described.

Changes to the IPCS BROWSE Panels

Following are several changes in the way you use the Release 1.2 BLSLDISP panels. Most of the changes are to make BLSLDISP more like the ISPF BROWSE and EDIT panels.

- To identify on entry panels the storage IPCS is to display, use one of the following keywords:

ACTIVE, MAIN, or STORAGE
DSNAME(dsname) or DATASET(dsname)
FILE(ddname) or DDNAME(ddname)

Because IPCS can now access main storage or data sets using a ddname, you can no longer simply specify a data set name.

- Instead of displaying on storage panels repetitive data or blanks for storage that cannot be obtained, IPCS inserts a one-line summary, either:

LENGTH(xxxxx)==> Storage not available
LENGTH(xxxxx)==> All bytes contain X'xx' (or C'c')
LENGTH(xxxxx)==> Same as above

- On pointer and storage panels, you can now use the following primary commands:

STACK Adds a symbol to the symbol stack.

RENUM Renumbers the symbol stack so that every numeric suffix between the first and last symbol is used.

FIND Locates and displays the storage to which the specified value points.

RFIND Repeats the last FIND command. RFIND is disabled in Release 1.2. You can enable it by creating a command table for IPCS, as described in the *IPCS User's Guide and Reference*.

- You can use two new operands, CURSOR and X, on the primary commands you enter on storage panels. CURSOR represents the fullword that the cursor precedes or is under. IPCS treats that fullword as the target address of a command. For example, LOCATE CURSOR% displays the storage beginning at the 24-bit address in the byte the cursor precedes or is under when the command is executed.

X represents the first byte of the displayed storage. IPCS treats the contents of that byte as the target address of the command.

- You can put address space keywords on STACK and LOCATE subcommands. Thus, you can display data from an address space other than the one currently displayed without leaving the storage panel.
- IPCS displays the output from IPCS subcommands and dump processing exits in full-screen mode rather than line mode.
- You can update the dump's symbol stack from more than one logical screen. Previously, if working in split-screen mode, you could update the stack from only one screen.

For more information, see either the *IPCS User's Guide and Reference* or the tutorial panels for the BLSLDISP dialog program. You can access the online tutorial by entering HELP on the command line of any BLSLDISP panel.

Changes to the Titles of IPCS Print Files

Release 1.2 IPCS print files have different default titles. Also, instead of DATE and TIME headings, the first line of each page contains a 17-character timestamp.

The default title is the title in the default dump data set. If no title is available, IPCS uses the old default, "IPCS PRINT LOG FOR userid." As in previous releases, you can override the default by specifying a different title on the OPEN subcommand that opens the print file.

Although the format of the date and time is changed, the default values are still the date and time problem analysis started. In Release 1.2, however, you can specify a different value (for example, the time the dump was taken) on the TITLE keyword of the OPEN subcommand.

Using the MVS/XA Versions of IPCS and PRDMP on Other Systems

To aid in migrating to MVS/XA, IBM allows you to execute MVS/XA versions of IPCS and PRDMP on certain MVS/SP 1.3 systems. You need the MVS/XA versions to view and print MVS/XA dumps. The MVS/370 versions of IPCS and PRDMP can process only MVS/370 dumps; the MVS/XA versions can process only MVS/XA dumps. In fact, PRDMP erases dumps taken on different versions of MVS.

IBM imposes some restrictions on running the MVS/XA modules on MVS/370 systems. The MVS/370 processor must be in a location where MVS/SP Version 2 and MVS/XA DFP are licensed. You can use IPCS and PRDMP on the MVS/370 system up to 18 months after the first shipment of MVS/XA program products to that location. The Agreement for IBM Licensed Programs (Z120-2800) defines the term "location."

The remainder of this topic describes how to obtain the modules and data sets required to run MVS/XA PRDMP and IPCS on an MVS/370 system. You might also want run the PRDMP and IPCS programs from one release of MVS/XA on an earlier release. Although you can use any MVS/XA release of PRDMP or IPCS to print or view dumps taken on another MVS/XA system, to format all control blocks correctly, the level of PRDMP or IPCS must match the level of the dump.

IBM provides the jobstreams required to copy the MVS/XA IPCS and PRDMP modules (except those supporting the EREP PRDMP exit) into a data set you can use on another system. In addition to creating that data set, you need to ensure that the IPCS modules access the correct IPCS/ISPF panel and message libraries. The way you perform these tasks depends on whether you are copying Release 1.2 or earlier levels of IPCS and PRDMP. This topic describes each method separately.

Finally, to obtain the EREP PRDMP exit, you must have EREP Version 2 installed on your MVS/370 system. The EREP PRDMP exit, which is contained in EREP Version 2 instead of in MVS/XA, is required to print MVS/XA LOGREC records. EREP Version 2 runs on both MVS/XA and MVS/370 and can process LOGREC records created on either.

Copying Release 1.2 IPCS and PRDMP Modules and Data Sets

The jobstreams for creating a data set containing Release 1.2 PRDMP and IPCS modules are in several members of SYS1.ASAMPLIB and, after sysgen, SYS1.SAMPLIB.

1. Combine the jobstreams into one member by running the jobstream in the MIGJOB01 member. MIGJOB02 will then contain the combined jobstreams.
2. Replace the data set specification on the SYSLMOD DD statement with the name of your target data set. The default name on the SYSLMOD statement is SYS1.MIGLIB.
3. Edit the JOB statement in MIGJOB02 to reflect your account's requirements.
4. Run the jobstream in MIGJOB02 to create the target data set.

Release 1.2 IPCS modules use panels and messages in two Release 1.2 data sets, SYS1.SBLSPNL0 and SYS1.SBLSMSG0, respectively. Earlier systems with IPCS installed might also have data sets with the same names. Therefore, to ensure that Release 1.2 IPCS uses the Release 1.2 copies of those data sets:

1. Allocate two data sets in which to copy the Release 1.2 panels and messages. Give the data sets names other than SYS1.SBLSPNL0 or SYS1.SBLSMSG0.
2. Copy the Release 1.2 SYS1.SBLSPNL0 and SYS1.SBLSMSG0 data sets into the new data sets.
3. When allocating the data sets required to run Release 1.2 IPCS, concatenate the new data sets in front of the ISPF message and panel data sets and, if included, the SYS1.SBLSPNL0 and SYS1.SBLSMSG0 data sets. (You can omit SYS1.SBLSPNL0 and SYS1.SBLSMSG0 from the concatenation.)

Note: To use the Release 1.2 IPCS dialogs on your MVS/370 system, you must have ISPF Version 2 installed on the system.

The PRDMP procedure for starting Release 1.2 PRDMP on either an MVS/XA or an MVS/370 system is different from the one for starting earlier PRDMP releases. In Release 1.2, PRDMP runs as a command processor under TSO. Therefore, the EXEC and DD statements in the procedure are changed. See "SYS1.PROCLIB Changes" for a listing of the Release 1.2 procedure.

Copying Release 1.0 and 1.1 IPCS and PRDMP Modules and Data Sets

The jobstreams for creating a data set that contains Release 1.0 or 1.1 PRDMP modules and compatible IPCS modules are in the PRDMPXA and BLSAMPLE members of SYS1.ASAMPLIB and, after sysgen, SYS1.SAMPLIB:

1. Replace the data set specification on the SYSLMOD DD statement with the name of your target data set.
2. Edit the JOB statements in each member to reflect your account's requirements.
3. Run both jobstreams to create the target data set. (The PRDMP jobstream copies component analysis routines that IPCS also uses into the target data set.)

Therefore, to use the MVS/XA level of IPCS on MVS/370, you must run both jobstreams.)

The Release 1.0 and 1.1 levels of IPCS use panels and messages contained in the SYS1.ABLSPNL0 and SYS1.ABLMSG0 data sets, respectively. To guarantee access to the panels and messages:

1. Allocate two data sets in which to copy SYS1.ABLSPNL0 and SYS1.ABLMSG0. Give the data sets names other than SYS1.SBLSPNL0 or SYS1.SBLMSG0, because systems with IPCS installed might already have data sets with those names.
 2. Copy SYS1.ABLSPNL0 and SYS1.ABLMSG0 into the new data sets.
 3. When allocating the data sets required to run MVS/XA IPCS, concatenate the new data sets in front of the ISPF message and panel data sets, and if included, the SYS1.SBLSPNL0 and SYS1.SBLMSG0 data sets. (You can omit SYS1.SBLSPNL0 and SYS1.SBLMSG0 from the concatenation.)
-

Debugging Considerations

Changes to the System Trace Facility

The MVS/XA system trace facility is significantly different from the MVS/370 version. The following list summarizes the differences:

- **Flexibility in selecting events to be traced**

The MVS/XA system trace facility can perform explicit tracing, address space tracing, and branch tracing. Explicit tracing records all of the normal system interrupt and dispatch events traced in MVS/370, plus the following:

- New I/O instructions
- Machine check interrupt (MCH)
- Restart interrupt (RST)
- Alternate CPU recovery interrupt (ACR)
- Lock suspension (SUSP)
- Trace options alteration (ALTR)
- User-defined event trace (USRn)

Address space tracing records successfully-executed PC, PT, and SSAR instructions. Branch tracing records successfully executed BALR, BASR, and BASSM instructions. (The system does not, however, trace branch instructions that do not branch out of line, for example, BALR x,0.)

The TRACE command is changed to allow installations to dynamically control which type of tracing is performed. Options are:

- Explicit and address space tracing on, branch tracing off
- All tracing on
- All tracing off

The system treats explicit and address space tracing as a single option. Also, the system can perform branch tracing only when the other trace options are active.

- **System trace is automatically activated**

The system automatically activates explicit and address space (but not branch) tracing at system initialization time. If you prefer other trace options, put an appropriate TRACE command in a COMMNDxx PARMLIB member or issue the command from the master operator console. In MVS/370, installations must use a TRACE command to keep system trace active after system initialization time.

- **Concurrent system and GTF tracing**

System and GTF tracing can be active at the same time on an MVS/XA system. Activating GTF trace no longer turns off system trace, as it does in MVS/370.

Explicit system tracing and GTF tracing record some of the same events. Therefore, if you activate both, you might want to tailor GTF trace to record only events that explicit system tracing does not record. *Diagnostic Techniques* lists the events that system trace records. *SPL: Service Aids* describes the events that GTF trace records and how to tailor GTF trace.

- **The structure, location, and format of the system trace table is changed**

The system trace table consists of queues of trace buffers, one queue for each processor sharing the operating system. The system trace table formatter merges the entries from the separate trace buffers into a single logical table. In MVS/370, the system trace table is a single buffer.

The trace buffers are located in the LSQA of a new TRACE address space. In MVS/370, the system trace table is located in SQA. Moving the trace data reduces the system's use of common virtual storage. It also isolates the trace data from the rest of the system, which provides a greater degree of data integrity.

System trace entries vary in length. MVS/370 entries have fixed length.

- **Installations can control the size of the trace table**

The TRACE command is changed to allow installations to dynamically change the size of the trace table. The default size is 16 K of trace buffers per online processor. The size of the MVS/370 trace table is fixed at IPL time.

- **Installations can create and format their own trace entries**

Installations can use a new macro, PTRACE, to create their own trace table entries. *System Macros and Facilities* describes PTRACE. *Diagnostic Techniques* describes how to create and format user entries.

- **Dumping trace data**

TRT, SUM, and SUMDUMP are the only dump options for including the system trace data in dumps. The SQA option no longer dumps system trace entries because the trace buffers have been moved to the TRACE address space. The system trace data printed in **user dumps** depends on the requestor's authorization. If the requestor is authorized, the dump includes the system

trace table entries for all address spaces. If the requestor is unauthorized, the dump includes only system trace entries from the current address space that were made after the job started. Dumping only job-related trace entries for unauthorized users improves system integrity and makes debugging problem programs easier.

SVC dumps include trace entries for all address spaces. The trace data always appears in the non-summary part of the dump, even when dumped in response to a SUM or SUMDUMP request.

Because the MVS/XA trace data is in separate buffers and the trace entries vary in length, it is not feasible to read unformatted dumps of the trace table. Installations need to use print dump (PRDMP) or SNAP/ABDUMP dump to format trace table entries. The system trace table formatter merges the entries from the separate trace buffers into a single logical table. The formatter merges timestamped entries (explicit trace events) from oldest to newest. It merges branch and address space trace entries, which are not timestamped, in relative order to the timestamped entries.

Note: To obtain formatted trace table entries you must include the new TRACE verb in the PRDMP procedure. The TRACE verb has operands that allow installations to limit the trace information printed. See "New and Changed PRDMP Control Statements."

SDWA Changes

The SDWA has increased in size. All of the additional storage is included in SDWA extensions. The additional storage contains data for I/O machine checks, new locks, new dump tailoring options that specify storage subpool lists, and new service data. The information is contained in the following extensions:

- The previously-existing recordable extension 1 (SDWARC1) contains additional service data.
 - The new recordable extension (SDWARC2) contains I/O machine check data.
 - The new recordable extension (SDWARC3) contains new lock and lockword information that can be specified in the FRELOCK keyword of the SETRP macro.
 - A new non-recordable extension (SDWANRC2) contains the SNAP dump tailoring information for storage subpools.
-

Addressing Mode Reflected in Dumps

When producing summary (SUM or SUMDUMP) dumps, dump routines use the addressing mode at the time of the error to determine whether the addresses in registers are 24-bit or 31-bit values. If a program is running in 31-bit addressing mode when an error occurs, the system treats addresses as 31-bit values. If a program is running in 24-bit mode, the system treats them as 24-bit values. If a program has 31-bit addresses in some registers and switches to 24-bit mode just before an error occurs, the dump routines consider the addresses to be 24-bit values. As a result, dumps at times might not include the desired areas.

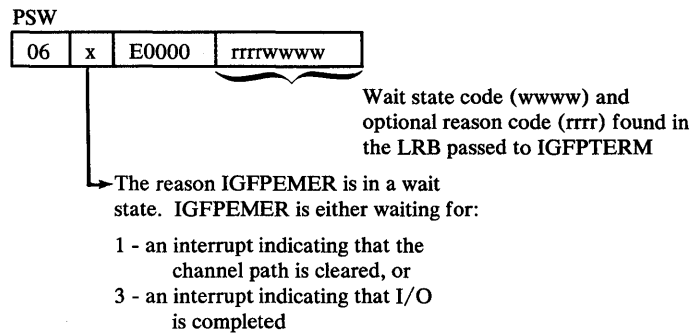
Specifying Reason Codes

Users can specify a reason code on ABEND, CALLRTM, and SETRP macros. The reason code supplements the completion code associated with abnormal termination. It allows users to uniquely identify the cause of abnormal termination. RTM propagates the reason code to each recovery exit and to the TCB and ASCB control blocks. Thus, the reason code appears in system messages.

System Termination Facility Wait State Codes

In MVS/370, the system termination facility (IGFPTERM and IGFPTREC) issues wait state code X'024' when IGFPTREC fails to receive an I/O interrupt while attempting either to write a LOGREC record or to issue a WTO message. Users are prevented from seeing the wait state code of interest, namely the wait state code indicating the error condition that caused system termination processing to begin (the wait state code in the LRB passed to IGFPTERM).

In MVS/XA, the system termination facility puts into the PSW the wait state code and the optional reason code found in the LRB, and a reason code indicating why IGFPEMER is in a wait state. (In MVS/XA, IGFPEMER replaces IGFPTREC.)



Programmers seeing a wait state code in a PSW with this format can locate the message that was to be displayed at the operator's console and the LOGREC record that was to be written to SYS1.LOGREC. At the time the wait state code is loaded, Register 1 points to a 2-word parameter list. The first word contains the address of the WTO message, the second word contains the address of the LOGREC record.

Exceeding the Region Limit

With MVS/XA, installations can use the SMF step initiation exit (IEFUSI) to specify region size and region limits. When IEFUSI changes the VSM region size and region limits, MVS/XA records the change in an SMF type 30 record. Therefore, if a job is cancelled for exceeding the limit when the JCL specified adequate space, check for an SMF record indicating that IEFUSI changed the limit. See "Limiting User Region Size using IEFUSI Instead of IEALIMIT" in Chapter 5 for a description of the IEFUSI enhancements.

Diagnosing Checkpoint/Restart Errors

If an internal error occurs in Release 1.2, checkpoint/restart puts diagnostic information into the SDWA for recording in SYS1.LOGREC. Checkpoint/restart also issues an SVC dump to store selected dump information in a SYS1.DUMPxx data set. The dump information includes:

- All storage currently allocated to checkpoint/restart
- Four K of storage on each side of each register
- Load modules
- The SDWA

Dumps obtained using SYSUDUMP or SYSABEND DD statements are not useful for solving problems in checkpoint/restart. Use PRDMP instead to print dumps that checkpoint/restart creates.

Chapter 7. Accounting

This chapter contains information pertaining to accounting procedures. In general, converting to MVS/XA does not require that you change your accounting programs significantly, if at all. You do, however, need to examine accounting programs to determine whether they will:

- Execute successfully in MVS/XA. Most SMF records are the same or compatibly expanded in MVS/XA. Therefore, in many cases, accounting programs will work unchanged.
- Bill jobs the same whether executed on MVS/XA or MVS/370. After MVS/XA is installed and operating, you can perform comparison runs between your present system and MVS/XA. Depending on the results, you might need to adjust your accounting or billing algorithms.

Although SMF reports additional measures of I/O activity and virtual storage use, it continues to report the old data as well. In most instances, the data is also derived the same way. Processor or CPU utilization times and EXCP counts for application data sets are calculated as in MVS/370. EXCP counts for program libraries, however, are slightly different. "Increases in EXCP Counts for Program Fetch Activity" describes the differences.

The topics in this chapter describe some changes that might affect existing accounting programs. Most of the information, however, describes new measurements you will want to use in the future, after your MVS/XA system is stabilized. The topics included are:

- "Device Connect Time"
- "New Fields Measuring Virtual Storage Use" on page 7-2
- "SMF30PRV and SMF30SYS Fields" on page 7-2
- "Type 22 SMF Record Changes" on page 7-3
- "Increases in EXCP Counts for Program Fetch Activity" on page 7-3
- "Summary of SMF Record Changes" on page 7-4
- "SMF Compatibility Between Release 1.0 and Later Releases" on page 7-5

Device Connect Time

In addition to the EXCP counts available in MVS/370, SMF accumulates device connect time for each data set defined by a DD statement, for each address space, and for each command issued during a TSO session. Device connect time is similar to channel busy time in MVS/370. It measures the amount of time during an I/O operation that the channel subsystem is transferring data or control commands (such as SEEK) on the channel path. Device connect time is a more accurate measure of actual device use than the number of physical blocks transferred (the EXCP count).

Type 30 and 32 SMF records include new fields for reporting device connect time. In type 30 records, the SMF30DCT field in the EXCP section indicates the device connect time for a data set. The SMF30TCN field shows the total device connect time for the address space. The SMF32TCT field in type 32 records reports the total device connect time used while executing a command during a TSO session.

If you currently obtain EXCP counts from type 4, 5, 34, or 35 records and plan to use device connect time in MVS/XA, consider modifying those programs to obtain EXCP counts from type 30 and 32 records instead. Changing the programs now might ease the transition later. Device connect time is not reported in the other records mentioned.

New Fields Measuring Virtual Storage Use

The storage and paging section of type 30 SMF records includes new fields that report virtual storage use above and below 16 Mb. Eventually, you might want to modify accounting routines that measure virtual storage to use the new data. Many system control blocks have moved to virtual storage above 16 Mb. Also, user programs will begin using storage above 16 Mb.

