

RSX-11M/M-PLUS

Task Builder Manual

Order No. AA-L680B-TC

RSX-11M Version 4.1
RSX-11M-PLUS Version 2.1

First Printing, June 1979
Revised, January 1982
Revised, April 1983

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may be used or copied only in accordance with the terms of such license.

No responsibility is assumed for the use or reliability of software on equipment that is not supplied by Digital Equipment Corporation or its affiliated companies.

Copyright © 1979, 1982, 1983 by Digital Equipment Corporation
All Rights Reserved.

Printed in U.S.A.

The postpaid READER'S COMMENTS form on the last page of this document requests the user's critical evaluation to assist in preparing future documentation.

The following are trademarks of Digital Equipment Corporation:

DEC	DIBOL	RSX
DEC/CMS	EduSystem	UNIBUS
DEC/MMS	IAS	VAX
DECnet	MASSBUS	VMS
DECsystem-10	PDP	VT
DECSYSTEM-20	PDT	digital
DECUS	RSTS	
DECwriter		

ZK2250

HOW TO ORDER ADDITIONAL DOCUMENTATION

In Continental USA and Puerto Rico call 800-258-1710
In New Hampshire, Alaska, and Hawaii call 603-884-6660
In Canada call 613-234-7726 (Ottawa-Hull)
800-267-6146 (all other Canadian)

DIRECT MAIL ORDERS (USA & PUERTO RICO)*

Digital Equipment Corporation
P.O. Box CS2008
Nashua, New Hampshire 03061

*Any prepaid order from Puerto Rico must be placed
with the local Digital subsidiary (809-754-7575)

DIRECT MAIL ORDERS (CANADA)

Digital Equipment of Canada Ltd.
940 Belfast Road
Ottawa, Ontario K1G 4C2
Attn: A&SG Business Manager

DIRECT MAIL ORDERS (INTERNATIONAL)

Digital Equipment Corporation
A&SG Business Manager
c/o Digital's local subsidiary or
approved distributor

Internal orders should be placed through the Software Distribution Center (SDC), Digital Equipment Corporation, Northboro, Massachusetts 01532

CONTENTS

	Page
PREFACE	xv
SUMMARY OF TECHNICAL CHANGES	xix
CHAPTER 1	INTRODUCTION AND COMMAND SPECIFICATIONS
1.1	TASK COMMAND LINE 1-2
1.1.1	Printing the Map File 1-2
1.1.2	Omitting Specific Output Files 1-3
1.2	MULTILINE INPUT 1-3
1.3	OPTIONS 1-4
1.4	MULTIPLE TASK SPECIFICATIONS 1-5
1.5	INDIRECT COMMAND FILES 1-5
1.6	COMMENTS IN LINES 1-8
1.7	FILE SPECIFICATIONS 1-8
1.8	SUMMARY OF SYNTAX RULES 1-10
CHAPTER 2	TASK BUILDER FUNCTIONS
2.1	LINKING OBJECT MODULES 2-1
2.1.1	Allocating Program Sections 2-2
2.1.1.1	Access-code and Allocation-code 2-5
2.1.1.2	Type-Code and Scope-Code 2-7
2.1.2	Resolving Global Symbols 2-7
2.2	THE TASK STRUCTURE 2-8
2.3	OVERLAYS 2-10
2.4	ADDRESSING CONCEPTS 2-13
2.4.1	Physical, Virtual, and Logical Addresses 2-13
2.4.2	Unmapped Systems 2-14
2.4.3	Mapped Systems 2-14
2.4.4	Regions 2-18
2.5	TASK MAPPING AND WINDOWS 2-20
2.5.1	Task Windows 2-20
2.6	RSX-11M-PLUS SUPERVISOR MODE 2-24
2.6.1	Supervisor-Mode Mapping 2-24
2.7	PRIVILEGED TASKS 2-25
2.8	MULTIUSER TASKS (RSX-11M-PLUS ONLY) 2-28
2.9	USER-MODE I- AND D-SPACE TASKS (RSX-11M-PLUS) 2-28
CHAPTER 3	OVERLAY CAPABILITY
3.1	OVERLAY STRUCTURES 3-1
3.1.1	Disk-Resident Overlay Structures 3-2
3.1.2	Memory-Resident Overlay Structures (Not Supported on RSX-11S) 3-5
3.2	OVERLAY TREE 3-15
3.2.1	Loading Mechanism 3-16
3.2.2	Resolution of Global Symbols in a Multisegment Task 3-16

CONTENTS

		Page
3.2.3	Resolution of Global Symbols from the Default Library	3-18
3.2.4	Allocation of Program Sections in a Multisegment Task	3-19
3.3	OVERLAY DATA STRUCTURES AND RUN-TIME ROUTINES	3-19
3.3.1	Overlaid Conventional Task Structures	3-20
3.3.2	Overlaid I- and D-Space Task Structures	3-21
3.4	OVERLAY DESCRIPTION LANGUAGE	3-23
3.4.1	.ROOT and .END Directives	3-23
3.4.2	.FCTR Directive	3-25
3.4.3	Arguments for the .FCTR and .ROOT Directives	3-25
3.4.3.1	Named Input File	3-25
3.4.3.2	Specific Library Modules	3-26
3.4.3.3	A Library to Resolve References Not Previously Resolved	3-26
3.4.3.4	A Section Name Used in a .PSECT Directive	3-26
3.4.3.5	A Segment Name Used in a .NAME Directive	3-26
3.4.4	Exclamation Point Operator	3-26
3.4.5	.NAME Directive	3-27
3.4.5.1	Example of The Use of The .NAME Directive	3-28
3.4.6	.PSECT Directive	3-29
3.4.7	Indirect Command Files	3-30
3.5	MULTIPLE-TREE STRUCTURES	3-30
3.5.1	Defining a Multiple-Tree Structure	3-30
3.5.1.1	Defining Co-trees With a Null Root by Using .NAME	3-31
3.5.2	Multiple-Tree Example	3-31
3.6	CREATING AN ODL FILE FROM A VIRTUAL ADDRESS SPACE ALLOCATION DIAGRAM	3-35
3.6.1	Creating a .ROOT Statement by Using a Virtual Address Space Allocation Diagram	3-37
3.6.2	Creating a .FCTR Statement by Using a Virtual Address Space Allocation Diagram	3-38
3.6.3	Creating an ODL Statement for a Co-Tree by Using a Virtual Address Space Diagram	3-39
3.7	OVERLAYING PROGRAMS WRITTEN IN A HIGH-LEVEL LANGUAGE	3-40
3.8	EXAMPLE 3-1: BUILDING AN OVERLAY	3-41
3.9	SUMMARY OF THE OVERLAY DESCRIPTION LANGUAGE	3-49

CHAPTER 4 OVERLAY LOADING METHODS

4.1	AUTOLOAD	4-1
4.1.1	Autoload Indicator	4-2
4.1.2	Path Loading	4-3
4.1.3	Autoload Vectors	4-4
4.1.4	Autoloadable Data Segments	4-7
4.2	MANUAL LOAD	4-7
4.2.1	MACRO-11 Manual Load Calling Sequence	4-7
4.2.2	MACRO-11 Manual Load Calling Sequence For I- and D-Space Tasks	4-8
4.2.3	FORTRAN Manual Load Calling Sequence	4-9
4.2.4	FORTRAN Manual Load Calling Sequence for I- and D-Space Tasks	4-11
4.3	ERROR HANDLING	4-11
4.4	GLOBAL CROSS-REFERENCE OF AN OVERLAID TASK	4-12
4.5	USE AND SIZE OF OVERLAY RUNTIME ROUTINES	4-14

CHAPTER 5 SHARED REGION CONCEPTS AND EXAMPLES

5.1	SHARED REGIONS DEFINED	5-1
5.1.1	The Symbol Definition File	5-4
5.1.2	Position-Independent Shared Regions	5-5

CONTENTS

	Page
5.1.2.1	Position-Independent Shared Region Mapping . . . 5-5
5.1.2.2	Specifying a Position-Independent Region . . . 5-5
5.1.3	Absolute Shared Regions 5-7
5.1.3.1	Absolute Shared Region Mapping 5-7
5.1.3.2	Specifying an Absolute Shared Region 5-7
5.1.3.3	Absolute Shared Region .STB File 5-9
5.1.4	Shared Regions with Memory-Resident Overlays . . 5-9
5.1.4.1	Considerations About Building an Overlaid Shared Region 5-9
5.1.4.2	Example of Building a Memory-Resident Overlaid Shared Region 5-10
5.1.4.3	Options for Use in Overlaid Shared Regions . 5-10
5.1.4.4	Autoload Vectors and .STB Files for Overlaid Shared Regions 5-11
5.1.5	Run-Time Support for Overlaid Shared Regions . 5-13
5.1.6	Linking to a Shared Region 5-13
5.1.7	Number and Size of Shared Regions 5-17
5.1.8	Example 5-1: Building and Linking to a Common in MACRO-11 5-17
5.1.9	Linking Shared Regions Together 5-25
5.1.10	Example 5-2: Building and Linking to a Device Common in MACRO-11 5-26
5.1.11	Example 5-3: Building and Linking to a Resident Library in MACRO-11 5-31
5.1.11.1	Resolving Program Section Names in a Shared Region 5-39
5.1.12	Example 5-4: Building a Task That Creates a Dynamic Region 5-40
5.2	CLUSTER LIBRARIES 5-43
5.2.1	Building the Libraries 5-44
5.2.1.1	Summary of Rules for Building the Libraries 5-44
5.2.1.2	Rule 1: All Libraries but the First Require Resident Overlays 5-45
5.2.1.3	Rule 2: User Task Vectors Indirectly Resolve all Interlibrary References 5-46
5.2.1.4	Rule 3: Revectoring Entry Point Symbols Must Not Appear in the "Upstream" .STB File . . . 5-47
5.2.1.5	Rule 4: A Called Library Procedure Must Not Require Parameters on the Stack 5-47
5.2.1.6	Rule 5: All the Libraries Must be PIC or Built for the Same Address 5-47
5.2.1.7	Rule 6: Trap or Asynchronous Entry Into a Library is not Permitted 5-47
5.2.2	Building Your Task 5-48
5.2.3	Examples 5-48
5.2.3.1	F77CLS -- Build the Default Library for the FORTRAN-77 OTS 5-48
5.2.3.2	FDVRES -- Build an FMS-11/RSX V1.0 Shareable Library 5-49
5.2.3.3	FDVRESBLD.ODL -- Overlay Description for FMS-11/RSX V1.0 Cluster Library 5-51
5.2.3.4	FCSRES Library Build 5-51
5.2.3.5	F77TST.CMD -- File to Build the FMS-11/RSX V1.0 FORDEM Test Task 5-51
5.2.4	Overlay Run-Time Support Requirements 5-51
5.3	VIRTUAL PROGRAM SECTIONS 5-53
5.3.1	FORTRAN Run-Time Support for Virtual Program Sections 5-56
5.3.2	Example 5-5: Building a Program that Uses a Virtual Program Section 5-58
CHAPTER 6 PRIVILEGED TASKS	
6.1	INTRODUCTION 6-1

6.2	PRIVILEGED AND NONPRIVILEGED TASK DISTINCTION . . .	6-1
6.3	PRIVILEGED TASK HAZARDS	6-1
6.4	SPECIFYING A TASK AS PRIVILEGED	6-2
6.5	PRIVILEGED TASK MAPPING	6-2
6.6	/PR:0 PRIVILEGED TASK	6-4
6.7	/PR:4 PRIVILEGED TASK	6-5
6.8	/PR:5 PRIVILEGED TASK	6-5
6.9	EXAMPLE 6-1: BUILDING A PRIVILEGED TASK TO EXAMINE UNIT CONTROL BLOCKS	6-6

CHAPTER 7	USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)	
7.1	USER TASK DATA SPACE DEFINED	7-1
7.2	I- AND D-SPACE TASK IDENTIFICATION	7-1
7.3	COMPARISON OF CONVENTIONAL TASKS AND I- AND D-SPACE TASKS	7-2
7.4	CONVENTIONAL TASK MAPPING	7-2
7.5	I- AND D-SPACE TASK MAPPING	7-3
7.6	TASK WINDOWS IN I- AND D-SPACE TASKS	7-5
7.7	SPECIFYING DATA SPACE IN YOUR TASK	7-5
7.8	OVERLAID I- AND D-SPACE TASKS	7-5
7.8.1	Autoload Vectors and .STB Files	7-9
7.9	I- AND D-SPACE TASK MEMORY ALLOCATION AND EXAMPLE MAPS	7-9
7.9.1	Virtual Memory Allocation for MAIN.TSK	7-10
7.9.2	Virtual Memory Allocation for MAINID.TSK	7-10
CHAPTER 8	SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)	
8.1	INTRODUCTION	8-1
8.2	MODE-SWITCHING VECTORS	8-1
8.3	COMPLETION ROUTINES	8-2
8.4	RESTRICTIONS ON THE CONTENTS OF SUPERVISOR-MODE LIBRARIES	8-2
8.5	SUPERVISOR-MODE LIBRARY MAPPING	8-3
8.5.1	Supervisor Mode Library Data	8-3
8.5.2	Supervisor Mode Libraries with I- and D-Space Tasks	8-3
8.6	BUILDING AND LINKING TO SUPERVISOR-MODE LIBRARIES	8-7
8.6.1	Relevant TKB Options	8-7
8.6.2	Building the Library	8-7
8.6.3	Building the Referencing Task	8-7
8.6.4	Mode Switching Instruction	8-8
8.7	CSM LIBRARIES	8-8
8.7.1	Building a CSM Library	8-8
8.7.2	Linking to a CSM Library	8-9
8.7.3	Example CSM Library and Linking Task	8-10
8.7.3.1	Building SUPER	8-18
8.7.3.2	Building TSUP	8-18
8.7.3.3	Running TSUP	8-19
8.7.4	The CSM Library Dispatching Process	8-19
8.8	CONVERTING SCAL LIBRARIES TO CSM LIBRARIES	8-20
8.9	USING SUPERVISOR-MODE LIBRARIES AS RESIDENT LIBRARIES	8-20
8.10	MULTIPLE SUPERVISOR-MODE LIBRARIES	8-20
8.11	LINKING A RESIDENT LIBRARY TO A SUPERVISOR-MODE LIBRARY	8-21
8.12	LINKING SUPERVISOR-MODE LIBRARIES	8-21
8.13	WRITING YOUR OWN VECTORS AND COMPLETION ROUTINES	8-21
8.14	OVERLAID SUPERVISOR-MODE LIBRARIES	8-21