The new fields report:

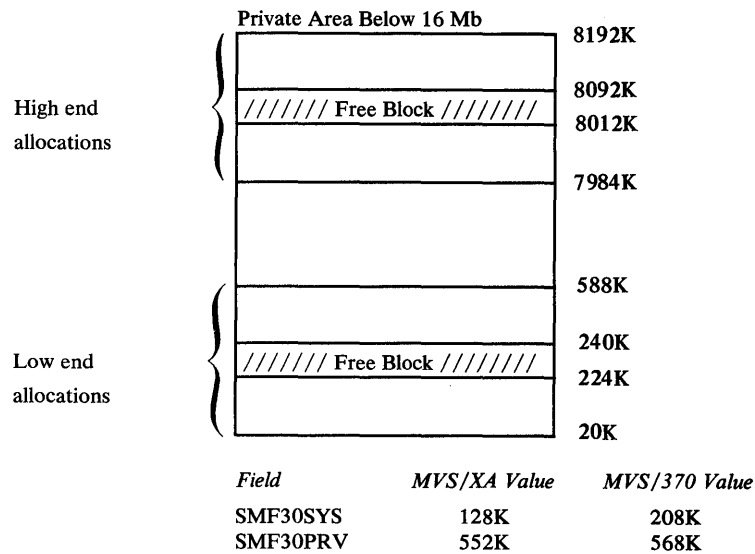
- The region size below and above 16 Mb (SMF30RGB and SMF30ERG)
 - The maximum amount of virtual storage allocated from the LSQA and SWA subpools below and above 16 Mb (SMF30ARB and SMF30EAR)
 - The maximum amount of virtual storage allocated from the user subpools below and above 16 Mb (SMF30URB and SMF30EUR)
-

SMF30PRV and SMF30SYS Fields

The SMF30PRV and SMF30SYS fields continue to report private area use below 16 Mb. However, MVS/XA calculates the source for the fields differently than MVS/370. In MVS/XA, SMF30SYS and SMF30PRV show the total number of bytes (in 1 K units) allocated from subpools in the high and low ends of the private area, respectively. The amounts do not include imbedded free blocks.

The MVS/370 values report the total number of bytes between the highest and lowest addresses allocated from subpools in the high and low ends, respectively. The amounts include any free blocks imbedded in the respective ranges. Therefore, the MVS/XA values might be lower than the MVS/370 values for the same job.

The following picture illustrates the differences between the MVS/XA and MVS/370 values.



Two new fields in the type 30 record (SMF30ARB and SMF30URB) report the same data as the MVS/XA SMF30SYS and SMF30PRV fields, but in bytes instead of 1 K units. SMF30ARB is equivalent to SMF30SYS. SMF30URB is equivalent to SMF30PRV.

Type 22 SMF Record Changes

MVS/XA replaces the channel section of type 22 records with channel path information. Therefore, you must at least reassemble existing programs that use type 22 records. You might also want to modify the programs to use the new channel path data.

Increases in EXCP Counts for Program Fetch Activity

The EXCP counts that SMF records for program fetch activity are likely to be higher in MVS/XA than in MVS/370. In both systems, SMF records EXCP counts in either the SMF30TEP field or, if a STEPLIB is used, in the SMF30BLK field or the equivalent SMF4EXCP field. (SMF30TEP and SMF30BLK are type 30 SMF records; SMF4EXCP is a type 4 record.) If your installation uses any of these fields to measure program fetch activity, you need to determine if the increase affects your accounting programs.

The higher counts result from program fetch changes described in “Ensuring Optimal Program Fetch Performance” in Chapter 8. MVS/XA records all fetch I/O activity, whereas MVS/370 misses some. (For example, it appears that MVS/370 does not count redrives caused by the need to fix additional storage.)

The MVS/XA Release 1.0 and 1.1 versions of program fetch count actual EXCPs. In systems with Release 1.2 program fetch or its equivalent (the version obtained by installing the PTF for APAR OZ75713 on Release 1.1), the EXCP counts for non-overlay modules with correct RLD count values report the number of text blocks transferred instead of actual EXCPs. These counts are likely to be the same as the actual EXCP counts obtained in earlier MVS/XA releases. For overlay

modules the SMF counts in Release 1.2 and its equivalent are likely to be less than in earlier MVS/XA releases, but more than in MVS/370.

Summary of SMF Record Changes

The following chart summarizes the SMF record changes. For more detail, see *SPL: SMF*.

SMF Record	Release			Description of the Change
	1.0	1.1	1.2	
Type 4 (Step Termination)			x	In Release 1.2, the SMF4RSHO field is increased from 2 bytes to 4 bytes to accommodate a region size greater than 16 Mb. It is moved from the beginning to the end of the storage and paging section to avoid changing the offsets of other fields in the record. You must recompile programs that use the SMF4RSHO field.
Type 6 (JES2 (Output Writer))			x	Includes several new fields at the end of the 3800 Printing Subsystem section. The new fields are meaningful only if the 3800 Printing Subsystem Model 3 is running under a functional subsystem (FSS).
Type 22 (Configuration Record)	x			The channel section of the configuration record is replaced by a channel path section. Also, the format of the storage section is changed to report 31-bit counters and addresses.
Type 30 (Common Address Space Work Record)	x			Includes additional fields described in "Device Connect Time" and "New Fields Measuring Virtual Storage Use." One other new field in the completion section, SMF30ARC, reports the abend reason code. Also, MVS/XA calculates the values in the SMF30PRV and SMF30SYS fields differently. See "SMF30PRV and SMF30SYS Fields."
				x
Type 32 (TSO User Work Accounting Record)	x			Includes the device connect time per TSO user, in addition to the existing data.
Type 34 (TSO Step Termination)			x	In Release 1.2, the TIVEFRGN field is increased from 2 bytes to 4 bytes to accommodate a region size greater than 16 Mb. It is moved from offset 74 to offset 82 to avoid changing the offsets of other fields in the record. You must recompile programs that use the TIVEFRGN field.
Types 4, 14, 15, 19, 30, 34, 40, 64, and 69	x			The adjacent fields containing channel addresses and unit addresses are combined to form a single field for a device number.
Type 4, 30, 34 and 40	x			VIO data sets are designated by the value X'7FFF' in the device address field. MVS/370 uses the value X'0FFF', which is a valid device number in MVS/XA.

Figure 7-1 (Part 1 of 2). SMF Record Changes

SMF Record	Release			Description of the Change
	1.0	1.1	1.2	
Type 70-79	x			<p>The formats of these records have changed. As a result:</p> <ul style="list-style-type: none"> - Installations must use the post processor and report writers shipped with RMF Version 3 to process Version 3 SMF records. - Installations must modify user programs that process SMF records 70-79. <i>RMF Reference and User's Guide</i> describes the new format. Most fields contain the same information in Version 3 as they did in Version 2. - The RMF Version 3 post processor can process SMF records written by RMF Version 2 Release 2.2 (SE2 support) or later levels. The post processor converts the old records to the new format before processing them. Installations cannot assume that all data in the old records appears in the converted records. Because the contents of some records are changed (particularly those dealing with I/O operations), the post processor omits some data from the old records. The meanings of other fields have changed. - Installations that require data in SMF records written by an earlier level of RMF than RMF Version 2 Release 2.2 need to keep both the records and the earlier level of the post processor.
All records	x			Bit 5 in the header is set to 1. The flag allows users to distinguish MVS/XA records from MVS/370 records.

Figure 7-1 (Part 2 of 2). SMF Record Changes

SMF Compatibility Between Release 1.0 and Later Releases

To improve performance, Release 1.1 changes SMF data set handling in two ways:

- When formatting an SMF data set, SMF fills it with dummy records instead of binary zeros, as before. The dummy records are shorter than any valid SMF record and contain the characters 'SMFEOFMARK'. The SMF dump program, IFASMFDP, recognizes dummy records and terminates processing when it encounters one. Thus, IFASMFDP no longer reads to the physical end of file when processing partially filled data sets.
- SMF uses a binary search rather than a linear search to find where to start recording in a partially full data set. A binary search reduces the number of control intervals read before finding the starting point.

The Release 1.1 changes are compatible. Release 1.0 and later versions of IFASMFDP can read data sets with or without the SMF EOF marks. Although you can use the same SMF data sets for all levels of the system, you need to consider the following points:

- If you use the CLEAR or ALL options when running IFASMFDP, SMF formats the SMF data sets according to the release level of IFASMFDP.
- When processing data sets that contain SMF EOF marks, the Release 1.0 level of IFASMFDP ignores the marks and reads every control interval to the end of the data set. There is no performance gain. The Release 1.1 and later levels of IFASMFDP recognize the SMF EOF marks and terminate processing. You see the most significant performance gain when processing a large data set that contains SMF EOF marks and is almost empty.

- When processing data sets that do not contain SMF EOF marks, all levels of IFASMFDP read every control interval to the end of the data set. There is no performance gain.
- If you IPL a Release 1.0 system using SMF data sets that contain SMF EOF marks, the data sets appear full during normal data set selection. The operator must dump and clear at least one data set before SMF can begin recording.
- To find where to begin recording, the Release 1.1 level of SMF performs a binary search, regardless of the data set's format. The Release 1.0 level always performs a linear search. In neither case are any records lost.
- SMF initialization does not reformat any data set unless:
 - The data set has a bad control interval that caused an I/O error during the previous IPL. The data set was taken out of service after the error occurred, but the control interval is still bad.
 - The data set has just been allocated and has not been formatted.

In summary, you probably want to use the Release 1.0 level of IFASMFDP until your more current system is in production. You thereby avoid the situation where all SMF data sets appear full and, when running Release 1.1, you still have the advantage of a binary search. You will not, however, see any performance improvements during SMF data set processing.

Chapter 8. Measurement and Tuning

This chapter includes topics related to performance:

- “Ensuring Optimal Program Fetch Performance”
- “Using a New Directory for LNKST Data Sets” on page 8-7
- “SMF Data Set Placement” on page 8-9
- “Using the ASM Backing Slot Function” on page 8-9
- “Using Residency Time to Calculate the Page-in Rate of an Address Space” on page 8-9
- “Changes to ASM’s Paging Algorithms” on page 8-9

Ensuring Optimal Program Fetch Performance

Program fetch is rewritten in MVS/XA and further modified in Release 1.2. All MVS/XA versions of program fetch can fetch the same modules as the MVS/370 version. However, for optimal performance, the Release 1.2 level of program fetch requires:

- A count value for each text block. The count value is the number of relocation dictionary (RLD), control, and RLD/control records associated with the text block.
- Text blocks as large as the linkage editor allows for the output device.

“Recommended Actions” later in this topic describes how to modify program libraries to attain optimal fetch performance. The changes have no effect on the MVS/370 fetch process. The programs you can use to insert count values and reblock modules are:

- The linkage editors supplied with MVS/XA DFP, MVS/370 DFP, and DFDS 1.4¹.
- The MVS/XA DFP and MVS/370 DFP versions of IEBCOPY. The IEBCOPY in both MVS/370 DFP Release 1.1 and MVS/XA DFP Release 1.1 can insert count values in and reblock only **non-overlay** modules (modules that are not in an overlay structure). The IEBCOPY in MVS/XA DFP Release 1.2 and Release 1.1 with the fix for APAR OZ75717 installed can insert count values in **both overlay and non-overlay** modules. However, you need to use the linkage editor to reblock overlay modules.

¹ You need to install the fix for APAR OZ57635 on DFDS 1.4 to obtain correct counts. Modules link edited using the DFDS 1.4 linkage editor without the required PTF installed might contain incorrect counts. Incorrect counts have no effect on the MVS/370 fetch process. However, they degrade fetch performance in MVS/XA. If MVS/XA program fetch encounters an incorrect count value, it issues message CSV300I and continues without using count values.

The steps described in “Recommended Actions” correct any incorrect counts. If you take those actions, you need not separately search for and link edit modules with incorrect counts.

Following are descriptions of how the new linkage editors and the MVS/XA Release 1.2 versions of IEBCOPY and program fetch work. You can obtain equivalent levels of IEBCOPY and program fetch on a system running MVS/XA Release 1.1 by installing PTFs for the following APARs:

- OZ75713 - Replaces program fetch
- OZ75717 - Changes the ALTERMOD and COPYMOD functions
- OZ76136 - Changes a sysgen macro to ensure that future sysgens correctly include the new version of program fetch

Performance Related Changes to the Linkage Editor and IEBCOPY

The linkage editor and IEBCOPY programs identified earlier record the number of relocation dictionary (RLD), control, and RLD/control records following each text block. They put the record count following the first text block in the load module's PDS directory entry. They record the counts for subsequent text blocks in the RLD/control or control record immediately preceding the text block.

Because the count values are located in existing fields that neither MVS/370 program fetch nor previous linkage editors use, load modules containing count values are downward compatible.

Performance Related Changes to Program Fetch

Release 1.2 program fetch reads in non-overlay modules differently from the way it reads in overlay modules.

Fetching Modules that Are Not in an Overlay Structure

If valid counts are available, Release 1.2 program fetch reads one text record and up to 48 associated RLD, control, or RLD/control records using a single I/O operation. Program fetch uses PCIs to dynamically chain additional read operations to the channel program whenever possible. The PCI processing in Release 1.2 involves less disabled time than the PCI processing in MVS/370.

When the count values are invalid or missing, program fetch issues one I/O request for each text record and the first RLD or control record that follows, and one I/O request for each additional RLD, control, or RLD/control record. Therefore, fetch performance in MVS/XA depends on:

- Whether or not valid count values are present.
- The size of each text block. It is best to have block sizes as large as the linkage editor allows for the device type. The larger the block size, the more time program fetch has to chain additional read requests to the currently executing channel program. Chaining read requests improves performance by eliminating the need to:
 - Initiate separate I/O requests
 - Perform SEEK operations if the access mechanism has been repositioned
 - Re-establish the rotational position required to begin the read operation

Fetching Modules that Are in an Overlay Structure

Release 1.2 program fetch reads each text record and the associated RLD, control, or RLD/control records of an overlay module with one I/O operation. However,

program fetch does not use PCIs to chain the read operations when fetching overlay modules. Therefore, the text block sizes in overlay modules have a greater effect on performance--the larger the text blocks, the fewer I/O operations required to fetch the overlay segment.

Recommended Actions

You can improve fetch performance by inserting count values in modules that lack them, and by reblocking modules. The method you choose for making the changes depends on whether you are changing modules in an overlay or non-overlay structure.

Changing Modules that are Not in an Overlay Structure

You can update non-overlay modules using new operations that the MVS/XA DFP or MVS/370 DFP versions of IEBCOPY provide:

- ALTERMOD simply inserts count values.
- COPYMOD copies modules from one library to another. In the process, it inserts count values and reblocks the modules.

Using IEBCOPY with the COPYMOD parameter produces a new data set. Therefore, after copying the modules, you need to scratch the original data set and rename the new one.

The primary candidates for reblocking are:

- SYS1.LINKLIB
- SYS1.COMDLIB
- Program libraries used by interactive applications (for example, CICS and IMS, provided those programs use the standard program fetch)

Reblock the system libraries after constructing the system.

When using the IEBCOPY COPYMOD statement, you need to consider two parameters, MAXBLK and MINBLK, which specify the maximum and minimum block sizes IEBCOPY can create.

- Take the default MAXBLK value to obtain the largest block sizes the linkage editor supports for the device type.
- Use a MINBLK value of 1 K. The initial default value for MINBLK is 1 K; however, your installation might have changed it. *MVS/XA Utilities* describes how to reset MINBLK.

Setting a small MINBLK default value might seem like a contradiction. However, the MINBLK value affects only the size of the last data record on a track. Because of the way program fetch chains read requests across tracks, that record can be small without negatively affecting program fetch performance.

You can also update modules by link editing them using any of the linkage editors identified earlier. Unless you need to link edit a module for other reasons, however, using IEBCOPY is easier.

Changing Modules that are in an Overlay Structure

As stated earlier, you can improve fetch performance by inserting count values in modules that lack them, and by reblocking modules. To insert count values in overlay modules, you can use the same techniques as described for non-overlay modules. However, the IEBCOPY COPYMOD function does not reblock overlay modules. It only copies overlay modules and inserts count values. The IEBCOPY ALTERMOD function works the same for modules in overlay and non-overlay structure.

To reblock and insert count values in overlay modules, you can use one of the linkage editors identified earlier. Note that, to link edit overlay modules, you must provide the link edit control statements required to create the overlay structure. Use the maximum block size the linkage editor allows for the device type.

Increasing the Size of the Page-fixed Area

Some MVS/370 installations improve fetch performance by increasing the amount of virtual storage program fetch fixes at one time. They make the change by adjusting a constant value within the page fix program. Because MVS/XA program fetch fixes 96 K at one time, the equivalent modification is not required in MVS/XA. (MVS/370 program fetch fixes 18 K. MVS/370 DFP program fetch fixes 64 K.)

Maintaining Count Values and Optimal Block Sizes

To maintain count values and optimal block sizes when link editing the modules you modify, always use one of the linkage editors listed earlier. In addition, ensure that the linkage editor constructs the largest possible block size for the device being used. "Assembling and Link Editing Programs" in Chapter 8 describes additional maintenance considerations. The following figure summarizes how different versions of program fetch, IEBCOPY, and the linkage editor handle modules with and without count values.

INPUT	PROGRAM	OUTPUT/COMMENTS
A load module without count values	MVS/XA DFP linkage editor	A load module with count values inserted. Depending on the JCL used and other linkage editor constraints, the text records might also be reblocked.
	MVS/370 DFP linkage editor	
	DFDS 1.4 linkage editor with the required PTF installed	A load module without count values. Depending on the JCL used and other linkage editor constraints, the text records might also be reblocked.
	Earlier versions of the linkage editor	
	The ALTERMOD or COPYMOD functions of the IEBCOPY program in MVS/XA DFP Release 1.2 or Release 1.1 with the required PTF installed	A load module with count values inserted. If COPYMOD is used, a non-overlay module might also be reblocked.
	The ALTERMOD or COPYMOD functions of the IEBCOPY program in MVS/370 DFP Release 1.1	A non-overlay load module has count values inserted; an overlay module does not. If COPYMOD is used, a non-overlay module might also be reblocked.
	Earlier versions of IEBCOPY	A load module without count values. The module is not reblocked.
	MVS/XA program fetch	In some cases, you might observe program fetch performance degradation.
MVS/370 program fetch	No change.	
A load module with count values	MVS/XA DFP linkage editor	A load module with count values inserted. Depending on the JCL used and other linkage editor constraints, the text records might also be reblocked.
	MVS/370 DFP linkage editor	
	DFDS 1.4 linkage editor with the required PTF installed	A load module without count values. Depending on the JCL used and other linkage editor constraints, the text records might also be reblocked.
	Earlier versions of the linkage editor	
	The ALTERMOD or COPYMOD functions of the IEBCOPY program in MVS/XA DFP Release 1.2 or Release 1.1 with the required PTF installed	A load module with count values inserted. If COPYMOD is used, a non-overlay module might also be reblocked.
	The ALTERMOD or COPYMOD functions of the IEBCOPY program in MVS/370 DFP Release 1.1	A non-overlay load module has count values inserted; an overlay module does not. If COPYMOD is used, a non-overlay module might also be reblocked.
	Earlier versions of IEBCOPY	The count values remain. The module is not reblocked.
	MVS/XA program fetch	Depending on the text record lengths, program fetch might perform at its best.
MVS/370 program fetch	No change.	