CONTENTS

Page

CHAPTER 9	MULTIUSER TASKS	
9.1	INTRODUCTION	9-1
9.1.1	Overlaid Multiuser Task	9-2
9.1.2	Disk Image of a Multiuser Task	9-2
9.1.3	I- and D-Space Multiuser Tasks	9-3
9.2	EXAMPLE 9-1: BUILDING A MULTIUSER TASK	9-5

CHAPTER 10	SWITCHES	
10.1	SWITCHES	10-1
10.1.1	/AC[:n] -- Ancillary Control Processor	10-5
10.1.2	/AL -- Allocate Checkpoint Space	10-6
10.1.3	/CC -- Concatenated Object Modules	10-7
10.1.4	/CM -- Compatibility Mode Overlay Structure	10-8
10.1.5	/CO -- Build a Common Block Shared Region	10-9
10.1.6	/CP -- Checkpointable	10-10
10.1.7	/CR -- Cross-Reference	10-11
10.1.8	/DA -- Debugging Aid	10-14
10.1.9	/DL -- Default Library	10-15
10.1.10	/EA -- Extended Arithmetic Element	10-16
10.1.11	/EL -- Extend Library	10-17
10.1.12	/FP -- Floating Point	10-18
10.1.13	/FU -- Full Search	10-19
10.1.14	/HD -- Header	10-20
10.1.15	/ID -- I- and D-space Task (RSX-11M-PLUS Only)	10-21
10.1.16	/IP -- Task Maps I/O Page	10-22
10.1.17	/LB -- Library File	10-23
10.1.18	/LI -- Build a Library Shared Region	10-25
10.1.19	/MA -- Map Contents of File	10-26
10.1.20	/MM[:n] -- Memory Management	10-27
10.1.21	/MP -- Overlay Description	10-28
10.1.22	/MU -- Multiuser (RSX-11M-PLUS Only)	10-29
10.1.23	/NM -- No Diagnostic Messages	10-30
10.1.24	/PI -- Position Independent	10-31
10.1.25	/PM -- Postmortem Dump	10-32
10.1.26	/PR[:n] -- Privileged	10-33
10.1.27	/RO -- Resident Overlay	10-34
10.1.28	/SE -- Send	10-35
10.1.29	/SG -- Segregate Program Sections	10-36
10.1.30	/SH -- Short Map	10-37
10.1.31	/SL -- Slave	10-44
10.1.32	/SP -- Spool Map Output	10-45
10.1.33	/SQ -- Sequential	10-46
10.1.34	/SS -- Selective Search	10-47
10.1.35	/TR -- Traceable	10-50
10.1.36	/WI -- Wide Listing Format	10-51
10.1.37	/XH -- External Header (RSX-11M-PLUS only)	10-52
10.1.38	/XT[:n] -- Exit on Diagnostic	10-53

CHAPTER 11	OPTIONS	
11.1	OPTIONS	11-1
11.1.1	ABORT -- Abort the Task-Build	11-4
11.1.2	ABSPAT -- Absolute Patch	11-5
11.1.3	ACTFIL -- Number of Active Files	11-6
11.1.4	ASG -- Device Assignment	11-7
11.1.5	CLSTR -- System-Owned Cluster of Resident Libraries or Commons	11-8
11.1.6	CMPRT -- Completion Routine -- RSX-11M-PLUS Only	11-10
11.1.7	COMMON or LIBR -- System-Owned Resident Common or System-Owned Resident Library	11-11

11.1.8	DSPPAT -- Absolute Patch for D-space	11-13
11.1.9	EXTSCT -- Program Section Extension	11-14
11.1.10	EXTTSK -- Extend Task Memory	11-15
11.1.11	FMTBUF -- Format Buffer Size	11-16
11.1.12	GBLDEF -- Global Symbol Definition	11-17
11.1.13	GBLINC -- Include Global Symbols	11-18
11.1.14	GBLPAT -- Global Relative Patch	11-19
11.1.15	GBLREF -- Global Symbol Reference	11-20
11.1.16	GBLXCL -- Exclude Global Symbols	11-21
11.1.17	LIBR -- System-Owned Library	11-22
11.1.18	MAXBUF -- Maximum Record Buffer Size	11-23
11.1.19	ODTV -- ODT SST Vector	11-24
11.1.20	PAR -- Partition	11-25
11.1.21	PRI -- Priority	11-27
11.1.22	RESCOM or RESLIB -- Resident Common or Resident Library	11-28
11.1.23	RESLIB -- Resident Library	11-30
11.1.24	RESSUP -- Resident Supervisor-Mode Library -- RSX-11M-PLUS only	11-31
11.1.25	ROPAR -- Read-Only Partition -- RSX-11M-PLUS Only	11-33
11.1.26	STACK -- Stack Size	11-34
11.1.27	SUPLIB -- Supervisor-Mode Library -- RSX-11M-PLUS Only	11-35
11.1.28	TASK -- Task Name	11-36
11.1.29	TSKV -- Task SST Vector	11-37
11.1.30	UIC -- User Identification Code	11-38
11.1.31	UNITS -- Logical Unit Usage	11-39
11.1.32	VSECT -- Virtual Program Section	11-40
11.1.33	WNDWS -- Number of Address Windows	11-41

APPENDIX A TASK BUILDER INPUT DATA FORMATS

A.1	DECLARE GLOBAL SYMBOL DIRECTORY RECORD	A-2
A.1.1	Module Name (Type 0)	A-4
A.1.2	Control Section Name (Type 1)	A-5
A.1.3	Internal Symbol Name (Type 2)	A-5
A.1.4	Transfer Address (Type 3)	A-6
A.1.5	Global Symbol Name (Type 4)	A-6
A.1.6	Program Section Name (Type 5)	A-7
A.1.7	Program Version Identification (Type 6)	A-10
A.1.8	Mapped Array Declaration (Type 7)	A-10
A.1.9	Completion Routine Definition (Type 10)	A-11
A.2	END OF GLOBAL SYMBOL DIRECTORY RECORD	A-11
A.3	TEXT INFORMATION RECORD	A-11
A.4	RELOCATION DIRECTORY RECORD	A-12
A.4.1	Internal Relocation (Type 1)	A-14
A.4.2	Global Relocation (Type 2)	A-15
A.4.3	Internal Displaced Relocation (Type 3)	A-15
A.4.4	Global Displaced Relocation (Type 4)	A-16
A.4.5	Global Additive Relocation (Type 5)	A-16
A.4.6	Global Additive Displaced Relocation (Type 6)	A-17
A.4.7	Location Counter Definition (Type 7)	A-17
A.4.8	Location Counter Modification (Type 10)	A-18
A.4.9	Program Limits (Type 11)	A-18
A.4.10	Program Section Relocation (Type 12)	A-19
A.4.11	Program Section Displaced Relocation (Type 14)	A-19
A.4.12	Program Section Additive Relocation (Type 15)	A-20
A.4.13	Program Section Additive Displaced Relocation (Type 16)	A-21
A.4.14	Complex Relocation (Type 17)	A-22
A.4.15	Resident Library Relocation (Type 20)	A-23
A.5	INTERNAL SYMBOL DIRECTORY RECORD	A-24
A.5.1	Overall Record Format	A-24

CONTENTS

Page

A.5.2 TKB Generated Records (Type 1) A-25
A.5.2.1 Start-of-Segment Item Type (1) A-25
A.5.2.2 Task Identification Item Type (2) A-26
A.5.2.3 Autoloadable Library Entry Point Item Type
(3) A-26
A.5.3 Relocatable/Relocated Records (Type 2) A-27
A.5.3.1 Module Name Item Type (1) A-27
A.5.3.2 Global Symbol Item Type (2) A-28
A.5.3.3 PSECT Item Type (3) A-29
A.5.3.4 Line-Number or PC Correlation Item Type (4) A-29
A.5.3.5 Internal Symbol Name Item Type (5) A-30
A.5.4 Literal Records (Type 4) A-30
A.6 END OF MODULE RECORD A-30

APPENDIX B DETAILED TASK IMAGE FILE STRUCTURE

B.1 LABEL BLOCK GROUP B-1
B.2 CHECKPOINT AREA B-9
B.3 HEADER B-10
B.3.1 Low-Memory Context B-10
B.3.2 Logical Unit Table Entry B-14
B.4 TASK IMAGE B-14
B.4.1 Autoload Vectors for Conventional Tasks B-17
B.4.2 Autoload Vectors for I- and D-Space Tasks B-17
B.4.3 Segment Descriptor B-18
B.4.4 Window Descriptor B-20
B.4.5 Region Descriptor B-21

APPENDIX C HOST AND TARGET SYSTEMS

C.1 INTRODUCTION C-1
C.2 EXAMPLE C-1: TRANSFERRING A TASK FROM A HOST TO A
TARGET SYSTEM C-2

APPENDIX D MEMORY DUMPS

D.1 POSTMORTEM DUMPS D-1
D.2 SNAPSHOT DUMP D-5
D.2.1 Format of the SNPBK\$ Macro D-6
D.2.2 Format of the SNAP\$ Macro D-7
D.2.3 Example of a Snapshot Dump D-8

APPENDIX E RESERVED SYMBOLS

APPENDIX F IMPROVING TASK BUILDER PERFORMANCE

F.1 EVALUATING AND IMPROVING TASK BUILDER THROUGHPUT . F-1
F.1.1 Table Storage F-2
F.1.2 Input File Processing F-6
F.1.3 Summary F-6
F.2 MODIFYING COMMAND SWITCH DEFAULTS F-7
F.3 THE SLOW TASK BUILDER F-11

APPENDIX G THE FAST TASK BUILDER

CONTENTS

Page

APPENDIX H ERROR MESSAGES

GLOSSARY

INDEX

EXAMPLES

EXAMPLE 3-1	Map File for OVR.TSK	3-44
3-2	Map File for RESOVR.TSK	3-47
4-1	Cross-Reference Listing of Overlaid Task	4-13
5-1	Part 1 Common Area Source File in MACRO-11	5-18
5-1	Part 2 Task Builder Map for MACCOM.TSK	5-19
5-1	Part 3 MACRO-11 Source Listing for MCOM1	5-21
5-1	Part 4 MACRO-11 Source Listing for MCOM2	5-22
5-1	Part 5 Task Builder Map for MCOM1.TSK	5-24
5-2	Part 1 Assembly Listing for TTCOM	5-27
5-2	Part 2 Task Builder Map for TTCOM	5-28
5-2	Part 3 Assembly Listing for TEST	5-29
5-2	Part 4 Memory Allocation Map for TEST	5-31
5-3	Part 1 Source Listing for Resident Library LIB.MAC	5-32
5-3	Part 2 Task Builder Map for LIB.TSK	5-34
5-3	Part 3 Source Listing for MAIN.MAC	5-35
5-3	Part 4 Task Builder Map for MAIN.TSK	5-37
5-4	Part 1 Source Listing for DYNAMIC.MAC	5-41
5-4	Part 2 Task Builder Map for DYNAMIC.TSK	5-43
5-5	Part 1 Source Listing for VSECT.FTN	5-59
5-5	Part 2 Task Builder Map for VSECT.TSK	5-61
6-1	Part 1 Source Code for PRIVEX	6-7
6-1	Part 2 Task Builder Map for PRIVEX	6-10
7-1	Map of Overlaid Task MAIN.TSK	7-12
7-2	Map of Overlaid I- and D-Space Task MAINID.TSK	7-14
8-1	Code for SUPER.MAC	8-10
8-2	Memory Allocation Map for SUPER	8-11
8-3	Completion Routine, SCMPCS, from SYSLIB.OLD	8-12
8-4	Code for TSUP.MAC	8-14
8-5	Memory Allocation Map for TSUP	8-17
9-1	Part 1 Source Listing for ROTASK.MAC	9-6
9-1	Part 2 Task Builder Map for ROTASK.TSK	9-8
10-1	Cross-Reference Listing for OVR.TSK	10-12
10-2	Memory Allocation File (Map) Example	10-37
C-1	Part 1 Task Builder Map for LIB.TSK	C-3
C-1	Part 2 Task Builder Map for MAIN.TSK	C-4
D-1	Sample Postmortem Dump (Truncated)	D-2
D-2	Sample Program That Calls for Snapshot Dumps	D-9
D-3	Sample Snapshot Dump (in Word Octal and Radix-50)	D-10
D-4	Sample Snapshot Dump (in Byte Octal and ASCII)	D-11

FIGURES

FIGURE 2-1	Relocatable Object Modules	2-2
2-2	Modules Linked for Mapped and Unmapped Systems	2-3
2-3	Allocation of Task Memory	2-6
2-4	Disk Image of the Task	2-9
2-5	Memory Image	2-9