Figure 8-1. Processing Load Modules

Factors Affecting Text Block Sizes

Several factors affect how the linkage editor determines text block sizes:

- The REGION parameter on the JOB and EXEC JCL statements and on the EXEC statements in SYS1.PROCLIB. The default REGION values defined for the installation can also affect the text block size.
- The values specified for the 'SIZE=(value1,value2)' parameter on the LKED EXEC statement. The SIZE values specify the amount of virtual storage the linkage editor is to use.
- The block size of the previously allocated output library identified on the SYSLMOD DD statement.
- The DCBS option on the PARM parameter of the LKED EXEC statement. Using DCBS allows you to override the block size originally specified for the output data set.
- The block size of the intermediate data set (the data set named on the SYSUT1 DD statement). The linkage editor determines the intermediate data set's block size based on several factors, including the device type.
- The block size of the:
 - Primary input data set (named on the SYSLIN DD statement)
 - Automatic call library (named on the SYSLIB DD statement)
 - The diagnostic output data set (named on the SYSPRINT DD statement)
- The DC option on the LKED EXEC statement. DC causes the linkage editor to construct text blocks of 1 K or less.
- The sizes of the control sections (CSECTs) and named common areas being combined into one load module. When building a text record, the linkage editor puts multiple CSECTs and named common areas into the same record, until it runs into a CSECT or named common area that does not completely fit. The linkage editor then truncates that text record and begins a new one. It never splits CSECTs or named common areas across text records that contain other CSECTs or named common areas.

That restriction also applies if a CSECT or named common area is larger than the maximum text block allowed. The linkage editor does not put any other CSECT or named common area in the last text record occupied by the large CSECT or named common area. Because of this restriction, text records are not always uniform in size or as large as the linkage editor allows for the output device.

The linkage editor uses all of these controls to determine the maximum block size. Poorly chosen values can force the linkage editor to build text blocks smaller than necessary. Therefore, you need to carefully consider any that you specify. The *Linkage Editor and Loader* describes how the linkage editor determines block size and how you can control it.

Following are some of the more common reasons less-than-optimal block sizes are produced:

- The REGION value is too small. Unlike the MVS/370 linkage editor, the MVS/XA version does not execute in an overlay environment. Therefore, it requires 32 K more virtual storage than the MVS/370 linkage editor. Because of the additional storage requirement, the link edit step might fail or the linkage editor might be forced to build smaller text blocks than would the MVS/370 linkage editor.
 - The values specified on the SIZE parameter cause an inadequate output buffer length.
 - The intermediate data set (SYSUT1) supports a smaller maximum record length than the output data set (SYSLMOD).
 - The data set was copied from a different type of device and not link edited again. For example, a data set was copied from a 3330 device to a 3350 device. (Text records cannot exceed 12 K on 3330 devices; 3350 devices allow 18 K records.)
 - The SYSLMOD DD statement specifies a less-than-optimal BLKSIZE value.
-

Using a New Directory for LNKLST Data Sets

A new LNKLST lookaside (LLA) function in Release 1.1 creates and maintains a directory of modules in the LNKLST concatenation. BLDL can use the new directory instead of the PDS directories or the BLDL table to locate modules in the LNKLST concatenation. Because the new directory is hashed and resides in the new LLA address space, using it has several advantages:

- You no longer need to tune the LNKLST concatenation for optimal performance, nor do you need to maintain BLDL lists. The order in which data sets are concatenated does not affect the time required to search hashed directories. Because the new directory is in storage, BLDL lists are unnecessary.
- The new directory eliminates the channel and device contention that occurs when searching PDS directories.
- Data sets in the LNKLST concatenation no longer have to be APF authorized. Consequently, you can include unauthorized data sets formerly included in STEP, JOB, and TASK libraries. The procedure for including unauthorized data sets is described later.
- The LNKLST concatenation can include up to 123 data sets. Earlier MVS releases allow a maximum of 16.
- You can update the LLA directory without performing an IPL. Adding or changing entries in the BLDL table requires an IPL. New commands for updating the directory are described later.
- You can control the amount of paging done for the LLA directory by putting the LLA address space in a separate SRM performance group and adjusting its working set size.

Starting the LLA Function

You might not have to do anything to start the LLA function. The IEACMD00 PARMLIB member shipped with Release 1.1 contains a new command, START LLA, which starts a new LLA procedure in SYS1.PROCLIB. The new procedure, in turn, causes the system to build and begin using the LLA directory.

If you use an IEACMDxx member other than IEACMD00, ensure that it includes a START LLA command. Also ensure that the SYS1.PROCLIB data set you use includes the LLA procedure. In addition, if you want to include unauthorized data sets in the LNKST concatenation, you must specify the new LNKAUTH system parameter, as described in the following topic. The default is to treat all modules fetched via the LNKST concatenation as APF authorized, which is consistent with earlier releases.

Including Data Sets that Are Not APF Authorized

The new LNKAUTH system parameter has two values:

- LNKAUTH=APFTAB - The system treats only those data sets named in the APF table as APF authorized.
- LNKAUTH=LNKST - The system treats all data sets in the LNKST concatenation as APF authorized, regardless of whether their names are in the APF table. LNKST is the default.

To include data sets that are not APF authorized in the LNKST concatenation:

- Either include LNKAUTH=APFTAB in the appropriate IEASYSxx PARMLIB member or have the operator specify it when prompted for system parameters.
- Include in the APF table all LNKST modules to be APF authorized.

Note that the APF authorization established at IPL time remains in effect for the duration of the IPL, even if the LLA function is stopped.

Updating the LLA Directory

To add or change an entry in the new directory, either:

- Issue a MODIFY LLA,REFRESH command to refresh the directory.
- Stop the LLA function by issuing the STOP LLA command, then build a new directory by issuing a START LLA command.

Unless your installation shares LNKST data sets among multiple systems, the first method is preferable. The system can refresh the directory without interrupting its use. Stopping the LLA function causes BLDL to search the PDS directories instead of the LLA directory, which can degrade performance.

If more than one system shares the LNKST data sets, the second method might be better. It allows you to synchronize directory updates. Operators stop the LLA function on all systems, then restart it via the LLA START command.

SMF Data Set Placement

If the device on which an SMF data set resides requires intervention, SMF can generate a large backlog of records while the device is unavailable. Changes in Release 1.1 can alleviate the backlog, provided you do not put all SMF data sets on the same device. If a system with Release 1.1 installed detects that SMF is not writing from buffers, it attempts to use another SMF data set. However, moving to another data set solves the problem only if that data set is on another device.

Using the ASM Backing Slot Function

In Release 1.2, the constants that control the number of slots ASM reserves as back up for each new address space or VIO data set are changed to prevent ASM from reserving any. The changes were made because most installations provide adequate paging space and prefer not to use the backing slot function to limit address space and VIO data set creation.

If your installation wants ASM to reserve backing slots, you need to change the constants. In Release 1.2, the constants are in the ASMSLOTC and ASMSLOTV fields in the ASMVT. ASM uses the ASMSLOTC value when calculating the number of slots to reserve for address spaces. It uses the ASMSLOTV value when calculating the number for VIO data sets. Earlier releases keep the same values in the nucleus CSECTs ILRSLOTC and ILRSLOTV. *Initialization and Tuning* describes how ASM uses ASMSLOTC and ASMSLOTV and how to change them.

Using Residency Time to Calculate the Page-in Rate of an Address Space

If your installation is at the Release 1.2 level, you can request that SRM use residency time instead of execution time when calculating the page-in rate for address spaces in a specified performance group. However, SRM continues to base the page-in rate for cross memory address spaces on elapsed time. Previously, the system used execution time only, except in the case of cross memory address spaces.

Basing the rate on residency time allows the system to decrease the target working set size of an address space while the address space is inactive. Because most installations prefer to maintain minimum working sets for swappable address spaces, requesting residency time calculations is an option mainly for address spaces that are non-swappable.

Basing the rate on execution time protects the frames in the working set while the address space is inactive. The system adjusts the target working set size only while the address space is active. While the address space is inactive, the target size remains the same as when last calculated.

To request that SRM use residency time, use the new IPS parameter, PPGRTR. PPGRTR specifies the high or low limit the rate must exceed before SRM adjusts the address space's working set size. Previously, PPGRT and CPGRT were the only parameters for specifying page-in thresholds.

Changes to ASM's Paging Algorithms

The Release 1.2 level of ASM uses different algorithms for selecting local page data sets and slots on page data sets. The changes are designed to make the paging process more efficient and might result in less tuning effort on your part.

Changes to the Data Set Selection Algorithm

The new data set selection algorithm distributes paging I/O more evenly among local page data sets. ASM continues to maintain the same three circular queues of control blocks representing local page data sets: one for local page data sets on cached auxiliary storage subsystems, one for data sets on fixed-head devices, and one for data sets on movable-head devices. As before, ASM tries to write first to a data set on a cached auxiliary storage subsystem. However, instead of picking the next available data set that contains free space, ASM now also considers the responsiveness of the device and might avoid unresponsive data sets.

ASM begins searching at the data set following the one last selected from the queue. ASM considers each data set on one queue before continuing to the next queue.

Changes to the Slot Selection Algorithm

The new slot selection algorithm tries to reduce device arm movement and seek time by concentrating paging I/O toward the front of the data set.

Chapter 9. Coexistence Considerations

Running both MVS/370 and MVS/XA in the same installation is referred to as coexistence. Installations maintain coexistence because they:

- Have processors that support only MVS/370
- Must use one type of operating system as backup for the other

Most installations will maintain some form of coexistence during the migration period. Many will continue to run both operating systems for some time after MVS/XA is established as a production system.

MVS/370 and MVS/XA can coexist either as independent operating systems running on different processors, as independent operating systems that alternately run on the same 308x processor, or as loosely-coupled operating systems.

In a loosely-coupled JES3 configuration or in a JES2 multi-access spool environment, MVS/370 and MVS/XA must have the same level of JES installed. The JES3 component in MVS/SP Version 1 Release 3.1 with PTFs installed is functionally equivalent to the JES3 component in MVS/SP Version 2. “Installing the JES2 Component of MVS/SP - JES2 Version 2” in Chapter 2 identifies functionally equivalent JES2 components.

In all types of coexistence, the major objectives are to:

- Maintain programs that can run on either system.
- In some cases, ensure that MVS/370 can run the MVS/XA workload or that MVS/XA can run the MVS/370 workload in backup situations.

When MVS/XA and MVS/370 systems are loosely-coupled, installations have some additional considerations, including:

- Ensuring that jobs that must run on a particular system are routed to that system
- Determining which data sets can be shared
- Reviewing DSI procedures

This chapter includes information related to these topics.

Maintaining Programs that Can Run on Both MVS/370 and MVS/XA Systems

Topics in this section describe:

- Instructions for assembling and link editing programs that must run on both MVS/370 and MVS/XA systems
- Criteria for ensuring that programs can run on both systems
- Ways to avoid unnecessary 24-bit dependencies in new programs

- Instructions for using the SPLEVEL macro to generate compatible expansions of fourteen downward incompatible macros
- Two methods of ensuring that programs using the SYNCH macro can run on both MVS/370 and MVS/XA systems

Assembling and Link Editing Programs

In a mixed installation, use Assembler H Version 2 to assemble all programs that use new 370-XA instructions or that are to be run in 31-bit addressing mode. Use the linkage editor in either DFDS Release 1.4 (with the PTF for APAR OZ57635 installed), MVS/XA DFP, or MVS/370 DFP to link edit programs that will be run on the MVS/XA system in 24-bit addressing mode. Using one of those linkage editors is important for fetch performance reasons, as described in Chapter 8. Use the MVS/XA linkage editor to link edit programs that are to be run in 31-bit addressing mode.

The MVS/XA linkage editor is the only one that inserts AMODE and RMODE indicators in CESD entries for load module CSECTs and in the partitioned data set (PDS) entries for load modules. The linkage editors in OS/VS2 MVS and DFDS do not support AMODE and RMODE indicators.

- They do not insert or retain the AMODE and RMODE indicators in the PDS directory entry. If a load module is link edited using one of those linkage editors, any AMODE or RMODE indicators already in the PDS directory entry are deleted.
- The same linkage editors ignore any AMODE or RMODE indicators in the ESD or CESD. Indicators already present remain unchanged. New ones are not inserted. (Assembler H Version 2 and selected HLL compilers, not the linkage editor, insert AMODE and RMODE indicators in the ESD entries of object modules.)

“Establishing a Program’s Addressing Mode” in Chapter 3 describes how AMODE and RMODE indicators are inserted and used in more detail.

Guidelines for Ensuring Program Compatibility

If a program is to run on both MVS/370 and MVS/XA systems, the program must:

- Perform the desired function on both systems. A program might execute without error on both systems, but not produce the desired results (for example, an SMF post processor that analyzes data that has different formats in MVS/XA and MVS/370).
- Use the MVS/370 expansion of macros whose MVS/XA expansions do not work on MVS/370 systems. There are fourteen such macros. “Handling Downward Incompatible Macros” lists them and describes ways of obtaining the appropriate expansion. In some cases, the SYNCH macro is also downward incompatible. See “Downward Incompatible SYNCH Macros” for details.
- Not use new MVS/XA function. New function includes new instructions, new macros, and new parameters, keywords, or options on existing macros. Exceptions are the new LOC, VRC, and VRU parameters on the GETMAIN

macro and the `AMODE=24` parameter on `SYNCH`. Those parameters generate object code that works on MVS/370 systems. See “New Parameters on the `GETMAIN` Macro Instruction” in Chapter 3.

- Use only system services that are supported in both MVS/370 and MVS/XA. Chapter 3 identifies functions not supported in MVS/XA.
- Provide dual paths for functions that are not compatible between MVS/370 and MVS/XA and dynamically select the proper path at execution time. Bit 0 in the `CVTDCB` field of the `CVT` indicates whether MVS/XA is executing. (If it is, bit 0 equals 1.) The MVS/XA `CVT` map defines the bit as `CVTMVSE`. Method 3 in “Handling Downward Incompatible Macros” shows the dual path section of a sample program.

If the program uses non-standard interfaces to system modules or uses system control blocks, you must also ensure that:

- Methods of invoking system services work in both MVS/370 and MVS/XA.
- Control block references can be used in both MVS/370 and MVS/XA.

Some programs require different versions to run in MVS/370 and MVS/XA (for example, `RMF` analysis routines). Installations can either:

- Keep each version of the program in a separate library.
- Keep both versions of the program in the same library, but give each a different name.
- Rewrite the program so that it has dual paths and dynamically selects the proper path at execution time, as mentioned earlier.

Guidelines for Developing New Programs

When designing new programs that must run on both MVS/370 and MVS/XA systems, avoid unnecessary 24-bit dependencies in programs that might be executed in 31-bit mode:

- Use fullword address fields, even if the value in the field is below 16 Mb.
- Avoid using the load address (`LA`) instruction to clear the high-order byte. In 31-bit mode, the `LA` instruction clears only the high-order bit, not the entire byte, as it does in 24-bit mode.
- When coding `BAL` or `BALR`, avoid using the information saved in the high-order byte of the first operand (the instruction length code, program mask, and condition code). When executed in 31-bit mode, `BAL` and `BALR` do not save that information. 370-XA processors provide a new instruction, `IPM` (Insert Program Mask), which saves the program mask and condition code when executing in 370-XA mode.
- Use `ESTAE` instead of `STAE`. `STAE` is not changed to support 31-bit addressing.
- When obtaining large amounts of storage, use the `VRU`, `VRC`, `RU`, and `RC` forms of `GETMAIN` and `FREEMAIN`. These forms support a new `LOC`

parameter, which allows users to specify from where virtual storage is to be obtained and how it is to be backed when fixed. See “New Parameters on the GETMAIN Macro Instruction” in Chapter 3.

In contrast, VSM satisfies the LC, LU, VC, VU, EC, EU, and R forms of GETMAIN requests with virtual storage below 16 Mb. Also, when fixing storage obtained via those forms of GETMAIN, RSM always uses real storage below 16 Mb.

The following macros are not changed to provide full 31-bit support. MVS/XA provides new services instead. You might want to use dual paths when using any of these services:

- PGFIX, PGFREE, PGRLE, PGLOAD, and PGOUT. A new PGSER macro provides the equivalent services and supports 31-bit addresses.
 - SPIE. ESPIE is the MVS/XA counterpart.
-

Handling Downward Incompatible Macros

Most of the MVS/XA expansions of previously existing macro instructions run on both MVS/370 and MVS/XA systems (that is, the macro instructions are downward compatible). The following macro instructions are exceptions. The MVS/XA expansions of these macros will not run on an MVS/370 system.

ATTACH	SETFRR INLINE=YES
CHKPT	SETLOCK RELEASE,TYPE=(reg)/ALL
ESTAE	SMFEXIT
EVENTS	STAX
FESTAE	STIMER
INTSECT	SYNCH, unless it specifies the parameter AMODE=24
SCHEDULE SCOPE=GLOBAL	WTOR
SDUMP if it specifies new parameters	

To share user-written programs among MVS/370 and MVS/XA systems and to have backup capability while migrating to MVS/XA, users must be able to override the downward incompatible macro expansions with macro expansions that will run on both MVS/370 and MVS/XA systems. MVS/XA provides that capability for all of the macros listed, except SYNCH. (See “Downward Incompatible SYNCH Macros” for instructions on maintaining SYNCH compatibility.)

The MVS/XA MACLIB contains two different expansions for all of the above macros, except SYNCH: an MVS/SP Version 1 Release 3 expansion and an MVS/XA expansion. Note that the source statements that invoke the macro instructions remain the same, only the expansions are different for the two environments. The Version 1 expansions run on both MVS/370 systems and MVS/XA systems executing programs in 24-bit addressing mode. The MVS/XA expansion is required when using any new parameters or options on the above macros. In most cases, the MVS/XA expansion is also required when executing in 31-bit addressing mode. (SCHEDULE, SDUMP, and SETLOCK are exceptions.)

The level of the macro expansion (MVS/370 or MVS/XA) that is generated during assembly depends on the value of an assembler language global SET symbol. When the SET symbol value is one, the system generates MVS/370 expansions. When the SET symbol value is two, the system generates MVS/XA expansions.

MVS/SP Version 2 includes a new macro, SPLEVEL, which allows programmers to change the value of the SET symbol. When SPLEVEL itself is assembled, it assigns a value to the SET symbol. That value becomes the default value for the entire installation.

The SPLEVEL macro shipped with MVS/SP Version 2 assigns a SET value of 2. **Therefore, unless a program specifically changes the SET value, the assembler generates MVS/XA macro expansions.**

Your installation can change the SET value shipped with MVS/SP Version 2, or individual programmers can override the SET value in particular programs:

- To change the SET value for the entire installation, after sysgen, modify the SPLEVEL source code in SYS1.MACLIB. Change the statement that assigns the SET value: '&DEFAULT SETC n', where 'n' is 1 or 2. Note that when assembling MVS/XA system programs, either at sysgen or when applying service, the SET value must be 2. (MVS/XA expansions are required.)
- Programmers can issue within a program the SPLEVEL SET=n macro, where n equals 1 to obtain MVS/370 expansions, or 2 to obtain MVS/XA expansions. The SPLEVEL macro sets the symbol to the specified value for that program's assembly only. Thus, issuing the SPLEVEL macro only affects expansions within the program being assembled. A single program can include multiple SPLEVEL macros to generate different macro expansions.