CONTENTS

	Page	
2-6	Simple 2-Segment, Disk-Resident Overlay Calling Sequence	2-11
2-7	Simple 2-Segment, Memory-Resident Overlay Calling Sequence	2-12
2-8	Virtual and Logical Address Space Coincidence in an Unmapped System	2-15
2-9	Memory Layout for Unmapped System	2-16
2-10	Task Relocation in a Mapped System	2-17
2-11	Memory Management Unit's Division of Virtual Address Space	2-18
2-12	Mapping for 4K-Word and 6K-Word Tasks	2-19
2-13	Window Block 0	2-21
2-14	Virtual to Logical Address Space Translation	2-23
2-15	Mapping for a Conventional User Task and a System Containing a Supervisor-Mode Library in an RSX-11M-PLUS System	2-26
2-16	Mapping for a Conventional User Task Using a Supervisor-Mode Library in an RSX-11M-PLUS System	2-27
2-17	Simplified APR Mapping for an I- and D-space Task	2-29
3-1	TK1 Built As a Single-Segment Task	3-4
3-2	TK1 Built As a Multisegment Task	3-5
3-3	TK1 Built with Additional Overlay Defined	3-7
3-4	TK2 Built As a Single-Segment Task	3-8
3-5	TK2 Built As a Memory-Resident Overlay	3-9
3-6A	Relationship Between Virtual Address Space and Physical Memory -- Time 1	3-11
3-6B	Relationship Between Virtual Address Space and Physical Memory -- Time 2	3-12
3-7A	Relationship Between Virtual Address Space and Physical Memory -- Time 3	3-13
3-7B	Relationship Between Virtual Address Space and Physical Memory -- Time 4	3-14
3-8	Overlay Tree for TK1	3-16
3-9	Resolution of Global Symbols in a Multisegment Task	3-17
3-10	Resolution of Program Sections for TK1	3-19
3-11	Typical Overlay Root Segment Structure	3-21
3-12	Typical Overlaid I- and D-Space Task with Up-Tree Segment	3-22
3-13	Tree and Virtual Address Space Diagram	3-24
3-14	Overlay Tree for Modified TK1	3-31
3-15	Virtual Address Space and Physical Memory for Modified TK1	3-33
3-16	Overlay Co-Tree for Modified TK1	3-34
3-17	Virtual Address Space and Physical Memory for TK1 As a Co-Tree	3-35
3-18	Virtual Address Space Allocation Diagram	3-36
3-19	Virtual Address Space Allocation for a Main Tree and Its Co-Tree	3-40
3-20	Overlay Tree of Virtual Address Space for OVR.TSK	3-43
3-21	Allocation of Virtual Address Space for OVR.TSK	3-46
3-22	Allocation of Virtual Address Space for RESOVR.TSK	3-48
4-1	Details of Segment C of TK1	4-2
4-2	Path-Loading Example	4-4
4-3	Autoload Vector Format for Conventional Tasks	4-4
4-4	Autoload Vector Format for I- and D-space Tasks	4-5
4-5	Example Autoload Code Sequence for a Conventional Task	4-5
4-6	Autoload Overlay Tree Example	4-12
5-1	Typical Resident Common	5-2
5-2	Typical Resident Library	5-3
5-3	Interaction of the /LI, /CO, and /PI Switches	5-4
5-4	Specifying APRs for a Position-Independent Shared Region	5-6

CONTENTS

Page

5-5	Mapping for an Absolute Shared Region	5-8
5-6	Windows for Shared Region and Referencing Task .	5-15
5-7	Allocation Diagram for MACCOM.TSK	5-20
5-8	Assigning Symbolic References within a Common .	5-23
5-9	Allocation of Virtual Address Space for MAIN.TSK	5-38
5-10	Example Library and Task Structure	5-44
5-11	Example of an Unbalanced Tree with Null Segment	5-45
5-12	Example of an Overlay Cluster Library Structure	5-45
5-13	Example of a Vectored Call Between Libraries . .	5-46
5-14	VSECT Option Usage	5-55
6-1	Privileged Task Mapping	6-3
6-2	Mapping for /PR:4 and /PR:5	6-4
6-3	Allocation of Virtual Address Space for PRIVEX .	6-11
7-1	Conventional Task Linked to a Region in an I- and D-space System	7-3
7-2	I- and D-space Task Mapping in an I- and D-space System	7-4
7-3	Simplified Disk Image of a Non-Overlaid I- and D-Space Task	7-6
7-4	Overlaid I- and D-Space Task Virtual Address Space	7-7
7-5	Example Overlay Tree for Overlaid I- and D-Space Task IAND	7-7
7-6	Simplified Disk Image of Overlaid I- and D-Space Task IAND	7-8
7-7	Memory Allocation Diagram for MAIN.TSK	7-10
7-8	Memory Allocation Diagram for MAINID.TSK I-Space	7-11
7-9	Memory Allocation Diagram for MAINID.TSK D-Space	7-11
8-1	Mapping of a 24K Conventional User Task That Links to a 16K Supervisor-Mode Library	8-4
8-2	Mapping of a 20K Conventional User Task that Links to a 12K Supervisor-Mode Library Containing 4K of Data	8-5
8-3	Mapping of a 40K I- and D-Space Task That Links to an 8K Supervisor-Mode Library	8-6
8-4	Overlay Configuration Allowed for Supervisor-Mode Libraries	8-22
9-1	Allocation of Program Sections in a Multiuser Task	9-2
9-2	Windows for a Multiuser Task	9-3
9-3	Example Allocation of Program Sections in an I- and D-Space Multiuser Task	9-4
9-4	Windows for an I- and D-Space Multiuser Task . . .	9-5
A-1	General Object Module Format	A-3
A-2	Global Symbol Directory Record Format	A-4
A-3	Module Name Entry Format	A-4
A-4	Control Section Name Entry Format	A-5
A-5	Internal Symbol Name Entry Format	A-5
A-6	Transfer Address Entry Format	A-6
A-7	Global Symbol Name Entry Format	A-6
A-8	Program Section Name Entry Format	A-8
A-9	Program Version Identification Entry Format . . .	A-10
A-10	Mapped Array Declaration Entry Format	A-10
A-11	Completion Routine Entry Format	A-11
A-12	End of Global Symbol Directory Record Format . .	A-11
A-13	Text Information Record Format	A-12
A-14	Relocation Directory Record Format	A-14
A-15	Internal Relocation Entry Format	A-15
A-16	Global Relocation Entry Format	A-15
A-17	Internal Displaced Relocaton Entry Format . . .	A-16
A-18	Global Displaced Relocation Entry Format	A-16
A-19	Global Additive Relocation Entry Format	A-17
A-20	Global Additive Displaced Relocation Entry Format	A-17
A-21	Location Counter Definition Entry Format	A-18
A-22	Location Counter Modification Entry Format . . .	A-18
A-23	Program Limits Entry Format	A-19
A-24	Program Section Relocation Entry Format	A-19