Obtaining the Appropriate Macro Expansions

Following are three ways programmers can use SPLEVEL to obtain the appropriate macro expansion within their programs. Methods 1 and 2 generate different expansions in different programs (for example, MVS/370 expansions in Program A and MVS/XA expansions in Program B). Method 3 generates different expansions within the same program:

Method 1 - Obtaining different expansions in different programs

Keep the SPLEVEL macro shipped with MVS/SP Version 2 in the SYS1.MACLIB macro library. Put a copy of SPLEVEL into another macro library by itself, and change the source code to establish SET = 1 as the installation default. When assembling programs, use JCL to access the appropriate macro library.

In the following example, the SPLEVEL macro that establishes SET = 1 as the installation default is by itself in the SET1MACS macro library.

To assemble the MVS/XA expansions in programs, use:

```
//SYSLIB DD DSN=SYS1.MACLIB,DISP=SHR
```

To assemble MVS/370 expansions, use:

```
//SYSLIB DD DSN=SET1MACS,DISP=SHR
// DD DSN=SYS1.MACLIB,DISP=SHR
```

You can, of course, switch the SPLEVEL macros and put the one that establishes SET = 1 as the installation default in SYS1.MACLIB.

Method 2 - Obtaining different expansions in different programs

Issue the SPLEVEL SET=*n* macro once at the beginning of the module to obtain the appropriate expansions:

```
MODULE      CSECT
            SPLEVEL  SET=1
            .
            .
            .
```

Method 3 - Obtaining different expansions within the same program

Assemble both levels of the macro and make an execution-time test to determine which level to execute. The following example invokes the correct level of the WTOR macro:

```
*   DETERMINE WHICH SYSTEM IS EXECUTING
            TM      CVTDCB,CVTMVSE  (CVTMVSE is bit 0 in
            BO      SP2             the CVTDCB field. It
                                   indicates whether
                                   MVS/XA is executing.)
*   INVOKE THE MVS/370 VERSION
*   OF THE WTOR MACRO INSTRUCTION
            SPLEVEL SET=1
            WTOR      .....
            B         CONTINUE
SP2      DS         0H
*   INVOKE THE MVS/XA VERSION
*   OF THE WTOR MACRO INSTRUCTION
            SPLEVEL SET=2
*
            WTOR      .....
CONTINUE DS         0H
```

Determining Which Level of the Macro Instruction to Use

- You can use a SET value of 2 when you assemble programs that will be run only on MVS/XA systems, regardless of whether they will run in 24- or 31-bit addressing mode.
- Programs that must run on both MVS/370 and MVS/XA systems must assemble at least MVS/370 macro expansions. (Programs with dual paths assemble the MVS/XA expansions as well.)
- Programs that are designed to run on both MVS/370 and MVS/XA systems and that require MVS/XA expansions on the MVS/XA system must obtain both expansions and determine at execution time which level to execute. (See Method 3 in the previous examples.)
- When assembling MVS/XA system programs, either at sysgen or when applying maintenance, the SET value must be 2 (that is, MVS/XA expansions are required).
- Programs that run on MVS/XA systems in 31-bit addressing mode must use the MVS/XA expansions of the following downward incompatible macros:

ATTACH	FESTAE	STAX
ESTAE	INTSECT	STIMER
EVENTS	SMFEXIT	WTOR

- Programs that specify any new parameters, keywords, or options on the macros must use the MVS/XA expansions.
-

Downward Incompatible SYNCH Macros

In some cases, programs that use the SYNCH macro and are assembled using the MVS/XA MACLIB will not run on an MVS/370 system. The downward incompatible programs are those that do NOT specify the parameter AMODE=24 on SYNCH (that is, programs that either omit the AMODE parameter or specify AMODE=31, AMODE=DEFINED, or AMODE=CALLER).

Unless SYNCH specifies AMODE=24, after assembly, the SYNCH macro expansion contains a nonzero value in a previously reserved field in the SYNCH parameter list. The MVS/XA expansion uses that field as an AMODE indicator. The MVS/370 SYNCH processor treats the nonzero field as an error and issues ABEND x'206'.

If a program that issues SYNCH must run on both MVS/XA and MVS/370 systems, you need to ensure that the AMODE indicator field in the SYNCH parameter list is zero. You can either:

- Use the MVS/370 MACLIB when assembling the program.
 - Specify the parameter AMODE=24 on the SYNCH macro and use the MVS/XA MACLIB for assembly.
-

Backup Considerations

Following are backup considerations for installations that must use an MVS/370 system as backup for an MVS/XA system, or use an MVS/XA system as backup for an MVS/370 system:

- For program products that have separate licenses for MVS/370 and MVS/XA, installations need both licenses for backup capability. Such programs include:
 - RMF Version 2 and RMF Version 3
 - MVS/SP Version 1 Release 3 or later releases and MVS/SP Version 2
- User-written programs that access system control blocks or that use authorized services and interfaces might not run on both MVS/370 and MVS/XA systems. If such programs are not compatible between systems, the installation is without backup capability.
- The size of the available private area increases in MVS/XA. Installations that use the additional private area will have to reevaluate the capability of using an MVS/370 system for backup.
- When using the same 308x processor for backup, installations need to ensure that the backup system can use the current system's IOCDS. The 370/370-XA IOCP can create an IOCDS that can be used in either 370 or 370-XA mode. See "Creating a New IOCDS" in Chapter 2 for more information.
- Installations need to have a procedure defined for changing processor modes and for changing the IPL volume.

- On the MVS/370 system, use DFDSS or an equivalent product instead of IEHDASDR or DRWDASDR to produce backup tapes that might need to be restored on an MVS/XA system. Neither IEHDASDR nor DRWDASDR are available in MVS/XA. You must use DFDSS or an equivalent function to perform dump/restore operations in MVS/XA. DFDSS cannot use the dump format that IEHDASDR or DRWDASDR produces.
 - MVS/XA can support a larger workload than MVS/370. An MVS/370 system might not be able to support the MVS/XA workload.
-

Routing Jobs in a Mixed JES2 or JES3 Complex

When MVS/370 and MVS/XA systems are loosely-coupled, installations must ensure that JES routes jobs that must run on a particular system to that system (for example, jobs that use new MVS/XA function).

If a job needs a device that is attached to only one processor, JES3 automatically routes the job to that processor. To ensure the proper job-system match in other situations, installations and programmers can use existing job routing procedures:

- An installation can define specific execution job classes for jobs that must run on MVS/XA, for jobs that must run on MVS/370, and for jobs that can run on either system. The installation then associates each job class with the appropriate processor. Users ensure that their jobs run on the appropriate system by specifying the CLASS= parameter on the job's JCL. JES2 users specify the CLASS= parameter on the JOB statement. JES3 users specify it on the `//*MAIN` or `//*JOB` statement.

Note: Routing by job class works only in situations where processors are always running the same operating systems when the job routing takes place. For example, if Job A specifies CLASS=XA, the processor associated with class XA must always be running MVS/XA when Job A executes.

- If a user knows which operating system is running on a processor, the user can specify on the job's JCL the processor on which the job is to run. JES2 programmers use the SYSAFF parameter on the `/*JOBPARM` statement. JES3 programmers use the SYSTEM parameter on the `//*MAIN` statement.

TSO users that require a particular system must log on and submit started tasks to that system.

See *SPL: JES2 Initialization and Tuning* or *SPL: JES3 Initialization and Tuning* for more information.

Using Global Resource Serialization

The global resource serialization components of MVS/SP Version 2 and MVS/SP Version 1 Release 3 and later releases are compatible. Therefore, loosely-coupled MVS/XA and MVS/370 systems can use global resource serialization to control data sharing. The RESERVE/DEQ functions are also compatible.

In general, the fact that a global resource serialization complex includes mixed systems does not impose additional restrictions on the types of data sets that can be shared. Depending on the level of the systems in the complex, you might need to modify the RNLs to ensure that VSAM data sets are shared properly. See “Serializing VSAM Data Sets” in Chapter 5 for more information. Also read “Updating SYSTEMS Exclusion RNLs” in the same chapter. *Global Resource Serialization* describes data set sharing in general.

MVS/XA systems and MVS/370 systems that have MVS/370 DFP installed can share VSAM and CVOL catalogs. If Data Facility Extended Function (5740-XYQ) is also installed on the MVS/370 system, the systems can share ICF catalogs as well.

System Data Sets that Cannot be Shared

MVS/XA and MVS/370 systems cannot share the following system data sets: SYS1.LINKLIB, SYS1.LPALIB, SYS1.NUCLEUS, and SYS1.SVCLIB. MVS/370 systems cannot use the MVS/XA SYS1.PARMLIB data set as shipped. Installations also need two versions of SYS1.MACLIB. The MVS/XA system requires the MVS/XA expansions of downward incompatible macros. Also, some mapping macros are unique to MVS/370 or MVS/XA.

Using SYS1.PROCLIB in a Loosely-coupled JES3 Configuration

All converter-interpretor (CI) processing in a loosely-coupled JES3 configuration occurs on one processor. When converting jobs that execute procedures, the system performing CI processing uses the procedures in its own PROCLIB, not necessarily those in the PROCLIB of the system that will run the job.

In a loosely-coupled configuration that includes both MVS/370 and MVS/XA systems, you need to ensure that the processor performing the CI service uses the procedure appropriate for the system that will run the job. If the procedure for starting a task in MVS/370 is different from the equivalent MVS/XA procedure (as is the RMF procedure), you need to either:

- Modify the procedure to work on both MVS/370 and MVS/XA systems.
- Maintain two procedures and change the name of at least one of them.

“RMF Procedure” in Chapter 2 describes how to create an RMF procedure that starts either RMF Version 2 or Version 3.

DSI Procedures in a Loosely-coupled JES3 Configuration

If one of the processors involved in dynamic system interchange (DSI) is running MVS/XA and the other processor is running MVS/370, you must ensure that:

- The JES3 global function has sufficient devices that are supported by both the MVS/370 and MVS/XA systems. “Devices Not Supported” in Chapter 2 lists devices supported in MVS/370 but not in MVS/XA.
 - Any user-written JES3 routine that must run on a particular system (either MVS/370 or MVS/XA) is disabled before using DSI.
-

Appendix A. Parameter Changes in Incompatible Macros

This appendix describes changes to the parameters that the following macros pass to their service routines:

- ATTACH
- ESTAE
- EVENTS
- SMFEXIT
- STAX
- STIMER
- SYNCH
- WTOR

You need to be concerned about the changes only if you have programs that invoke the associated service routines directly (for example, by branch entry) instead of using the macros. You need to modify the parameter lists that those programs build.

To maintain compatibility, the MVS/XA service routines for all of macros listed except SYNCH accept MVS/370 or MVS/XA parameters. In most cases, the service routines check a flag bit (identified as FLAG BIT in the following figures) to determine which format (MVS/370 or MVS/XA) the input parameters are in. If the bit is 0, the parameters are in MVS/370 format. If the bit is 1, the parameters are either in MVS/XA format or in the format indicated by a format number somewhere in the parameter list. (The only defined format number is 1, which indicates MVS/XA format.)

If you build your own parameters, ensure that the flag bit and, in some cases, the format number correctly specify which version of the parameters you are passing to the service routine.

Note: In the following figures, blank fields represent fields that are not changed.

ATTACH Parameter List Changes

The FLAG BIT is the high-order bit of byte 8. The FORMAT NUMBER is byte 61.

MVS/370

Flags	Entry address	
	DCB address or zeros	
Flags	ECB address	
Flags	Give subpool number or list address	
Flags	Share subpool number or list address	
Flags	End-of-task exit routine address	
Resv.	JSCB address	
Task ID	STAI or ESTAI parameter list address	
Flags	STAI or ESTAI routine address	
Resv.	TASKLIB DCB address	
Flags	Resv.	Length of parameter list
Subpool number(s)		

MVS/XA

Entry address		
DCB address or zeros		
ECB address		
← FLAG BIT		
Give subpool number or list address		
Share subpool number or list address		
End-of-task exit routine address		
JSCB address		
STAI or ESTAI parameter list address		
STAI or ESTAI routine address		
TASKLIB DCB address		
Flags	Task ID	Length of parameter list
Subpool number(s)		
Flags	FORMAT NUMBER	Reserved

ESTAE Parameter List Changes

The FLAG BIT is the low-order bit of byte 13.

MVS/370

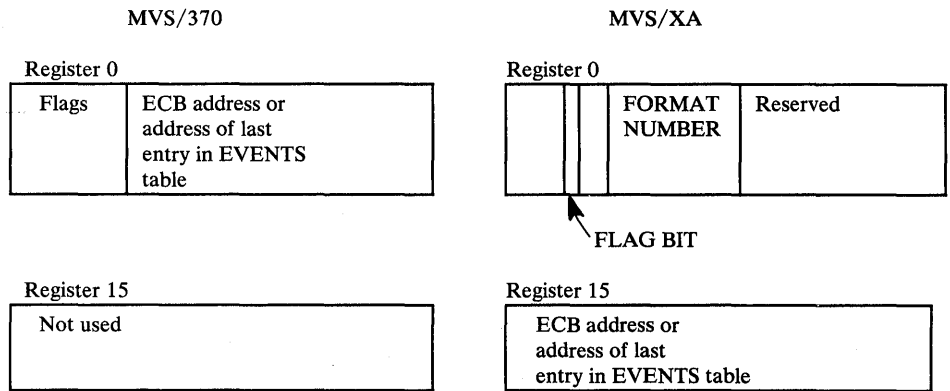
	ESTAE routine address
	Reserved

MVS/XA

	Reserved
	Reserved
	← FLAG BIT
	ESTAE routine address

EVENTS Parameter Changes

The FLAG BIT is the fourth bit of byte 0 in Register 0. The FORMAT NUMBER is the second byte of Register 0.




SMFEXIT Parameter List Changes

If the user specifies a work register on the SMFEXIT macro, the macro sets bit six of the parameter list to 1. In MVS/370, the bit is reserved. SMF uses the work register to save and restore the caller's addressing mode.

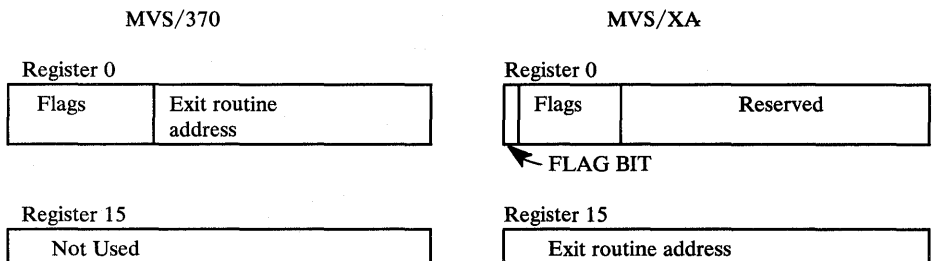
STAX Parameter List Changes

The FLAG BIT is the bit 6 of byte 16. The FORMAT NUMBER is byte 17.




STIMER Parameter Changes

The FLAG BIT is the high-order bit of byte 0.



SYNCH Parameter List Changes




AMODE flags:

- 00 AMODE=24
- 01 AMODE=DEFINED
- 10 AMODE=31
- 11 AMODE=CALLER

For more information, see “Downward Incompatible SYNCH Macros” in Chapter 9.

WTOR Parameter List Changes

The FLAG BIT is the high-order bit of byte 0.



Appendix B. Control Block Changes

Figure B-1 lists the control blocks that are new, changed, or deleted or that can reside anywhere in virtual storage (above or below 16 Mb). If a control block can reside anywhere, the figure indicates it might be above 16 Mb. If a control block's virtual storage location depends on the caller, the notes column indicates that the location is specified by the user. The figure does not include control blocks that are not changed and that must reside below 16 Mb.

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
ABDA (IHAABDA)	x				x		x		
			x		x		x		
ABDPL (IHAABDPL)	x				x			x	
			x		x			x	
ABEPL (IHAABEPL)	x			x			x		
AE (IHA AE)	x			x			x		
AMDDATA	x				x		x		
		x			x		x		
AMRQ (IHA CTM)		x			x			x	
AQAT (IHA AQAT)	x			x			x		
AQE (IHA AQE)	x					x			Replaced by IHA AE
ASCB (IHA ASCB)	x				x			x	
		x			x			x	
ASMHD (ILRASMHD)	x				x		x		
ASMVT (ILRASMVT)	x				x			x	
		x			x			x	
			x		x			x	
ASTE (IHA ASTE)	x				x			x	
ASVT (IHA ASVT)	x				x			x	
			x		x			x	
ASXB (IHA ASXB)	x				x			x	
ATECB (IEFZB432)		x			x			x	
ATTCH (IEZAT TCH)	x				x		x		User-specified
		x			x		x		
AWA (IEFVMAWA)		x			x			x	
BASEA (IEEBASEA)	x				x			x	Resides in the nucleus
		x			x			x	
BEB (IECDBEB)		x			x			x	
BLSRDTDT			x	x				x	
BLSRDUT			x	x				x	
BLSRESSY			x	x				x	
BLSRRDSY			x	x				x	
BLSRSSY			x	x				x	
CAT (IECDCAT)	x					x			

Figure B-1 (Part 1 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
CAW	x					x			
CCT (IRACCT)	x				x		x		
CCW (IOSDCCW)	x			x				x	New Format 1 CCW. Format 0 (the MVS/370 format) is unchanged.
CDA (IGFCDA)	x					x			
CDE (IHACDE)	x				x			x	
			x		x			x	
CHPT (IHAICHPT)	x			x			x		
CHRB (IOSDCHRB)			x		x		x		
CHT (IECDCHT)	x					x			
CIB (IEZCIB)	x				x			x	
		x			x			x	
CIFP (IEFCIFP)		x			x			x	
CIMP (IEFVMMWA)		x		x				x	
CLRATT (IEEVC102)	x			x				x	
CMCT (IRACMCT)	x			x			x		In the extended nucleus
COMMON	x				x			x	
		x			x			x	
			x		x			x	
COMWA (IEFCOMWA)		x			x			x	
CPT (IECDCPT)	x					x			
COB (IHACTM)		x			x			x	
CQE (IHACTM)		x			x			x	
CRCA (IECDCRCA)	x					x			
CSCB (IEECHAIN)		x			x			x	
			x		x			x	
CSD (IHACSD)	x				x			x	
CST (IECDCST)	x					x			
CSWK (IECDCSWK)	x					x			
CTM (IHACTM)	x				x				Only the XVSAV and CXSA maps have changed. The DOMPL can be above 16 Mb. The others must be below 16 Mb.
CTXT (IEZVX100)			x	x					Mapping macro
CVRWA (IEFCVRWA)		x			x			x	
CVT	x				x			x	
		x			x			x	
CXSA (IHACTM)	x				x			x	
		x			x			x	
DACB (IKJDACB)		x			x			x	
DALDDNAM (IEFZB4D2)		x			x		x		
DBOX (IOSDBOX)		x			x		x		
			x		x		x		

Figure B-1 (Part 2 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
DCQ (IHADCQ)	x			x				x	
DDP (IGFDDPRM)			x		x			x	
DDR (IHADDR)	x				x			x	
DDT (IECDDT)	x				x			x	
			x		x			x	
DEB (IEZDEB)	x				x			x	
DEVTAB (IFDEVTAB)			x		x			x	
DFE (IHADFE)	x			x			x		
DFLM (ADYDFLM)		x		x				x	
DMDT (IRADMDT)	x						x		
DOMC (IHADOMC)	x				x		x		
DOMPL (IHACTM)	x						x		
		x			x		x		
DQE (IHADQE)	x				x		x		
DSAB (IHADSAB)	x				x			x	
			x		x			x	
DSCA (ADYDSCA)		x		x				x	
DSPD (ADYDSPD)		x		x			x		
DSTAT (ADYDSTAT)		x		x				x	
DSVCB (IHADSVCB)	x			x			x		
		x			x			x	
DSX (ADYDSX)		x		x			x		
ECB (IHAECB)	x				x		x		User-specified
ECBE	x				x		x		User-specified
ENFCT (IEFENFCT)	x				x			x	
ENV (IGFENV)	x				x		x		
EPCB (IECDEPCB)		x			x			x	
EPIE (IHAETD)	x			x				x	
ERPIB (IGFERPIB)	x					x			
ESPI (IHAESPI)	x			x				x	
ESTA (IHAESTA)	x				x		x		User-specified
ESW (IHAESW)	x				x		x		
			x		x		x		
ETD (IHAETD)	x				x		x		User-specified
ETE (IHAETE)	x				x			x	
EVNT (IHAEVNT)	x				x			x	
EWA (EWAMAP)	x				x			x	
			x		x			x	
EXTD (ADYEXTD)		x		x			x		
FBQE (IHAFBQE)	x				x		x		

Figure B-1 (Part 3 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
FIX (IECDFIX)		x			x			x	
FQE (IHAFQE)	x				x		x		
FRRS (IHAFRRS)	x				x			x	
FTPT (IHACTM)		x			x			x	
GDA (IHAGDA)	x				x		x		
GENX (IEEZB816)			x	x			x		
GETPTWT (IEFZB600)	x					x			Replaced by the VSM parameter list
GTF records:									
CCW	x				x				
EOS	x				x				
IO	x				x				
I/O instructions	x			x					
PCI	x				x				
SIO	x					x			Replaced by new I/O instructions
UIO	x					x			
GSDA (IHAGSDA)			x		x			x	
GVT (ISGGVT)	x				x			x	
		x			x			x	
			x		x			x	
GWT (IHAGWT)	x				x		x		
HCL (IHAHCLOG)			x	x					Mapping macro
ICHPT (IHAICHPT)	x			x			x		
ICSE (IRAICSE)	x						x		
ICSM (IRAICSM)	x						x		
ICSS (IRAICSS)	x						x		
ICT (IRAICT)	x				x		x		
IDAL (IECDIDAL)		x			x			x	
IEAPMNP	x				x			x	
IEAPPNP	x				x			x	
IEAPQSR	x					x			
IEAPXNP	x					x			
IEAVNPB	x					x			
IEAVPSA	x				x			x	
IECDCPT	x					x			
IEEMBBQE		x			x			x	
IEEMBWKA		x			x			x	
IEESMFID		x			x			x	
IEEVC101	x				x			x	
IEFVMSWA		x			x			x	
IEFZB4D2	x				x			x	

Figure B-1 (Part 4 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
IFASMFR		x			x		x		
IFASMFRA		x			x			x	
IFASMFR2	x				x			x	
IFASMFR3	x				x			x	
IHAGSDA		x			x		x		
IHARCT (IHARCT)			x		x			x	
IHARRRA		x			x			x	
IHASNAP	x				x		x		User-specified
IHSA (IHAIHSA)	x				x			x	
INTWA (IEFVMIWA)		x			x			x	
			x		x			x	
IOCM (IECDIOCM)	x				x			x	
IOCOM (IECDIOCM)		x			x			x	
IOCX (IECDIOCX)	x				x			x	
IOE (ILRIOE)	x					x			
IOQ (IECDIOQ)	x				x			x	
		x			x			x	
IORB (ILRIORB)	x						x		
IOSB	x				x		x		The IOS driver determines the location of the IOSB.
		x			x		x		
			x		x		x		
IQE (IHAIQE)	x				x			x	
			x		x			x	
IRB (IHARB)	x				x			x	See RB in this table.
ITK (IEFVKEYS)		x			x			x	
IVT (IHAIVT)	x			x			x		
JCTX (IEFJCTX)		x			x			x	
			x		x			x	
JESCT (IEFJESCT)	x				x			x	
		x			x			x	
JFCB (IEFJFCBN)		x			x			x	
			x		x			x	
JMR (IEFJMR)		x			x			x	
JSBVT (IEFJSBVT)		x			x			x	
JSCB (IEZJSCB)	x				x			x	
LCCA (IHALCCA)	x				x			x	Fetch-protected
		x			x			x	
LCH (IECDLCH)	x					x			
LCT (IEFALLCT)	x				x			x	
		x			x			x	

Figure B-1 (Part 5 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
LDA (IHALDA)	x				x		x		
			x		x		x		
LLCB (IHALLCB)		x		x			x		
LLPM (IHALLPM)		x		x				x	
LMSG (IECDLMSG)	x				x			x	
LPBT (IRALPBT)	x			x			x		
LPDE (IHALPDE)	x				x			x	
LRB (IHALRB)	x						x		
MCHTRACE	x			x			x		
MCT (IRAMCT)	x				x		x		
MIHW (IGFDMIHW)	x					x			
MPL (IHAMPL)	x				x		x		
MSG (IGFMSG)	x						x		
MSRASDCA (IEEZB808)	x			x				x	
			x		x			x	
NEL (IEFNEL)		x			x			x	
NVT (IHANVT)	x				x			x	
ORB (IHAORB)	x			x			x		
OUCB (IRAUCB)	x				x		x		
OUXB (IHAUXB)	x				x		x		
PCCA (IHAPCCA)	x				x			x	
PEL (ISGPEL)		x			x		x		
			x		x		x		
PEXB (ISGPEXB)		x			x			x	
PIPL (IEZPIPL)	x				x			x	User-specified
PQA	x					x			
PQE (IHAPQE)	x					x			Replaced by IHARD
PRULE (IEZPRULE)	x				x			x	User-specified
PSA (IHAPSA)	x				x			x	The upper 2 K is fetch-protected.
		x			x			x	
			x		x			x	
PTUD (IOSDPTUD)	x						x		
PVT (IHAPVT)	x				x			x	The old PVT is split among IHAPVT, IARRIT, and IARRCE.
		x			x			x	
PWA (IGFPWA)	x				x		x		
P25C (EWAP25C)		x			x		x		
QCB (ISGQCB)			x		x			x	
QEL (ISGQEL)			x		x			x	
QHT (ISGQHT)		x			x			x	

Figure B-1 (Part 6 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
QRO (IHACTM)		x			x			x	
QWA (ISGQWA)			x		x			x	
QWB (ISGQWB)			x		x			x	
QXB (ISGQXB)			x		x			x	
RB (IHARB)	x				x			x	
RCE (IARRCE)	x			x			x		
		x			x		x		
RCT (IRARCT)	x				x		x		
RCTD (IEARCTD)	x				x			x	
RD (IHARD)	x			x			x		
RDCM (IEERDCM)	x				x			x	
			x		x			x	
RESV (IOSDRESV)		x			x			x	
RGR (IHARGR)	x			x			x		
RMCA (IRARMCA)	x				x		x		
RMCT (IRARMCT)	x				x		x		
		x			x		x		
RMEP (IRARMEP)	x				x		x		
RMEX (IRARMEX)	x				x		x		
RMPL (IHARMPL)	x				x			x	
RMPT (IRARMPT)	x				x		x		
		x			x		x		
RMQH (IRARMQH)	x						x		
RMSB (IRARMSB)	x				x		x		
RNLE (ISGRNLE)			x		x			x	
RPT (ISGRPT)		x			x			x	
RQE (IECDRQE)	x				x			x	
		x			x			x	
RQSV (IRARQSRV)	x						x		
RRQ (IECDRRQ)		x			x			x	
RSMHD (IHARSMHD)	x					x			Replaced by IARRAB
RTCT (IHARTCT)	x				x			x	
		x			x			x	
			x		x			x	
RTSD (IHARTSD)	x				x			x	
RT1W (IHART1W)	x				x			x	
RVT (IHARVT)	x							x	In the non-page protected nucleus
SCA (IHASCA)	x				x			x	
SCB (IHASCB)	x				x			x	
SCD (IECDSCD)	x					x			
SCHIB (IHASCHIB)	x			x			x		

Figure B-1 (Part 7 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
SCL (IEEZB815)			x	x					Mapping macro
SCSR (IEZVG100)		x		x				x	
SCT (IEFASCTB)			x		x			x	
SCVT (IHASCVT)	x				x			x	
			x		x			x	
SCWA (IHASCWA)	x			x			x		
		x			x		x		
			x		x		x		
SDEPL (IHASDEPL)	x			x				x	
		x			x			x	
SDRSB (IHASDRSB)	x				x			x	
SDUMP (IHASDUMP)	x				x		x		User-specified
SDWA (IHASDWA)	x				x			x	
		x			x			x	
SDWORK (IHASDWRK)		x			x			x	
SHDR (IHASHDR)	x				x			x	
			x		x			x	
SIAB (IECDSIAB)	x					x			
SIOT (IEFASIOT)		x			x			x	
SJDFP (IEFSJDFP)		x		x				x	
SJDLP (IEFSJDLP)		x		x				x	
SJEXP (IEFSJEXP)		x		x				x	
SJFNP (IEFSJFNP)		x		x				x	
SJGEP (IEFSJGEP)		x		x				x	
SJINP (IEFSJINP)		x		x				x	
SJJD P (IEFSJJD P)		x		x				x	
SJPRFX (IEFSJPRFX)		x		x				x	
SJPUP (IEFSJPUP)		x		x				x	
SJREP (IEFSJREP)		x		x				x	
SJRUP (IEFSJRUP)		x		x				x	
SJWRP (IEFSJWRP)		x		x				x	
SLCA (IFASLCA)		x		x				x	
SLFP (IHASLFP)	x				x		x		
		x			x		x		
			x		x		x		
SLPL (IHASLPL)	x				x		x		
		x			x		x		
SLWA (IHASLWA)	x				x		x		
		x			x		x		
SMCA (IEESMCA)	x				x			x	
		x			x			x	
SMDLR (IHASMDLR)	x				x			x	

Figure B-1 (Part 8 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
SMWK (IHASMWK)	x				x			x	
			x		x		x		
SPCT (IHASPCT)	x					x			Replaced by IARSFTE
SPQA (IHASPQA)	x			x			x		
SPQE (IHASPQE)	x				x		x		
SPT (IHASPT)	x			x			x		
SQAT (IHASQAT)	x			x			x		
SRB (IHASRB)	x				x		x		User-specified
SRCD (ADYSRCD)		x		x				x	
SRIO (IOSDRIO)			x		x			x	
SRPL (IEEZB814)	x			x				x	
SSAL (IEFSSAL)		x			x			x	
SSCM (IEFSSCM)		x			x			x	
SSCVT (IEFJSCVT)		x			x			x	
SSDA (IEFSSDA)		x			x			x	
SSJS (IEFSSJS)		x			x			x	
SSL (IHASSL)	x			x			x		User subpool and key
SSOB (IEFJSSOB)	x				x		x		If the subsystem(s) given control runs in 31-bit addressing mode, the SSOB can be above 16 Mb.
SSOBH (IEFSSOBH)	x				x			x	
SSRB (IHASSRB)	x				x		x		
SSVS (IEFSSVS)	x				x			x	
SSWT (IEFSSWT)		x			x			x	
STKE (IHASTKE)	x				x			x	
SUB (IEECSUB)	x				x			x	Resides in the nucleus
SVCS (ADYSVCS)		x		x				x	
SVCTABLE (IHASVC)	x				x			x	
SVT (IHASVT)	x				x			x	
SWB (IEFSWB)		x		x				x	
SYMPQ (ADYSYMP)		x		x				x	
TAXE (IKJTAXE)	x				x			x	
TBUF (IHATBVT)	x			x			x		
TBVT (IHATBVT)	x			x			x		
			x		x		x		
TCB (IKJTCB)	x				x			x	
TCCW (IECDTCCW)	x				x			x	
		x			x			x	
TCT (IEFTCT)	x				x			x	
		x			x			x	
			x		x			x	

Figure B-1 (Part 9 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
TDCM (IEETDCM)	x				x			x	
		x			x			x	
			x		x			x	
TFWA (IHATFWA)			x		x		x		
TOB (IHATOB)	x			x			x		
TOT (IHATOT)	x			x				x	
TPC (IEAVVTPC)			x		x			x	
TQE (IHATQE)	x				x		x		User-specified
			x		x		x		
Trace table entries:									Mapped by IHATTE
ACR	x			x			x		
ALTR	x			x			x		
BRANCH TRACE	x			x			x		
CALL	x				x		x		
CLKC	x				x		x		
DSP	x				x		x		
EMS	x				x		x		
EXT	x				x		x		
IO	x				x		x		
I/O instructions	x			x			x		
MCH	x			x			x		
PC trace	x				x		x		
PGM	x				x		x		
PT trace	x				x		x		
RST	x			x			x		
SIO	x					x			
SRB	x				x		x		
SS	x				x		x		
SSAR trace	x				x		x		
SSRB	x				x		x		
SUSP	x			x			x		
SVC	x				x		x		
SVCR	x				x		x		
SVCE	x				x		x		
USRn	x			x			x		
WAIT	x				x		x		
TRBP (IHATRBPL)	x						x		User-specified
TREP (IHATREPL)	x						x		User-specified
TRNSQ (ADYTRNM)		x		x				x	
TROB (IHATROB)	x			x			x		
TRQE (IRATROEL)	x						x		User-specified

Figure B-1 (Part 10 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
TRVT (IHATRVT)	x			x				x	
			x		x			x	
TTCH (IHATTCH)	x			x			x		
TTE (IHATTE)	x			x			x		See trace table entry in this table.
			x		x		x		
TXTFT (IEFTXTFT)		x			x			x	
UCB (IEFUCBOB)	x				x			x	
		x			x			x	
			x		x			x	
UCB look-up table	x					x			
UCDX (IEEUCDX)			x		x			x	
UCM (IEECUCM)	x				x			x	Resides in the nucleus
		x			x			x	
			x		x			x	
USECB (IOSUSECB)	x				x				
UXIR (IHACTM)			x	x				x	
VCB (IHAVCB)	x				x		x		User subpool and key
VFPM (IHAVFPM)	x						x		User subpool and key
VFQE	x				x			x	
VFVT (IHAVFVT)			x		x			x	
VOID (IECDVOID)	x				x			x	
		x			x			x	
			x		x			x	
VRAMAP (IHAVRA)		x			x			x	
VSL (IHAVSL)	x						x		User subpool and key
VSRI (IEEZB812)			x	x				x	
VTPC (IEAVVTPC)	x				x			x	
VTSP (IEFVTSP)		x			x			x	
WMST (IRAWMST)	x				x		x		
WPGDT (IRAWPGDT)	x						x		
			x		x		x		
WPL (IEZWPL)	x				x		x		User-specified
			x		x		x		
WQE (IHAWQE)			x		x			x	
		x			x			x	
WSAC			x		x		x		
WSAVTC (IHAWSAVT)	x				x		x		Component-specified, fetch protected
		x			x		x		
			x		x		x		

Figure B-1 (Part 11 of 12). Control Block Changes

Control Block	Release 1.0	Release 1.1	Release 1.2	New	Changed	Deleted	Might Be Above 16 Mb	Must Be Below 16 Mb	Notes
WSAVTG (IHAWSAVT)	x				x		x		Component-specified, fetch protected
WSAVTL (IHAWSAVT)			x				x		Unchanged
WWB (IHACTM)		x			x			x	
			x		x			x	
XCPS (IECDXCPS)		x			x			x	
			x		x			x	
XDBA (IECDXDBA)	x				x			x	
		x			x			x	
XFRR (IECDXFRR)		x			x			x	
XSA (IEEXSA)	x				x			x	
			x		x			x	
XSB (IHAXSB)	x				x		x		User-specified
XTLST (IHAXTLST)	x				x			x	
XV (IHACTM)		x			x			x	
XVSAV (IHACTM)	x				x			x	
YSTAK (IHAYSTAK)	x				x		x		
		x			x		x		
			x		x		x		

Figure B-1 (Part 12 of 12). Control Block Changes

Index

\$

\$HASPEQU macro 3-12, 5-5
\$HASPGEN macro 3-12, 5-5

A

ABEND

B37 6-8
D37 6-8
E37 6-8
macro 3-5, 3-39, 6-23
out-of-space 6-8
16E 3-17
504 3-10
538 3-13
80A 6-8

access methods 3-25
accessing control blocks above 16 Mb 3-35
accounting procedures 7-1
ACF/VTAM 2-7
ACONs 3-17
action message retention facility 5-7
address
 constants 3-17
 space tracing 6-20
 space vector table (ASVT) 2-15, 2-22
 spaces
 DUMPSRV 2-24
 full-function 2-24
 master scheduler 2-9
 page-in rate 2-20, 8-9
 SMF 2-23, 2-24
 working set size 2-20, 8-9

addresses

device 2-4
in the PSW 3-17
processor 3-17
real 3-20
size of 3-25
UCB 3-15

addressing mode

changing 3-27
definition of 3-25
establishing 3-27
parameter on CIRB 3-39
parameter on SYNCH 3-43
31-bit considerations 3-2

ADYDFLT load module 6-10

ADYHDFMT DAE header exit 5-6, 6-11
ADYSETxx PARMLIB member 2-18, 6-9
ADYSET00 PARMLIB member 6-9
ADYSET01 PARMLIB member 6-9
ADYSET02 PARMLIB member 6-9

algorithms

data set selection 8-9
slot selection 8-10

ALLNUC dump option 6-3

ALLOCATE modules 3-15

allocating SYS1.DUMPxx data sets 1-1

allocation space defaults CSECT 3-14

ALLVNUC dump option 6-3

ALT system parameter 2-21

alternate nucleus, specifying 4-2

AMASPZAP service aid program

changing global resource serialization thresholds 3-13
changing hot I/O recovery actions 4-4, 5-4
changing hot I/O thresholds 5-4
overriding segment protection 3-19

AMBLIST utility 3-29, 6-11

AMDPRDMP service aid

 See PRDMP

AMDPRECT module 5-1

AMDSADMP macro 2-27

AMDSARDM module 2-27

AMODE

 compared to current addressing mode 3-29
 description of 3-27
 determining 3-29
 flags in the CESD 3-28, 9-2
 flags in the ESD 3-28
 flags in the PDS directory entry 3-28, 9-2
 parameter on CIRB 3-39
 parameter on SYNCH 3-3, 3-43, 9-7
 specifying 3-28

APARs

 for DFDS 1.4 8-1
 for IEBCOPY ALTERMOD and COPYMOD 8-1, 8-2
 for program fetch 7-3, 8-2
 for SVCDUMP 3-15

APF authorization 2-22, 2-23, 3-17, 8-7

ASM (auxiliary storage manager)

 backing slot function 8-9
 control blocks in PRDMP output 6-11
 data set selection algorithm 8-9
 lock 3-18
 slot selection algorithm 8-10

ASMDATA PRDMP statement 6-11

ASMVT

 ASMSLOT field 8-9
 ASMSLOTV field 8-9

Assembler H Version 1 2-26

Assembler H Version 2

 indicating AMODE/RMODE 3-28
 sysgen requirement 2-2
 when to use 9-2

assembler language global SET symbol 9-1

assembling programs

 after installing BTAM/SP 3-9
 after installing compatibility PTFs 3-3, 3-8
 containing SYNCH macros 9-7
 MVS/XA system programs 9-6
 to run in MVS/XA 9-2

ASVT (address space vector table) 2-15, 2-22

ATTACH macro

 changing the addressing mode 3-27
 differences 3-5
 incompatible MVS/XA expansion 9-4
 parameter list changes A-1

authorized programs, changes affecting 3-1

B

backing slot function in ASM 8-9

backup considerations

 copying MVS/370 DLIBs 2-2
 producing backup tapes 9-8
 program incompatibilities 9-7
 program product licenses 9-7
 size of private area used 9-7