CONTENTS

Page

A-25	Program Section Displaced Relocation Entry Format	A-20
A-26	Program Section Additive Relocation Entry Format	A-21
A-27	Program Section Additive Displaced Relocation Entry Format	A-21
A-28	Complex Relocation Entry Format	A-23
A-29	Resident Library Relocation Entry Format	A-23
A-30	General Format of All ISD Records	A-25
A-31	General Format of a TKB Generated Record	A-25
A-32	Format of TKB Generated Start-of-Segment Item (1)	A-26
A-33	Format of TKB Generated Task Identification Item (2)	A-26
A-34	Format of an Autoloadable Library Entry Point Item (3)	A-27
A-35	Format of a Module Name Item Type (1)	A-28
A-36	Format of a Global Symbol Item Type (2)	A-28
A-37	Format of a PSECT Item Type (3)	A-29
A-38	Format of a Line-Number or PC Correlation Item Type (4)	A-30
A-39	Format of an Internal Symbol Name Item Type (5)	A-31
A-40	Format of a Literal Record Type	A-31
A-41	End-of-Module Record Format	A-32
B-1	Image on Disk of Non-Overlaid Conventional Task	B-2
B-2	Image on Disk of Conventional Non-Overlaid Task Linked to Overlaid Library	B-2
B-3	Image on Disk of Conventional Overlaid Task	B-3
B-4	Image on Disk of Overlaid I- and D-Space Task	B-4
B-5	Label Block 0 -- Task and Resident Library Data	B-7
B-6	Label Blocks 1 and 2 -- Table of LUN Assignments	B-9
B-7	Label Block 3 -- Segment Load List	B-9
B-8	Task Header, Fixed Part	B-11
B-9	Task Header, Variable Part	B-12
B-10	Vector Extension Area Format	B-13
B-11	Logical Unit Table Entry	B-14
B-12	Task-Resident Overlay Data Base for a Conventional Overlaid Task	B-15
B-13	Task-Resident Overlay Data Base for an I- and D-Space Overlaid Task	B-16
B-14	Autoload Vector Entry for Conventional Tasks	B-17
B-15	Autoload Vector Entry for I- and D-Space Tasks	B-18
B-16	Segment Descriptor	B-19
B-17	Window Descriptor	B-21
B-18	Region Descriptor	B-21
D-1	Snapshot Dump Control Block Format	D-6

TABLES

TABLE	2-1	Program Section Attributes	2-4
	2-2	Program Sections for Modules IN1, IN2, and IN3	2-6
	2-3	Individual Program Section Allocations	2-6
	2-4	Resolution of Global Symbols for IN1, IN2, and IN3	2-7
	4-1	Comparison of Overlay Run-Time Module Sizes	4-16
	5-1	Comparison of Overlay Run-Time Module Sizes	5-52
	7-1	Mapping Comparison Summary	7-2
	10-1	Task Builder Switches	10-2
	10-2	Files for SEL.TSK	10-47
	11-1	Task Builder Options	11-2
	A-1	Symbol Declaration Flag Byte -- Bit Assignments	A-7
	A-2	Program Section Name Flag Byte -- Bit Assignments	A-8
	A-3	Relocation Directory Command Byte -- Bit Assignments	A-13
	B-1	Task and Resident Library Data	B-4
	B-2	Resident Library/Common Name Block Data	B-8

CONTENTS

Page

F-1	Task File Switch Defaults	F-8
F-2	Map File Switch Defaults	F-10
F-3	Symbol Table File Switch Defaults	F-10
F-4	Input File Switch Defaults	F-11

PREFACE

MANUAL OBJECTIVES

This manual describes the concepts and capabilities of the RSX-11M/M-PLUS Task Builder.

Working examples are used throughout this manual to introduce and describe features of the Task Builder. Because RSX-11M systems support a large number of programming languages, it is not practical to illustrate the Task Builder features in all of the languages supported. Instead, most of the examples in the main text of this manual are written in MACRO-11.

INTENDED AUDIENCE

Before reading this manual, you should be familiar with the fundamental concepts of your operating system (RSX-11M or RSX-11M-PLUS) and with the operating procedures described in the RSX-11M/M-PLUS MCR Operations Manual. In addition, you should be familiar with the programming concepts described in the RSX-11M/M-PLUS Guide to Program Development.

STRUCTURE OF THIS DOCUMENT

This manual has 11 chapters. Their contents are summarized as follows:

- Chapter 1 describes the Task Builder command sequences that you use to interact with the Task Builder.
- Chapter 2 describes the basic Task Builder functions, including the Task Builder's allocation of virtual address space and the resolution of global symbols. It also contains an introduction to supervisor-mode libraries, privileged tasks, and multiuser tasks.
- Chapter 3 describes the Task Builder's overlay capability and the language you use to define an overlay structure.
- Chapter 4 describes the two methods available to you to load overlay segments.
- Chapter 5 describes some typical Task Builder features, including tasks that access shared regions and device commons, tasks that create dynamic regions, and virtual program sections.

PREFACE

- Chapter 6 defines privileged tasks, describes their mapping, and shows how to build a privileged task to examine unit control blocks.
- Chapter 7 describes user-mode I- and D-space, the mapping of these spaces, and the advantages of using I- and D-space in user mode.
- Chapter 8 describes supervisor-mode libraries. The chapter defines and shows how to build and use supervisor-mode libraries.
- Chapter 9 describes and shows how to build multiuser tasks.
- Chapter 10 lists and describes the Task Builder switches. The switches are listed in alphabetical order.
- Chapter 11 lists and describes the Task Builder options. The options are listed in alphabetical order.

This manual also contains eight appendices. Their contents are summarized as follows:

- Appendix A contains a detailed description of the Task Builder input data structures.
- Appendix B contains a detailed description of the task image file structure.
- Appendix C describes the considerations for building a task on one system to run on a system with a different hardware configuration.
- Appendix D describes two memory dumps: postmortem and snapshot.
- Appendix E contains a list of the symbols and program section names reserved for Task Builder use.
- Appendix F contains information on improving Task Builder performance.
- Appendix G describes the fast Task Builder.
- Appendix H contains the Task Builder error messages.

A Task Builder glossary follows the appendices.

ASSOCIATED DOCUMENTS

Other manuals closely allied with this document are described in the Information Directory and Master Index for your operating system. This directory defines the intended audience of each manual in the documentation set and provides a brief synopsis of each manual's contents.

PREFACE

CONVENTIONS USED IN THIS DOCUMENT

In this manual, horizontal ellipses (...) indicate that additional, optional arguments in a statement format have been omitted. For example:

```
input-spec,...
```

means that one or more input-spec items, separated by commas, can be specified.

Vertical ellipses mean that lines in an example, command lines, or lines in a Task Builder map file that are not pertinent to an example have been omitted. For example:

```
TKB>input-line  
.  
.  
.
```

means that one or more of the indicated TKB items have been omitted.

The words "Task Builder" in this manual have been abbreviated to the acronym TKB.