- switching between 370 and 370-XA 9-7
- using a common IOCDs 9-7
- workload 9-8
- BAL instruction 3-29, 9-3
- BALR instruction 3-29, 9-3
- BAS instruction 3-30
- BASR instruction 3-30
- BASSM instruction
 - description of 3-32
 - example of using 3-33, 3-34
- BLDL
 - lists 2-20, 8-7
 - system parameter 2-20, 2-22
 - using LLA directory 8-7
- BLDLF system parameter 2-20, 2-22
- BLSABDPL mapping macro 3-5, 3-39, 5-6
- BLSAMPLE member of SYS1.SAMPLIB 6-19
- BLSPDISE IPCS panel 6-16
- BLSPDSLE IPCS panel 6-16
- BLSQMDEF macro 3-5, 3-39
- BLSQMFLD macro 3-5, 3-39
- BLSRESSY macro 3-5, 3-39
- BQEs for SMF 3-17
- branch
 - instructions 3-29, 3-30, 3-31, 3-32
 - tracing 6-20
- BROWSE IPCS panel 6-17
- BSM instruction
 - description of 3-31
 - example of using 3-33, 3-34, 3-35
- BTAM RESETPL macro 3-5, 3-9
- BTAM/SP program product 3-9
- buffer queue elements for SMF 3-17
- buffers
 - I/O 3-37
 - obtained by access methods 3-21
 - real addresses of 3-20
 - SMF 2-23
 - used with EXCP 3-37
 - used with EXCPVR 3-20
 - used with VSAM services 3-36
- bypassing VSM's storage availability check 5-3
- B37 ABEND code 6-8

C

- CALL macro 3-5
- CALLDISP macro 3-7
- CALLRTM macro 3-7, 3-39, 6-23
- CANCEL command 4-7
- catalogs
 - CVOL 9-9
 - ICF 9-9
 - sharing 9-9
 - VSAM 9-9
- CB dump option 6-3
- CCWs
 - used with EXCP 3-37
 - used with EXCPVR 3-20
- CDEs (contents directory entries) 2-21
- cell pool 3-39
- CESD (composite external symbol dictionary) 3-28, 9-2
- channel
 - command words
 - See CCWs
 - numbers 2-4
 - sets 2-4
- CHANNEL macro 2-6
- checking for available storage 5-3

- checkpoint/restart diagnostic data 6-24
- CHKPT macro 3-5, 3-7, 3-39, 9-4
- CHNGDUMP command 4-7
- CHPID macro 2-4
- CI (converter-interpreter) processing 9-9
- CIRB macro 3-7, 3-39
- CLCL (Compare Logical Long) instruction 3-30
- CMB system parameter 2-15, 2-21
- CNTLUNIT macro 2-4
- coexistence
 - considerations 9-1
 - definition 9-1
 - DSI procedures 9-9
 - handling downward incompatible macros 9-4
 - programming considerations 9-2
 - routing jobs 9-8
 - sharing data sets 9-8
- color consoles
 - controlling message traffic 4-5
 - display options 4-5
 - specifying color attributes 2-23, 4-5
- commands
 - CANCEL 4-7
 - CHNGDUMP 4-7
 - CONFIG 4-7, 4-11
 - CONTROL
 - M 4-7
 - S 4-7
 - V 4-8, 5-8
 - DISPLAY
 - CONSOLES 4-9
 - DUMP 4-9
 - GRS 4-10
 - M 4-10
 - MPF 4-10
 - DUMP 4-10
 - DUMPDS 4-11
 - FORCE 4-11
 - incompatible 4-6
 - MODE 4-11
 - MODIFY 4-11, 8-8
 - MONITOR 4-11, 4-13
 - MSGRT 4-11
 - SET
 - DAE 4-12, 6-9
 - MPF 4-5, 5-7
 - SMF 4-12
 - SLIP 4-12, 4-13
 - START 4-13, 8-8
 - STOP 4-13, 8-8
 - STOPMN 4-11, 4-13
 - summary of changes to 4-6
 - TRACE 4-14
 - TRACK 4-5
 - VARY
 - CH 4-14
 - CPU 4-14
 - PATH 4-14
 - STOR 4-14
- COMMNDxx PARMLIB member 2-16
- Compare Logical Long (CLCL) instruction 3-30
- composite external symbol dictionary (CESD) 3-28, 9-2
- concatenating data sets
 - to SYS1.LINKLIB 2-22, 2-23, 8-7
 - to SYS1.LPALIB 2-9
- CONFIG
 - command 4-7, 4-11
 - frame 4-1

- Configuration (CONFIG) frame 4-1
- CONFIGxx PARMLIB member
 - summary of changes 2-18
 - using with the CONFIG command 4-7
- console
 - changing specifications for 4-7
 - clusters 4-5
 - color 4-5
 - frames
 - OPRCTL (operator control) 4-2
 - SYSCTL (SCP manual CNTL) 4-1
 - reestablishing console specifications 4-8
 - requesting status information 4-9
 - 3279 4-4
- CONSOLE macro 2-5
- contents directory entries (CDEs) 2-21
- contents supervision (CSV) modules 3-15
- control blocks
 - See also specific control block name
 - formatter service 5-9
 - in PRDMP output 6-10
 - list of differences B-1
 - retrieving data above 16 Mb 3-35
- CONTROL command
 - M 4-7
 - S 4-7
 - V 4-8, 5-8
- control records 8-2
- converter-interpreter (CI) processing 9-9
- copying
 - DLIBs 2-2
 - IOCP CSECTs 2-3
 - IOCP deck 2-3
 - IPCS modules 6-19
 - modules for fetch performance 8-3
 - PRDMP modules 6-19
- count values in load modules
 - how used 8-2
 - inserting 8-2, 8-3, 8-4
 - maintaining 8-4
- CPENABLE parameter in IEAOPTxx 2-21
- CPOOL macro 3-5, 3-39
- CPU
 - addresses 3-17
 - lock 3-18
 - timer 3-14
- CPUTIMER macro 3-5, 3-14, 3-39
- cross memory entry table entries 3-23
- CSA
 - dump option 6-3
 - specifying the size of 2-13, 2-14, 2-21
 - system parameter 2-13, 2-14, 2-21
- CSV (contents supervision) modules 3-15
- CSVLLCRE module 2-7
- CSV300I message 8-1
- CTRLPROG macro
 - specifying CSA and SQA size 2-13
 - unsupported parameters and options 2-6
- CVOL catalogs 9-9
- CVT
 - CVTDCB field 9-3, 9-6
 - CVTEFLPE field 3-16
 - CVTEFLPS field 3-16
 - CVTFLPAE field 3-16
 - CVTFLPAS field 3-16
 - CVTMVSE bit 9-3, 9-6
 - CVTNUCB field 3-16
- DAE (dump analysis and elimination)
 - command 2-20
 - components using 6-9
 - controlling 6-9
 - description 6-8
 - header exit, ADYHDFMT 6-11
 - symptom data 5-6, 6-8, 6-11
 - SYS1.DAE data set 2-8
- DAEALLOC member of SYS1.SAMPLIB 2-9, 6-10
- DAEDATA PRDMP statement 5-6, 6-11
- DAM 3-25
- DASD
 - initializing 2-10
 - operand on DUMP system parameter 2-15
 - use of space 8-3
- DASDR service routine 3-18
- DAT-off
 - modifying programs that run 3-22
 - module 3-23
 - nucleus 3-23
- data extent block 3-17
- Data Facility Data Set Services (DFDSS) 2-2, 9-8
- Data Facility Extended Function (DFEF) 9-9
- data management services 3-25
- data sets
 - See also system data sets
 - selection algorithms 8-9
 - SMF 8-9
- DATASET macro 2-5, 2-6, 2-9
- DATOFF macro 3-7
 - example 3-23
 - function of 3-39
- DC option on LKED EXEC statements 8-6
- DCBS option on LKED EXEC statements 8-6
- DD statements
 - for SYS1.DUMPxx data sets 1-1
 - in the PRDMP procedure 2-24
- DEB
 - for fetch-protected areas 3-17
 - for the LNKLST concatenation 3-17
- DEBAPFIN bit in the LNKLST DEB 3-17
- DEBCHK service routine 3-17
- debugging considerations 6-20
- deleting messages 5-7
- DEMF (Display Exception Monitor Facility) 2-3
- descriptor codes for messages 5-7
- device
 - address on IODEVICE 2-4
 - addresses for stand-alone dump 2-27
 - allocation load module 3-14
 - allocation tables
 - changing programs that access 3-17, 3-45
 - removing references to 2-16
 - allocation, non-specific 2-21
 - connect time
 - definition of 7-1
 - in calculating I/O service 2-20
 - in SMF records 7-1, 7-5
 - number 2-4
- Device Support Facilities Release 6 2-10
- devices
 - DASD, initializing 2-10
 - for page data sets 2-8
 - for system data sets 2-8
 - maximum allowed 2-3
 - not supported 2-6
 - specifying MIH intervals for 2-17

- DEVMSKMT table 2-16, 3-45
- DEVNAMET table 2-16, 3-45
- DFDSS (Data Facility Data Set Services) 2-2, 9-8
- Dfef (Data Facility Extended Function) 9-9
- direct access storage devices
 - See DASD
- DISP
 - JCL parameter 4-3
 - lock 3-18
- DISPLAY
 - command
 - CONSOLES 4-9
 - DQMP 4-9
 - GRS 4-10
 - M 4-10
 - MPF 4-10
 - PRDMP statement 6-11
- Display Exception Monitor Facility (DEMF) 2-3
- disposition of SYS1.DUMPxx data sets 1-1
- Divide Extended (DXR) instruction 3-31
- DLIBs 2-2
- downward incompatible macros 9-4
- DRWDASDR 9-8
- DSI (dynamic system interchange) procedures 9-9
- DSP (dynamic support program) 5-5
- dual paths in programs
 - example 9-6
 - selecting path 9-3
 - when required 9-3
- dummy SMF records 7-5
- dump
 - command 4-10
 - data sets
 - accessing via IPCS 6-16
 - allocating 1-1, 4-3
 - associating with a specific processor 2-15
 - clearing 4-11
 - defining 2-8, 4-11
 - deleting 4-11
 - eligible devices for 2-8
 - location of 6-16
 - moving 6-16
 - number and size of 2-8
 - scanning 6-11
 - exit routines 5-4, 5-8
 - format changes 6-6
 - headers
 - in formatted user dump 6-6
 - in SVC dumps 6-7
 - in SYSMDUMP 6-6, 6-7
 - indexes 6-6
 - options
 - ALLNUC 6-3
 - ALLVNUC 6-3
 - in IEAABD00 2-19
 - in IEADMP00 2-20
 - in IEADMR00 2-20
 - in SNAP parameter list 6-2
 - NOSYM 6-5
 - on SNAP macro 3-39, 3-42
 - on the DUMP command 6-2
 - on the SDUMP macro 3-41, 6-2
 - on the SNAP macro 3-42
 - SPLS 6-4
 - SQA 6-4
 - SUBPLST 6-4
 - SUBTASKS 6-4
 - SUM 6-4, 6-5
 - summary of new and changed 6-2
 - TRT 6-4

- stand-alone (See stand-alone dump)
- suppressing 6-7
- SVC (See SVC dumps)
- symptom 6-4
- SYSMDUMP (See SYSMDUMP)
- system parameter 2-15, 2-21
- SYSUDUMP (See SYSUDUMP)
- user summary 6-5
- dump analysis and elimination
 - See DAE (dump analysis and elimination)
- DUMP system parameter 2-15
- DUMPDS command 4-11
- DUMPSRV
 - address space 4-3
 - procedure in SYS1.PROCLIB 2-24
- DXR (Divide Extended) instruction 3-31
- dynamic address translation
 - See DAT-off
- dynamic allocation interface routine 3-14
- dynamic allocation user exit 3-14
- dynamic support program (DSP) 5-5
- dynamic system interchange (DSI) procedures 9-9
- D37 ABEND code 6-8

E

- ECBs, extended 3-17
- ECT (print dump exit control table) 5-1, 5-6
- ECT service 5-10
- Edit and Mark (EDMK) instruction 3-30
- EDMK (Edit and Mark) instruction 3-30
- EDTs (eligible device tables) 2-7
- eligible device tables (EDTs) 2-7
- ENQ macro
 - for VSAM data sets 5-1
 - limiting concurrent requests 3-12
 - summary of changes 3-5, 3-40
- entry points in IEFW21SD 3-38
- Environmental Recording, Editing, and Printing Program
 - See EREP
- EREP (Environmental Recording, Editing and Printing Program)
 - PRDMP exit 6-18
 - program produce co-requisite 6-18
- ESD (external symbol dictionary) 3-28
- ESPIE
 - macro 3-5, 3-40, 9-4
 - service routine 3-18
- ESTAE macro 3-5, 9-3, 9-4, A-2
- EVALDEF IPCS subcommand 6-14
- EVALDUMP IPCS subcommand 6-14
- EVENTS macro 3-5, 9-4, A-3
- EXCP
 - counts 2-20, 7-1, 7-3
 - macro
 - backing I/O buffers 3-37
 - parameter requirements 3-36
 - performing I/O above 16 Mb 3-37
 - performing I/O in 31-bit mode 3-36
 - using virtual IDAWs 3-37
- EXCPVR macro 3-7
 - parameter requirements 3-36
 - performing I/O above 16 Mb 3-20, 3-37
 - performing I/O in 31-bit mode 3-36
 - using the PGFX appendage 3-20
- EXEC statement
 - in the PRDMP procedure 2-24
 - LKED 8-6

execution time 2-20, 8-9
exit routines
 See user exit routines
exit services router 5-8
explicit tracing 6-20
extended
 color support 4-5
 CSA 2-13, 2-21
 ECBs 3-17
 region size, obtaining 5-3
 SQA 2-13, 2-21
external symbol dictionary (ESD) 3-28
External Writer 2-3
E37 ABEND code 6-8

F

FDPs (Field Developed Programs) 1-2
FESTAE macro 3-7, 9-4
fetch
 See also program fetch
 protection in PSA 3-20
Field Developed Programs (FDPs) 1-2
FIX system parameter 2-21
FLPA (fixed link pack area)
 building 2-9
 changes 3-16
 page protection 3-19
FORCE command 4-11
format model processor service 5-9
FORMAT PRDMP statement 6-11
frames
 See console frames and configuration frames
full-function address spaces, starting 2-24

G

generalized trace facility
 See GTF
GENERATE macro 2-5, 2-6, 2-10
generating
 an MVS/XA system 2-1, 2-2
 stand-alone dump 2-26
GETMAIN macro 3-5
 differences in processing 3-10
 in summary of changes 3-40
 limit on requests
 determining 5-3
 exceeding 6-23
 LOC parameter 3-44
 overlapping parameters 3-10
 suggestions for using 9-3
 VRC parameter 3-43
 VRU parameter 3-43
global RESERVE requests 2-26
global resource serialization
 displaying information about 4-10
 in a mixed environment 9-8
 limiting concurrent requests 3-12
 RNLs (resource name lists) 2-16
 serializing VSAM data sets 5-1
global SET symbol 9-1
GQSCAN macro
 limiting concurrent requests 3-12
 summary of changes 3-5, 3-40
GRS
 See global resource serialization
GRSRNL system parameter 2-16, 2-22

GRSRNLxx PARMLIB members 2-16, 2-18, 2-22
GRSRNL00 PARMLIB member 2-17
GTF (generalized trace facility)
 modules 3-15
 records 3-16
 tracing USR events 3-45
GTFPARAM PARMLIB member 2-18
GTRACE macro 3-5, 3-40, 3-45
GVT
 GVTCREQ field 3-13
 GVTCREQA field 3-13

H

HAM (JES2 spool access method) 5-5
hardcopy log records 3-13
header exits for PRDMP 5-6
highlighting messages 2-23
hot I/O interrupts
 controlling processing 2-23, 4-3
 recovery actions 4-4, 5-4
 thresholds 4-3, 5-4
HOTIO statement in IECIOSxx 2-23, 5-4

I

I/O

 above 16 Mb 3-37
 configuration data set
 See IOCDs
 configuration requirements 2-5
 event recording 2-18
 hot 2-23, 4-3
 in 31-bit mode 3-36
 instructions 3-12, 3-31
 interrupt processing 2-21, 4-3
 load balancing 2-21
 service, calculating 2-20
 to FLPA 3-16
 to real storage above 16 Mb 3-20
 using access methods 3-25, 3-36
 using EXCP 3-36, 3-37
 using EXCPVR 3-36, 3-37
IARUTVR module 3-15
IAR004I message 2-12
IATYMOD macro 3-12, 5-5
ICCLPB parameter in IEAOPTxx 2-21
ICF catalogs 9-9
ICP404I message 2-4
IDAWs (indirect addressing words)
 used with EXCPVR 3-20
 virtual 3-37
IEAABD00 PARMLIB member 2-19
IEABLDxx PARMLIB member 2-20
IEACMD00 PARMLIB member
 fixed storage allocations when executed 2-12
 SET DAE command 6-9
 START LLA command 2-24, 8-8
 summary of changes 2-20
IEADMP00 PARMLIB member 2-20
IEADMR00 PARMLIB member 2-20
IEAFIXxx PARMLIB member 2-16, 2-20
IEAIPL04 module 2-13
IEAIPSxx PARMLIB member 2-20
IEALIMIT exit 5-2
IEALOD00 PARMLIB member 2-21

- IEALPAxx PARMLIB member 2-16, 2-21
- IEAOPTxx PARMLIB member 2-21
- IEAPAKxx PARMLIB members 2-21, 2-22
- IEAPAK00 PARMLIB member 2-21
- IEASMFEX module 3-17
- IEASYSxx PARMLIB member
 - parameters
 - ALT 2-21
 - BLDL 2-20, 2-22
 - BLDLF 2-20, 2-22
 - CMB 2-21
 - CSA 2-13, 2-14, 2-21
 - DUMP 2-15, 2-21
 - FIX 2-21
 - GRSRNL 2-16, 2-22
 - LNKAUTH 2-22, 8-8
 - LPA 2-9, 2-22
 - MAXUSER 2-15, 2-22
 - MLPA 2-21
 - MSTRJCL 2-22
 - PAK 2-22
 - RSU 2-12
 - RSVNONR 2-15, 2-22
 - RSVSTRT 2-15, 2-22
 - SQA 2-13, 2-21
 - summary of changes to 2-21
- IEAVEDAT, DAT-off nucleus 3-23
- IEAVEURn, DAT-off module 3-23
- IEAVMXIT user exit 4-7, 5-7
- IEAVNIP0 module 2-13
- IEAVTABX module 5-4
- IEAVTRV module 3-15
- IEAVTSEL module 5-4
- IEA240I message 2-22
- IEA838I message 6-8
- IEBCOPY
 - ALTERMOD parameter 8-3
 - changes to 8-2
 - COPYMOD parameter 8-3
- IECIOSxx PARMLIB member
 - HOTIO statement 5-4
 - specifying MIH intervals 2-17
 - summary of changes 2-23
- IECVHIDT module 5-4
- IEECVXIT user exit 4-5, 5-7
- IEEMB846 module 3-12
- IEESYSAS procedure 2-24
- IEE978E message 2-23
- IEFAB4DC entry point in IEFW21SD 3-38
- IEFAB4UV entry point in IEFW21SD 3-38
- IEFAB4UV load module 3-45
- IEFAB445 entry point in IEFW21SD 3-38
- IEFAB445 load module 3-14
- IEFDB401 load module 3-14
- IEFDEVPT table 2-16, 3-45
- IEFEB4UV load module 3-45
- IEFGB4DS entry point in IEFW21SD 3-38
- IEFGB4UV entry point in IEFW21SD 3-38
- IEFPARM statement in the RMF procedure 2-25
- IEFRDER statement in the RMF procedure 2-25
- IEFSCAN module 3-45
- IEFUSI exit
 - bypassing the storage availability check 5-3
 - changes recorded in SMF 6-23
 - limiting user region size 5-2
- IEFW21SD load module 3-14, 3-38
- IEFXVNSL load module 3-14
- IEHDASDR 2-2, 2-10, 9-8
- IEWFETCH module 2-16
- IEWMBOSV, alias for IEWFETCH 2-16

- IEWMSEPT, alias for IEWFETCH 2-16
- IFASMF module 2-7
- IFASMFDP module 7-5
- IGC0001F module 2-16
- IGC0004F module 2-16
- IGC0004G module 2-16
- IGC046 module 2-16
- IGC047 module 2-16
- IGFPEMER module 6-23
- IGFPTERM module 6-23
- IGFPTREC module 6-23
- IHAHCLOG macro 3-14
- IKJDAIR interface routine 3-14
- IKJEFD00 load module 3-14
- IKJEFT01 module 2-24
- IKJEGSCD table 3-11
- IKJEGSTA module 3-11
- IKJEGSUB macro 3-11
- ILRSLOT nucleus CSECT 8-9
- ILRSLOTV nucleus CSECT 8-9
- IMPL 4-1
- INDEX DD statement in PRDMP procedure 2-24
- indirect address notation on SLIP commands 4-12
- initializing
 - DASD 2-10
 - MVS/XA 2-12
- Insert Program Mask (IPM) instruction 3-31
- Insert Storage Key (ISK) instruction 3-12
- Installed User Programs (IUPs) 1-2
- installing MVS/XA
 - IPCS modules on other systems 6-19
 - JES2 component 2-11
 - performing a full sysgen 2-1, 2-2
 - PRDMP modules on other systems 6-19
 - using SMP Release 4 2-1
 - using SMP/E 2-1, 2-2
- instructions
 - changed 3-29
 - deleted 3-12
 - LPSW 3-27
 - new 3-31
 - recording 4-3
 - stepping through 4-3
 - STNSM 3-23
 - STOSM 3-23
- Interactive Problem Control System
 - See IPCS
- Interactive System Productivity Facility (ISPF) 6-1
- interface routines
 - IKJDAIR 3-14
 - IKJEFD00 3-14
- interfaces
 - to access methods 3-25
 - to JES2 5-5
 - to routines above 16 Mb 3-34
 - to system services 3-23
- interval timer 3-14
- INTSECT macro 3-7, 9-4
- IOCDs (I/O configuration data set)
 - creating for MVS/XA 2-3
 - using in 370 mode 2-4
- IOCP (I/O configuration program)
 - creating an IOCDs 2-3
 - I/O configuration requirements 2-5
 - macros 2-4
- IODEVICE macro 2-4, 2-5, 2-6
- IOHALT macro 3-5, 3-8
- IOS

- control blocks in PRDMP output 6-11
- modules 3-15
- unit control block
 - See UCB
- IOSCAT lock 3-18
- IOSDATA PRDMP statement 6-11
- IOGEN UCBLook macro 3-5, 3-8, 3-15
- IOSLCH lock 3-18
- IOSLOOK macro 3-7, 3-8
- IOSRHIDT module 5-4
- IOSRVC parameter in IEAIPSxx 2-20
- IOSVSUCB module 3-8
- IOS201E message 4-6
- IPCS (Interactive Problem Control System)
 - BLSPDISE panel 6-16
 - BLSPDSLE panel 6-16
 - BROWSE panel 6-17
 - dump processing exits 5-8
 - installing the MVS/XA version on other systems 6-18
 - ISPF co-requisite 6-1
 - migration aid 6-18
 - new panels 6-16
 - specifying the dump source 6-16
 - subcommands 6-13
 - titles of print files 6-18
- IPL
 - options on SYSCTL frame 4-2
 - text 2-10
- IPM (Insert Program Mask) instruction 3-30, 3-31
- IPS parameters
 - IOSRVC 2-20
 - PPGRTR 2-20, 8-9
- ISAM 3-25
- ISGGRNL0 load module 2-16
- ISG0661 message 2-26
- ISK (Insert Storage Key) instruction 3-12
- ISPF (Interactive System Productivity Facility) 6-1
- IUPs (Installed User Programs) 1-2

J

- JCL
 - allocating SYS1.DAE 2-9, 6-10
 - allocating SYS1.DUMP data sets 4-3
 - copying PRDMP and IPCS modules 6-19
 - DISP parameter 4-3
 - routing jobs 9-8
 - starting the master scheduler address space 2-9
- JES2
 - functionally equivalent levels 1-2, 2-11
 - hardcopy log records 3-13
 - interfaces 5-5
 - multi-access spool environment 9-1
 - spool access method 5-5
 - user exits 5-4
 - user modifications to 3-12
- JES3
 - converter-interpreter (CI) processing 9-9
 - DSPs (dynamic support programs) 5-5
 - functionally equivalent levels 1-2
 - hardcopy log records 3-13
 - loosely-coupled configuration 9-1
 - user exits 5-5
 - user modifications to 3-12
 - using SYS1.PROCLIB 9-9
- jobstreams for copying PRDMP and IPCS modules 6-19

K

- KEYLIST dump option 6-3

L

- LA instruction 3-30, 9-3
- labeled tapes, using for stand-alone dumps 4-3
- LCH parameter in IECIOSxx 2-23
- library
 - (See the Preface for a list of related publications)
 - changes 1-3
- limiting
 - dump output 6-10
 - user region size 5-2
- link editing
 - allocation user routines 3-14
 - programs for optimal fetch performance 8-3, 8-4
 - programs to run in MVS/XA 9-2
- LINK macro 3-5, 3-27
- link pack area modules in PRDMP output 6-11
- linkage assist routine
 - description of 3-33
 - example of 3-34
- linkage editor
 - changes to 8-2
 - in DFDS 1.4 8-2, 9-2
 - in DFP 8-2, 8-3, 9-2
 - indicating AMODE/RMODE 3-28
 - overlay structure 3-29
 - performance related changes 8-2
 - REGION parameter 8-6
 - sysgen requirement 2-2
 - text block sizes 8-6
- LISTDUMP IPCS subcommand 6-14
- LKED EXEC statement 8-6
- LLA (LNKLST lookaside)
 - directory 2-20, 2-22, 4-11, 8-7, 8-8
 - function 8-7
 - procedure 2-20, 2-24, 8-8
- LLT
 - LLTAPFIN field 3-17
 - LLTAPFTB extension 3-17
- LNKAUTH system parameter 2-22, 8-8
- LNKLST
 - concatenation 2-22, 2-23, 5-7, 8-7
 - lookaside function (See LLA)
- LOAD macro
 - changes 3-5, 3-40
 - using to determine addressing mode 3-29, 3-40
- load modules
 - See also specify module names
 - copying 8-3
 - count values 8-2, 8-4
 - link editing 9-2
 - program fetch considerations 8-1
 - reblocking 8-3
 - text block sizes 8-3, 8-6
- loading microcode 2-10, 4-1
- LOC parameter on GETMAIN 3-44
- locks
 - changes to structure of 3-18
 - determining hierarchy position 3-18
 - determining locks held 3-18
 - on the SVCTABLE macro 2-5
- LOG command 3-13
- log records
 - for system commands 3-43

- hardcopy 3-13
 - in SYSLOG data sets 3-13
 - in SYS1.LOGREC data sets 3-16, 6-18
 - prefixes 3-43
- logical
 - control units 2-5
 - path utilization 2-21
- LOGREC records
 - format changes 3-16
 - in PRDMP output 6-18
- loosely-coupled configuration
 - DSI procedures 9-9
 - JES component 9-1
 - MVS/XA and MVS/370 9-1
 - routing jobs 9-8
 - sharing data sets 9-8
- low address protection 3-19
- LPA (link pack area)
 - directory entries (LPDEs) 2-21
 - system parameter 2-9, 2-22
- LPALST concatenation 2-21, 2-23
- LPALSTxx PARMLIB members 2-9, 2-22, 2-23
- LPAMAP PRDMP statement 6-11
- LPDEs (LPA directory entries) 2-21
- LPSW instruction, changing the addressing mode 3-27
- LRA instruction 3-30
- LSQA dump option 6-3

M

- macro instructions
 - See also specific macro names
 - \$HASPEQU 3-12, 5-5
 - \$HASPGEN 3-12, 5-5
 - BTAM RESETPL 3-9
 - CHKPT 3-7
 - comprehensive list of 3-3
 - downward incompatible 9-4
 - IATYMOD 3-12, 5-5
 - incompatible parameter changes A-1
 - IOCP
 - See IOCP, macros
 - IOHALT 3-8
 - IOGEN UCBLK 3-8
 - IOSLOOK 3-8
 - new downward compatible parameters 3-38
 - RESETPL 3-9
 - SPIE 3-9
 - SPLEVEL 3-42, 9-4
 - STATUS STOP,SYNCH 3-10
 - summary of new and changed 3-38
 - SYNCH 3-3, 3-27, 3-38, 3-43, 9-4, 9-7
 - sysgen
 - See sysgen, macros
 - when MVS/XA expansions are required 9-6
- Mass Storage Subsystems 2-10
- master catalog
 - LPALIB entries 2-9
 - making a backup copy 2-3
 - SYS1.LOGREC entry 2-9
- master scheduler address space 2-9
- master trace table in PRDMP output 6-11
- MAXBLK parameter on IEBCOPY 8-3
- MAXUSER system parameter 2-15, 2-22
- message processing facility (MPF) 4-5, 5-7
- messages
 - controlling traffic 4-5, 4-8, 5-7
 - CSV300I 8-1
 - displaying in color 4-5

- format of display 4-7
- from PRDMP processing 2-24
- highlighting 2-23, 4-5
- IAR004I 2-12
- ICP404I 2-4
- IEA240I 2-22
- IEA838I 6-8
- IEE978E 2-23
- IOS201E 4-6
- ISG066I 2-26
- modifying processing of 5-7
- reporting suppressed dumps 6-8
- retaining 2-23
- routing 4-5, 4-8
- suppressing 2-23
- MF/1 (System Activity Measurement Facility) 2-3
- MGCR macro 3-7, 3-40
- microcode, loading 2-10, 4-1
- migration aids
 - IPCS 6-18
 - PRDMP 6-18
- MIH (missing interrupt handler)
 - intervals, specifying 2-17, 2-23
 - parameter in IECIOSxx 2-17
- MINBLK parameter on IEBCOPY 8-3
- missing interrupt handler
 - See MIH
- MLPA
 - building 2-9
 - page protection 3-19
 - system parameter 2-21
- MODE command 4-11
- MODESET macro 3-7
- MODIFY command 4-11, 8-8
- modules
 - See load modules and specific module names
- MOD88 service routine 3-18
- MONITOR command 4-11, 4-13
- Move Long (MVCL) instruction 3-30
- MPF (message processing facility) 4-5, 5-7
- MPFLSTxx PARMLIB member 2-23, 4-5, 5-7
- MSCTC (MSC table create) utility 2-10
- MSGRT command 4-11
- MSS, loading microcode EC tapes for 2-10
- MSTRJCL system parameter 2-9, 2-22
- MSTRJCLxx member of SYS1.LINKLIB 2-9, 2-22
- MSTRJCL00 member of SYS1.LINKLIB 2-9
- MTRACE PRDMP statement 6-11
- MVCL (Move Long) instruction 3-30

N

- new function for programs 3-2
- non-specific device allocation 2-21
- non-standard tape label routine 3-14
- NOPROT option on system parameters 2-21, 3-19
- NOSYM parameter in PARMLIB members 6-5
- NUC dump option 6-3
- nucleus
 - dumping 2-19, 6-3
 - in PRDMP output 6-11
 - read-only 3-18
- NUCLKUP macro 3-7, 3-41
- NUCMAP
 - area in the nucleus 3-19
 - PRDMP statement 6-11
- NVESQA fields 2-13

NVSQA fields 2-13
NVTNVSQA field 2-13

O

OPEN

- IPCS subcommand 6-14
- processing, VSAM 5-1
- operating system
 - initializing 2-12
 - installing 2-1
 - IPL option 4-2
 - restarting 4-2
- operator control (OPRCTL) console frame 4-2
- operator response to message IOS201E 4-6
- OPRCTL (operator control) console frame 4-2
- overlay modules
 - fetching 8-2
 - inserting count values 8-3
 - reblocking 8-3
 - restrictions on using 3-29

P

- page data sets
 - defining 2-8
 - eligible devices for 2-8
- page fix appendage 3-20
- page protection
 - areas protected 3-18
 - of IEAFIXxx modules 2-20, 2-21
 - of IEALPAXx modules 2-21
 - turning off 3-19
- page-in rate for address spaces 2-20, 8-9
- PAK
 - lists 2-21
 - system parameter 2-22
- PAM 3-25
- parameters
 - See system parameters
- PARMLIB members
 - See SYS1.PARMLIB data set
- patch area 3-20
- PC instruction, changing the addressing mode 3-27
- PDS directory entry, AMODE/RMODE specifications in 3-28, 9-2
- performance considerations
 - page-in rate of an address space 8-9
 - paging algorithms 8-9
 - program fetch processing 8-1
 - SMF data set placement 8-9
 - SMF data set processing 7-5
 - using the ASM backing slot function 8-9
- performing
 - a full sysgen 2-1, 2-2
 - I/O in 31-bit addressing mode 3-36
 - IMPL 4-1
- PGFIX macro 9-4
- PGFREE macro 9-4
- PGFX (page fix appendage) 3-20
- PGLOAD macro 9-4
- PGOUT macro 9-4
- PGRLSE macro 9-4
- PGSER
 - macro 3-5, 3-41, 9-4
 - service routine 3-18
- PLPA
 - building 2-9
 - page protection 3-19

- positioning tasks
- POST exits 3-17
- post-dump exit routines 5-4
- power-on-reset function, performing 4-1
- PPGRTR parameter in IEAIPSxx 2-20, 8-9
- PPT (programming properties table) 2-7
- PRDMP
 - command processor 2-24
 - control statements 6-10
 - EREP exit 6-18
 - exit control table 5-1, 5-6
 - exit routines 5-6, 5-8
 - header exits 5-6
 - index
 - inserting user entries 6-12
 - obtaining before the dump 2-25, 6-12
 - length of output lines 6-13
 - migration aid 6-18
 - output buffer 5-6
 - procedure in SYS1.PROCLIB 2-24, 6-19
- PRDMPXA member of SYS1.SAMPLIB 6-19
- pre-dump exit 5-4
- preferred path 2-4
- preparing for MVS/XA 1-2
- print dump
 - See also PRDMP
 - exit control table 5-1
 - macro (PRDMP) 6-11
- printer requirements for PRDMP 6-13
- private area storage
 - minimizing amount lost because of rounding 2-14
 - reporting use of 7-2
- procedures
 - See SYS1.PROCLIB data set
- processor addresses 3-17
- program fetch
 - amount of virtual storage fixed 8-4
 - differences 8-2
 - performance 8-1
- program mask, obtaining 3-30
- program products
 - BTAM/SP 3-9
 - DEMF 2-3
 - Device Support Facilities Release 6 2-10
 - DFDSS 2-2
 - DFEF 9-9
 - EREP 6-18
 - IPCS 6-1
 - ISPF 6-1
 - licenses 9-7
 - MF/1 2-3
 - RMF
 - See RMF
- program status word
 - See PSW (program status word)
- programming considerations 3-1
- programming properties table (PPT) 2-7
- programs requiring modification
 - authorized 3-1
 - unauthorized 3-1
- PSA
 - changes to 3-20
 - fetch protection 3-20
 - low address protection 3-19
 - patch area 3-20
 - PSACLHS field 3-18
 - PSAHLHI field 3-18
 - work/save area locations 3-20

PSACLHS field 3-18
PSAHLHI field 3-18
PSW (program status word)
 addresses in 3-17
 addressing mode bit 3-25
PT instruction, changing the addressing mode 3-27
PTFs
 for GTF module AHLTSVC 3-8
 for IOHALT 3-8
 for IOSGEN UCBLOOK 3-9
 for JES2 1-2
 for program fetch 8-2
 for SVC 33 3-8
 for the DFDS 1.4 linkage editor 8-5
PTRACE macro 3-7, 3-41
publications
 (See the Preface for a list of related publications)
 changes 1-3
PURGE service routine 2-16, 3-18

Q

queuing messages 5-7

R

RACROUTE macro 3-41
real addresses
 in IDAWs 3-20
 in LSQA 3-22
 in SQA 3-22
 in the nucleus 3-22
 topics related to using 3-20
 using with EXCPVR 3-20
real storage dump module (AMDSARDM) 2-27
real time interval, setting 3-42
reason codes 3-39, 6-23
reblocking modules 8-3
reconfigurable storage, specifying 2-12
reconfiguring the system 2-18
records
 control 8-2
 GTF 3-16
 hardcopy log 3-13
 in SYSLOG 3-13, 3-43
 LOGREC 3-16
 RLD 8-2
 RLD/control 8-2
 SMF 3-12, 7-5
 system trace 3-16
 text 8-1, 8-2, 8-6
recovery actions for hot I/O 4-3
recovery termination manager (RTM) 3-15
REGION parameter
 compared to using IEFUSI 5-2
 on link edit jobs 8-6
 specifying more than 16 Mb 5-3
 storage availability check 5-3
region size
 exceeding 6-23
 extended 5-3
 specifying 5-2
relocation dictionary (RLD) records 8-2
RENUM IPCS subcommand 6-14
reports, RMF 4-6
RESERVE macro
 limiting concurrent requests 3-12
 summary of changes 3-5, 3-41

RESERVE requests, global 2-26
reserving ASVT entries 2-15, 2-22
RESETPL macro 3-1, 3-5, 3-9
residency time 2-20, 8-9
resident BLDL list 3-19
Resource Measurement Facility
 See RMF
resource name lists
 See RNLs
restart processing 4-2
restarting
 options 4-2
 processors 4-6
 SMF 4-12
retaining messages 2-23, 5-7
retrieving data above 16 Mb 3-35
RETURN macro 3-5, 3-41
RLD (relocation dictionary) records 8-2
RMF (Resource Measurement Facility)
 duration of initialization process 2-26
 modules 3-15
 Monitor II reports 4-6
 obtaining storage for I/O data 2-15
 post processors 7-5
 procedure in SYS1.PROCLIB 2-25
 SMF records 70-79 7-5
 starting 2-25
 user exit routines 5-4
RMODE
 description of 3-27
 determining 3-29
 flags in the CESD 3-28, 9-2
 flags in the ESD 3-28
 flags in the PDS directory entry 3-28, 9-2
 specifying 3-28
RNLDEF statements 2-16
RNLs (resource name lists)
 defining 2-17
 in GRSRNLxx PARMLIB members 2-16
 SYSTEM inclusion 5-2
 SYSTEMS exclusion 2-26, 5-2
 using defaults 2-26
routing
 codes for messages
 altering 5-7
 using 4-5
 jobs 9-8
 messages 4-8
RPQs for devices and features 1-3, 2-6
RSM
 backing virtual storage 3-21
 control blocks in PRDMP output 6-11
 modules 3-15
RSMDATA PRDMP statement 6-11
RSU system parameter 2-12
RSVNONR system parameter 2-15, 2-22
RSVSTRT system parameter 2-15, 2-22
RTM (recovery termination manager) 3-15
RUNCHAIN IPCS subcommand 6-14