Unless otherwise stated, references to tasks, their mapping, and their structure imply a nonprivileged task in an RSX-11M mapped system.

In the examples of Task Builder command sequences, the portion of the command sequence that you type is printed in red. The Task Builder's responses and prompts are printed in black.

Shading in the manual has the following meanings:

pink Indicates that the text describes features appearing only on RSX-11M operating systems.

gray Indicates that the text describes features appearing only on RSX-11M-PLUS operating systems.

SUMMARY OF TECHNICAL CHANGES

This manual contains the changes for RSX-11M Version 4.1 and RSX-11M-PLUS Version 2.1. This manual has been extensively revised. A study of the Table of Contents and this Summary of Technical Changes is recommended before you look for information in the manual.

GENERAL CHANGES

Editorial changes were made throughout the manual to correct typographical errors.

Small technical changes were made throughout the manual as a result of ongoing development, SPR responses, and readers' comments.

The major technical changes to the manual are listed below.

TECHNICAL CHANGES

NEW OPTIONS

DSPPAT -- Allows object-level patching of a conventional task or the D-space part of an I- and D-space task.

CHANGED OPTIONS

ABSPAT -- Allows object-level patching of a conventional task or the I-space part of an I- and D-space task.

COMMON -- The COMMON option causes the common to be mapped with D-space APRs. Therefore, for I- and D-space tasks, the common can contain data only.

EXTTSK -- Extends the D-space portion of an I- and D-space task. Because libraries are mapped with both I-space APRs and D-space APRs, extending the D-space of I- and D-space tasks may cause unmapping of the library's D-space APRs, which causes the library to be mapped in I-space only.

LIBR -- The LIBR option causes the library to be mapped with both I-space and D-space APRs when linked to an I- and D-space task.

RESCOM -- The RESCOM option causes the common to be mapped with D-space APRs. Therefore, for I- and D-space tasks, the common can contain data only.

RESLIB -- The RESLIB option causes the library to be mapped with both I-space and D-space APRs when linked to an I- and D-space task.

SUMMARY OF TECHNICAL CHANGES

NEW ERROR MESSAGES

Module module-name contains incompatible autoload vectors

CHANGED ERROR MESSAGES

Lookup failure resident library file
changed to
Lookup failure resident library file - filename.ext

MISCELLANEOUS TECHNICAL CHANGES

Autoload vectors for conventional tasks have changed. The call to \$AUTO is now made indirectly through .NAUTO in the overlay impure area.

I- and D-space tasks may be overlaid by using either autoload or manual load.

Autoload vectors for I- and D-space tasks have a format different from those of conventional tasks. The autoload vectors for I- and D-space tasks contain an I-space part located in the task's I-space and a D-space part located in the task's D-space.

Memory allocation diagrams may be used as an aid to create .ODL files.

Overlay Run-time System routines have changed size from the previous release.

MACRO-11 and FORTRAN manual load calling sequences for overlays in I- and D-space tasks may not use asynchronous loading.

For versions of TKB that support I- and D-space tasks and that were used to build libraries, TKB allocates autoload vectors in the root of the task only for those autoloadable entry points in the library referenced by the task.

I- and D-space tasks may link to commons, conventional libraries, and supervisor-mode libraries.

Loading I- and D-space tasks into memory requires two disk accesses. Overlaid I- and D-space tasks may require, in addition, two disk accesses for loading each segment if the segment contains both I-space and D-space.

Segment descriptors for I- and D-space tasks contain an extension for the D-space part.

Only one level of overlay is allowed in supervisor-mode libraries.

I- and D-space multiuser tasks are allowed. TKB uses four window blocks to map these tasks.

Internal Symbol Directory Records, along with their formats, are described in Appendix A. They consist of:

- Type 1 records, generated by TKB and output to the .STB file
- Type 2 records, generated by language processors

SUMMARY OF TECHNICAL CHANGES

- Type 3 records, created from type 2 records and output to the .STB file
- Type 4 records, written to the .STB file without modification

A new bit called LD\$TYP distinguishes between a library or common. See offset R\$LFLG in the resident library name block data in Appendix B.

The first library in a cluster may be overlaid and contain a non-null root.

New Task Builder reserved symbols have been added to Appendix E.

The Fast Task Builder supports the /EA switch and the TASK= option.

The map format for an I- and D-space task shows both I- and D-space contributions to a segment and the disk blocks that contain data sections.

Other minor technical and editorial changes have been made also.

CHAPTER 1
INTRODUCTION AND COMMAND SPECIFICATIONS

The basic steps in developing a program are as follows:

1. You write one or more routines in an RSX-11M/M-PLUS supported source language and enter each routine as an ASCII text file, through an editor.
2. You submit each text file to the appropriate language translator (an assembler or compiler), which converts it to a relocatable object module.
3. You specify the object modules as input to the Task Builder (TKB), which combines the object modules into a single task image output file.
4. You install and run the task.

If you find errors in the task when you run it, you make corrections to the text file using the editor, and then repeat steps 2 through 4.

The Task Builder's main function is to convert relocatable object modules (.OBJ files) into a single task image (.TSK file) that you can install and run on a RSX-11M or RSX-11M-PLUS system. The task is the fundamental executable unit in both systems.

If your program consists of a single object module, using the Task Builder (TKB) is appropriately simple. You specify as input only the name of the file containing the object module produced from the translation of the program, and specify as output the task image file.

Typically, however, programs consist of more than a single object module. In this case, you name each of the object module files as input. TKB links the object modules, resolves references between them, resolves references to the system library, and produces a single task image ready to be installed and executed.

TKB makes a set of assumptions (defaults) about the task image based on typical usage and storage requirements. You can override these assumptions by including switches and options in the task-building terminal sequence. Thus, you can build a task that is tailored to its own input/output and storage requirements.

TKB also produces (upon request) a memory allocation (map) file that contains information describing the allocation of address space, the modules that make up the task image, and the value of all global symbols. In addition, you can request that a list of global symbols, accompanied by the name of each referencing module, be appended to the file (global cross reference).

INTRODUCTION AND COMMAND SPECIFICATIONS

Note that the examples in this manual use MCR as the operating system language. Refer to the RSX-11M-PLUS Command Language Manual and, in particular, to the command in that manual for DIGITAL Command Language equivalence.

The following example shows a simple sequence for building a task:

```
>MAC PROG=PROG
>TKB PROG=PROG
>INS PROG
>RUN PROG
```

The first command (MAC) causes the MACRO-11 assembler to translate the source code of the file PROG.MAC into a relocatable object module in the file PROG.OBJ. The second command (TKB) causes TKB to process the file PROG.OBJ and to produce the task image file PROG.TSK. The third command (INS) causes the INSTALL processor to add the task to the Executive's directory of executable tasks (System Task Directory). The fourth command (RUN) causes the task to execute.

The example just given includes the command

```
>TKB PROG=PROG
```

This command illustrates the simplest use of TKB. It gives the name of a single file as output and the name of a single file as input.