S

SADMP
 See stand-alone dump
SALLOC lock 3-18
SAM 3-25
save areas in the PSA 3-20

SCHEDULE macro 3-7, 9-4
 SCHEDULR macro 2-5, 2-6
 SCP manual CNTL (SYSTCL) console frame 4-1
 SDUMP
 DAE function 6-8
 macro 3-7, 3-41, 3-44, 9-4
 SDWA
 additional information in 6-22
 changes to structure 3-10
 checkpoint/restart data 6-24
 SDWAVRA 6-9
 segment protection in PLPA 3-19
 select ASID service 5-10
 serializing VSAM data sets 5-1
 SET
 command
 DAE 4-12, 6-9
 MPF 4-5, 5-7
 SMF 4-12
 symbol 9-1
 Set Storage Key (SSK) instruction 3-12
 SETDEF IPCS subcommand 6-14
 SETFRR macro 3-7
 SETLOCK macro 3-7
 example of new function 3-18
 incompatible MVS/XA expansion 9-4
 new parameters 3-41
 specifying RELEASE,TYPE=(reg) | ALL 3-44
 SETRP macro 3-5, 3-42, 6-23
 SIZE parameter for LKED EXEC 8-6
 SLIP command
 in IEACMD00 2-12, 2-20
 in summary of commands 4-13
 MOD 4-12
 SET 4-12
 suppressing dumps 6-7
 31-bit indirect address notation 4-12
 slot selection algorithm 8-10
 slots, backup 8-9
 SMF
 address space 2-23, 2-24
 BQEs 3-17
 buffers 2-23
 compatibility between releases 7-5
 data set placement 8-9
 EOF marks 3-17, 7-5
 format of data sets 3-17
 recording IEFUSI changes 6-23
 recording TSO commands 3-12
 records 3-16, 7-2, 7-3, 7-4, 7-5
 reporting device connect time 7-1
 reporting virtual storage use 7-2
 step initiation exit 5-2
 SMFEOFMARKs in SMF records 3-17, 7-5
 SMFEXIT macro 3-5, 9-4, A-3
 SMFIOCNT macro 3-6, 3-17, 3-42
 SMF30ARB field 7-2, 7-3
 SMF30BLK field 7-3
 SMF30DCT field 7-1
 SMF30EAR field 7-2
 SMF30ERG field 7-2
 SMF30EUR field 7-2
 SMF30PRV field 7-2
 SMF30RGB field 7-2
 SMF30RGN field 7-4
 SMF30SYS field 7-2
 SMF30TCN field 7-1
 SMF30TEP field 7-3
 SMF30URB field 7-2, 7-3
 SMF32TCT field 7-1
 SMF4EXCP field 7-3
 SMF4RSHO field 7-4
 SMP Release 4 2-1, 2-2
 SMP/E 2-1, 2-2
 SNAP
 dump headers 6-6
 dump indexes 6-6
 dump processing exits 5-8
 macro 3-6, 3-42
 SPIE macro 3-6, 3-9, 9-4
 SPLEVEL macro 3-6
 examples 9-5
 function of 3-42, 9-5
 in JES system programs 3-12
 use in JES modifications 3-12
 use in JES2 user exits 5-5
 use in JES3 user exits 5-5
 SPLS dump option 6-4
 SQA
 dump option 6-4
 increasing minimum allocation for 2-13
 specifying the size of 2-13, 2-21
 system parameter 2-13, 2-21
 SRB
 SRBEP field 3-15
 SRBRMTR field 3-15
 SRM (system resources manager)
 calculating I/O service 2-20
 calculating page-in rate 2-20, 8-9
 collecting I/O data 2-15
 I/O interrupt processing 2-21
 I/O load balancing 2-21
 IOSRVC parameter 2-20
 modules 3-15
 non-specific device allocation 2-21
 PPGRTR parameter 2-20, 8-9
 SSI (subsystem interface) routines 5-5
 SSK (Set Storage Key) instruction 3-12
 STACK IPCS subcommand 6-14
 STAE macro 9-3
 stand-alone dump
 generating 2-26
 IPL option 4-2
 macro (AMDSADMP) 2-27
 real storage dump module (AMDSARDM) 2-27
 requesting 2-27
 storing status 4-3
 using labeled tapes 4-3
 START command
 LLA keyword 2-20, 2-24, 4-13, 8-8
 SUB keyword 4-13
 starting
 DUMPSRV address space 2-24
 full-function address spaces 2-24
 LLA function 2-24, 8-8
 master scheduler address space 2-9
 PRDMP 2-24
 RMF 2-25
 SMF address space 2-24
 STATUS STOP,SYNCH macro 3-6, 3-10
 STAX macro 3-6, 9-4, A-3
 STIMER
 macro 3-6, 3-14, 3-42, 9-4, A-3
 service routine 2-16, 3-18
 STIMERM macro 3-6, 3-42
 STNSM instruction 3-23
 STOP command 4-13, 8-8
 STOPMN command 4-11, 4-13

- storage
 - availability check 5-3
 - management locks 3-18
 - specifying reconfigurable 2-12
- storing status before taking a stand-alone dump 4-3
- STOSM instruction 3-23
- SUBPLST dump option 6-4
- subpools 3-17, 3-22
- subsystem interface (SSI) routines 5-5
- SUBTASKS dump option 6-4
- SUM dump option 6-4
- SUMMARY IPCS subcommand 6-15
- SUMMARY PRDMP statement 6-11
- suppressing
 - dumps
 - preventing 6-9
 - using DAE 6-8
 - using SLIP commands 6-7
 - messages 2-23, 5-7
- SVC
 - changing the addressing mode 3-27
 - dumps
 - checkpoint/restart data 6-24
 - DAE options for 6-8
 - format changes 6-6
 - suppressing 6-7
 - issued by WTO/WTOR user exits 5-7
 - Router 3-18
 - screening table addresses 3-15
 - table
 - changes 3-18
 - updating 3-43
 - 109 3-18
 - 138 3-18
 - 16 3-18
 - 33 3-8
 - 46 3-18
 - 47 3-18
 - 61 3-11
 - 82 3-18
 - 88 3-18
 - 97 3-11
- SVCDUMP modules 3-15
- SVCTABLE macro 2-5
- SVCUPDTE macro 3-7, 3-43
- SVT
 - SVTDACTV field 3-17
 - SVTPWAIT field 3-17
- swap data sets
 - defining 2-8
 - eligible devices for 2-8
- symptom
 - data from DAE 5-6, 6-8, 6-11
 - dump
 - description 6-4
 - obtaining via TSO 6-5
 - suppressing 6-5
- SYNCH macro 3-3, 3-6, 3-27, 3-38, 3-43, 9-4, 9-7, A-4
- SYSABEND dumps
 - headers 6-6
 - indexes 6-6
 - options 2-19
 - summary dump 6-5
 - suppressing 6-7
- SYSCTL (SCP Manual CNTL) console frame 4-1
- sysgen
 - creating an IOCDS 2-3
 - defining devices 2-3
 - defining system data sets 2-8
 - DLIB changes 2-3
 - functions deleted 2-3
 - initializing DASD 2-10
 - IPL text required 2-10
 - macros
 - CHANNEL 2-6
 - CONSOLE 2-5
 - CTRLPROG 2-6, 2-13
 - DATASET 2-5, 2-6, 2-9
 - GENERATE 2-5, 2-6, 2-10
 - incompatible differences 2-5
 - IODEVICE 2-5, 2-6
 - SCHEDULR 2-5, 2-6
 - SVCTABLE 2-5
 - UNITNAME 2-6
 - performing 2-1, 2-2
 - requirements for 2-2
 - SYS1.LOGREC placement 2-9
- SYSLIB DD statements 8-6
- SYSLIN DD statements 8-6
- SYSMOD DD statements 8-6
- SYSLOG data set records 3-13, 3-43
- SYSMDUMP dumps
 - DAE options for 6-8
 - format changes 6-6
 - headers 6-6
 - options 2-20
 - summary dump 6-6
 - suppressing 6-7
 - symptom dumps 6-5
- SYSPRINT DD statement 2-24, 8-6
- system data sets
 - defining during sysgen 2-8
 - dump data sets 2-8
 - eligible device types for 2-8
 - incompatible 9-9
 - page data sets 2-8
 - sharing 6-10
 - swap data sets 2-8
 - SYS1.DAE 2-8, 6-8
 - SYS1.DUMPx 1-1, 2-8, 2-15
 - SYS1.LOGREC 2-9
 - SYS1.PARMLIB 2-10
 - SYS1.PROCLIB 2-23
 - using the SYS1 qualifier 2-10
- system generation
 - See sysgen
- SYSTEM inclusion RNLs 5-2
- system log 3-43
- system parameters
 - ALT 2-21
 - BLDL 2-20, 2-22
 - BLDLF 2-20, 2-22
 - CMB 2-15, 2-21
 - CSA 2-13, 2-14, 2-21
 - DUMP 2-15, 2-21
 - FIX 2-21
 - GRSRNL 2-16, 2-22
 - LNKAUTH 2-22, 8-8
 - LPA 2-9, 2-22
 - MAXUSER 2-15, 2-22
 - MLPA 2-21
 - MSTRJCL 2-9, 2-22
 - PAK 2-22
 - RSU 2-12
 - RSVNONR 2-15, 2-22
 - RSVSTR 2-15, 2-22
 - SQA 2-13, 2-21
- system patch area 3-20

- system resources manager
 - See SRM (system resources manager)
- system services
 - interfaces to 3-23
 - parameter list changes A-1
- system termination facility 6-23
- system trace
 - activating 6-21
 - buffers 6-21
 - changes to 6-20
 - creating entries 6-21
 - data in dumps 2-19, 6-21
 - modules 3-15
 - records 3-16
 - selecting events 6-20
 - table 4-14, 6-12
 - types of 6-20
- SYSTEMS exclusion RNLs 2-26, 5-2
- SYSTSIN DD statement 2-24
- SYSTSPRT DD statement 2-24
- SYSUDUMP dumps
 - headers 6-6
 - indexes 6-6
 - options 2-20
 - summary dump 6-5
 - suppressing 6-7
- SYSUT1 DD statements 8-6
- SYS1.AOSC5 data set 2-3
- SYS1.DAE data set 2-8, 2-9, 6-8, 6-10
- SYS1.DUMPxx data sets
 - See dump, data sets
- SYS1.LINKLIB data set
 - location of RNLs 2-16, 2-22
 - MSTRJCLxx members 2-9, 2-22
 - sharing 9-9
- SYS1.LOGREC data sets
 - checkpoint/restart data 6-24
 - increasing the size of 2-9
 - placement 2-9
 - recording suppressed dumps 6-8
- SYS1.LPALIB data set
 - concatenation 2-9
 - sharing 9-9
- SYS1.MACLIB data set
 - different expansions of same macro 9-4
 - sharing 9-9
- SYS1.NUCLEUS data set
 - allocating 2-10
 - sharing 9-9
- SYS1.PARMLIB data set
 - characteristics 2-10
 - members
 - ADYSETxx 2-18, 6-9
 - ADYSET00 6-9
 - ADYSET01 6-9
 - ADYSET02 6-9
 - COMMNDxx 2-16
 - CONFIGxx 2-18
 - GRSRNLxx 2-16, 2-18, 2-22
 - GRSRNL00 2-17
 - GTFPARM 2-18
 - IEAABD00 2-19
 - IEABLDxx 2-20
 - IEACMD00 2-20, 8-8
 - IEADMP00 2-20
 - IEADMR00 2-20
 - IEAFIXxx 2-20
 - IEAIPSxx 2-20
 - IEALOD00 2-21
 - IEALPAXx 2-16, 2-21
 - IEAOPTxx 2-21
 - IEAPAKxx 2-21
 - IEAPAK00 2-21
 - IEASYSxx 2-21
 - IECIOSxx 2-23, 5-4
 - LPALSTxx 2-9, 2-23
 - MPFLSTxx 2-23, 4-5, 5-7
- sharing 9-9
- summary of changes to 2-17
- SYS1.PROCLIB data set
 - DUMPSRV procedure 2-24
 - IEESYSAS procedure 2-24
 - in a JES3 configuration 9-9
 - in converter-interpreter (CI) processing 9-9
 - LLA procedure 2-24, 8-8
 - PRDMP procedure 2-24
 - RMF procedure 2-25
 - summary of changes 2-23
- SYS1.SAMPLIB data set
 - BLSAMPLE member 6-19
 - DAEALLOC member 2-9, 6-10
 - MIGJOB01 and MIGJOB 02 members 6-19
 - PRDMPXA member 6-19
- SYS1.SBLSMG0 data set 6-19, 6-20
- SYS1.SBLSPNL0 data set 6-19, 6-20
- SYS1.SVCLIB data set, sharing 9-9

T

- tapes, labeled 4-3
- TCBSVCA2 field 3-15
- TCOMTAB control block 3-11
- terminating
 - jobs 4-7, 4-11
 - started processes 4-7, 4-11
 - time-sharing users 4-7, 4-11
- text records, sizes of 8-1, 8-6
- thresholds
 - for hot I/O interrupts 2-23
 - for I/O interrupt processing 2-21
 - for I/O load balancing 2-21
 - for limiting concurrent global resource serialization requests 3-13
 - for logical path utilization 2-21
- time
 - execution 2-20, 8-9
 - residency 2-20, 8-9
- time interval, real 3-42
- timer
 - CPU 3-14
 - interval 3-14
- titles of IPCS print files 6-18
- TIVEFRGN field in type 34 SMF records 7-4
- TRACE
 - command 4-14, 6-20
 - instruction 3-31
 - lock 3-18
 - PRDMP statement 6-12
- tracing
 - See also system trace
 - USR events via GTF 3-45
- TRACK command 4-5
- Translate and Test (TRT) instruction 3-30
- translating real addresses to virtual addresses 3-15
- TRT
 - dump option 6-4
 - instruction 3-30

TSO

- command package 2-3
- obtaining symptom dump output in 6-5
- terminal monitor program 2-24
- TEST command 2-3, 3-11
- TEST subcommand table 3-11

TSO/E for MVS/XA 2-3, 3-11

TTIMER

- macro 3-14, 3-39
- service routine 2-16, 3-18

U

UCB

- addresses 3-15
- look-up routine 3-15
- scan routine (IOSVSUCB) 3-8

UNALLOC command 3-11

unauthorized programs, changes affecting 3-1

unit control block

See UCB

unit verification 3-45

UNITNAME macro 2-6

user exit routines

- dump processing 5-8
- dynamic allocation 3-14
- EREP PRDMP 6-18
- IEALIMIT 5-2
- IEAVMXIT 4-7, 5-7
- IEECVXIT 4-5, 5-7
- IEFDB401 3-14
- IEFUSI 5-2
- JES2 5-4
- JES3 5-5
- post-dump 5-4
- PRDMP 5-6
- PRDMP header 5-6
- pre-dump 5-4
- print dump 5-1
- RMF 5-4
- SMF step initiation exit 5-2
- using ECT entries 5-1
- WTO 4-5
- WTO/WTOR 2-23, 4-5, 5-7

V

V=R programs 3-16

VARY command

- CH 4-14
- CPU 4-14
- PATH 4-14
- STOR 4-14

virtual IDAWs 3-37

virtual storage

- amount program fetch fixes 8-4
- changes in use of 7-2
- for RMF I/O measurements 2-15
- for the SLIP command processors 2-12
- map 2-14
- obtaining information about 3-43
- obtaining via GETMAIN 3-10
- reporting use of 7-2
- rules for backing 3-21

VRADAE key in the SDWAVRA 6-9

VRADATA macro 6-10

VRC parameter on GETMAIN 3-43

VRU parameter on GETMAIN 3-43

VSAM

- catalogs 9-9
- data sets 5-1
- interfaces to 3-25
- OPEN processing 5-1
- performing I/O in 31-bit mode 3-36
- record management load modules 3-15

VSM

- control blocks in PRDMP output 6-12
- GETMAIN limit
 - calculating 5-3
 - exceeding 6-23
 - specifying 5-2
- modules, residence of 3-15
- region size limit
 - calculating 5-3
 - exceeding 6-23
 - specifying 5-2
- storage availability check 5-3
- VSMDATA PRDMP statement 6-12
- VSMLIST macro 3-7, 3-43
- VSMLOC macro 3-7, 3-43
- VSMREGN macro 3-7, 3-43

W

wait state codes

- 0C4 6-23
- 081 2-10
- 114 4-6

work/save areas in the PSA 3-20

working set size 2-20, 8-9

WTL macro 3-6, 3-13, 3-43

WTO macro 3-6, 3-43

WTO/WTOR user exits 2-23, 4-5, 5-7

WTOR macro 3-6, 3-43, 9-4, A-4

X

XCTL macro

- changing the addressing mode 3-27
- differences 3-6

0

0C4 wait state code 6-23

081 wait state code 2-10

1

114 wait state code 4-6

16E ABEND code 3-17

2

24-bit dependencies in programs 9-3

3

31-bit addressing

- description of 3-25
- impact on programmers 3-26

indirect addresses on SLIP commands 4-12
list of related topics 3-2
modules running in 3-15
3279 MCS consoles 4-5
370 I/O instructions 3-12
370-XA mode
 initializing processor in 4-1
 instruction addresses 3-25
 IOCP differences 2-4
 switching to 4-1

5
504 ABEND code 3-10
538 ABEND code 3-13

8
80A ABEND code 6-8



This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you. Your comments will be sent to the author's department for whatever review and action, if any, are deemed appropriate.

Note: *Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.*

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

Please use pressure sensitive or other gummed tape to seal this form.

Cut or Fold Along Line

What is your occupation? _____

How do you use this publication? _____

Number of latest Newsletter associated with this publication: _____

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Cut or Fold Along Line

Fold and tape

Please Do Not Staple

Fold and tape



BUSINESS REPLY MAIL
FIRST CLASS PERMIT 40 ARMONK, NEW YORK

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



POSTAGE WILL BE PAID BY ADDRESSEE:

International Business Machines Corporation
Department D58, Building 920-2
PO Box 390
Poughkeepsie, New York 12602

Fold and tape

Please Do Not Staple

Fold and tape



You can use this form to submit conversion hints for possible publication in this book. It is understood that IBM and its affiliated companies shall have the nonexclusive right, in their discretion, to use, copy and distribute all submitted information or material, in any form, for any and all purposes, without any obligation to the submitter, and that the submitter has the unqualified right to submit such information or material upon such basis. When submitting conversion hints, indicate the system and release level from which you are

converting and the system and release level to which you are converting.

Your views about this publication might help improve usefulness; this form will be sent to the author's department for appropriate action. Using this form to request system assistance or additional publication will delay response, however. For more direct handling of such requests, please contact your IBM representative or the IBM Branch Office serving your locality.

What is your occupation? _____

Number of latest Technical Newsletter (if any) concerning this publication: _____

Please indicate your name, address, and phone number in the space below.

IBM representative. _____

State problem and suggested solution for resolving problem. Attach examples, etc., when available.

Please use pressure sensitive or other gummed tape to seal this form.
Cut or Fold Along Line

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments).

Input Form

Cut or Fold Along Line

Fold and tape

Please Do Not Staple

Fold and tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT 40 ARMONK, NEW YORK



POSTAGE WILL BE PAID BY ADDRESSEE:

International Business Machines Corporation
Department D58, Building 920-2
PO Box 390
Poughkeepsie, New York 12602

Fold and tape

Please Do Not Staple

Fold and tape