The following sections describe basic Task Builder command forms and sequences.

1.1 TASK COMMAND LINE

The task command line contains the output file specifications, followed by the input file specifications; they are separated by an equal sign (=). You can specify up to three output files and any number of input files.

The task command line has the following form:

```
task-image-file,map-file,symbol-definition-file=input-file,...
```

You must give the output files in a specific order: the first file you name is the image (.TSK) file; the second is the memory allocation (.MAP) file; and the third is the symbol definition (.STB) file. The map file lists information about the size and location of components within the task. The symbol definition file contains the global symbol definitions in the task and their virtual or relocatable addresses in a format suitable for reprocessing by TKB. You specify this file when you are building a resident library or common. (Resident libraries and commons are described in Chapter 3.) TKB combines the input files to create a single task image that can be installed and executed.

1.1.1 Printing the Map File

If you create a map file by specifying one in the TKB command line, there are a number of ways that you can print the file. The following examples show you ways that you may print the map file.

1. With the following two command lines, you can create a map file and then print it later. The TKB command line tells TKB to create a task file, a map file without printing it (by use

INTRODUCTION AND COMMAND SPECIFICATIONS

of the switch /-SP), and a symbol definition file. The PRINT command line tells the system to print the map file.

```
>TKB INV.TSK,INV.MAP/-SP,INV.STB=INV.OBJ
>PRINT INV.MAP
```

2. With the next command line, you can print the map file directly as it is created. In this case, TKB tells the system to print the file by use of the switch /SP. However, the system task PRT... or ...PRT must be installed for this method to work.

```
>TKB INV.TSK,INV.MAP/SP,INV.STB=INV.OBJ
```

3. With the next command line, you can print the map file on a line printer that you specify. It is best to use this command line on an RSX-11M-PLUS system because that system uses transparent spooling. Using this command line on an RSX-11M system may cause the printer to be unavailable to other tasks. See your system manager for specific details about using the following command line.

```
>TKB INV.TSK,LPn:,SY:INV.STB=INV.OBJ
```

1.1.2 Omitting Specific Output Files

You can omit any output file by replacing the file specification with the delimiting comma that would normally follow it. The following commands illustrate the ways in which TKB interprets the output file names.

Command	Output Files
>TKB IMG1,IMG1,IMG1=IN1	The task image file is IMG1.TSK, the memory allocation (map) file is IMG1.MAP, and the symbol definition file is IMG1.STB.
>TKB IMG1=IN1	The task image file is IMG1.TSK.
>TKB ,IMG1=IN1	The map file is IMG1.MAP.
>TKB ,,IMG1=IN1	The symbol definition file is IMG1.STB.
>TKB IMG1,,IMG1=IN1	The task image file is IMG1.TSK and the symbol definition file is IMG1.STB.
>TKB =IN1	This is a diagnostic run with no output files.

1.2 MULTILINE INPUT

Although you can specify a maximum of three output files, you can specify any number of input files. When you specify several input files, a more flexible format is sometimes necessary -- one that consists of several lines. This multiline format is also necessary when you want to include options in your command sequence (see Section 1.3).

INTRODUCTION AND COMMAND SPECIFICATIONS

If you type TKB, the Monitor Console Routine (MCR) activates the Task Builder. TKB then prompts for input until it receives a line consisting only of the terminating slash characters (//). For example:

```
>TKB
TKB>IMG1,IMG1=IN1
TKB>IN2,IN3
TKB>//
```

This sequence produces the same result as the single line command

```
>TKB IMG1,IMG1=IN1,IN2,IN3
```

Both command sequences produce the task image file IMG1.TSK and the map file IMG1.MAP from the input files IN1.OBJ, IN2.OBJ, and IN3.OBJ.

You must specify the output file specifications and the equal sign (=) on the first command line. You can begin or continue input file specifications on subsequent lines.

When you type the terminating slash characters (//), TKB stops accepting input, builds the task, and returns control to MCR.

1.3 OPTIONS

You use options to specify the characteristics of the task you are building. To include options in a task, you must use the multiline format. If you type a single slash (/) following the input file specification, TKB requests option information by displaying ENTER OPTIONS: and prompting for input. For example:

```
>TKB
TKB>IMG1,IMG1=IN1
TKB>IN2,IN3
TKB/
Enter Options:
TKB>PRI=100
TKB>COMMON=JRNAL:RO
TKB>//
```

In this sequence there are two options: PRI=100 and COMMON=JRNAL:RO. The two slashes end option input, initiate the task build, and return control to MCR upon completion.

NOTE

When you are building an overlaid task, there are exceptions to the use of the single slash (/). Overlaid tasks are described in Chapter 4.

The RSX-11M/M-PLUS Task Builder provides numerous options, which are described in Chapter 11. The general form of an option is a keyword followed by an equal sign (=) and an argument list. The arguments in the list are separated from one another by a colon (:). In the example above, the first option consists of the keyword PRI and a single argument indicating that the task is to be assigned the priority 100. The second option consists of the keyword COMMON and an argument list, JRNAL:RO, indicating that the task accesses a resident

INTRODUCTION AND COMMAND SPECIFICATIONS

common region named JRNAL and that the access is read-only. You can specify more than one option on a line by using an exclamation point (!) to separate the options. For example, the command

```
TKB> PRI=100!COMMON=JRNAL:RO
```

is equivalent to the two lines:

```
TKB> PRI=100
TKB> COMMON=JRNAL:RO
```

Some options accept more than one set of argument lists. You use a comma (,) to separate the argument lists. For example, in the command

```
TKB> COMMON=JRNAL:RO,RFIL:RW
```

the first argument list indicates that the task has requested read-only access to the resident common JRNAL. The second argument list indicates that the task has requested read/write access to the resident common RFIL.

The following three sequences are equivalent:

```
TKB> COMMON=JRNAL:RO,RFIL:RW

TKB> COMMON=JRNAL:RO!COMMON=RFIL:RW

TKB> COMMON=JRNAL:RO
TKB> COMMON=RFIL:RW
```

1.4 MULTIPLE TASK SPECIFICATIONS

If you intend to build more than one task, you can use the single slash (/) following option input. This directs TKB to stop accepting input, build the task, and request information for the next task build. For example:

```
>TKB
TKB> IMG1=IN1
TKB> IN2,IN3
TKB> /
Enter Options:
TKB> PRI=100
TKB> COMMON=JRNAL:RO
TKB> /
TKB> IMG2=SUB1
TKB> //
```

TKB accepts the output and input file specifications and the option input; it then stops accepting input upon encountering the single slash (/) during option input. TKB builds IMG1.TSK and then returns to accept more input for building IMG2.TSK.

1.5 INDIRECT COMMAND FILES

You can enter commands to TKB directly from the keyboard, or indirectly through the indirect command file facility. To use the indirect command file facility, you prepare a file that contains the TKB commands you want to be executed. Later, after you invoke TKB, you type an at sign (@) followed by the name of the indirect command file.

INTRODUCTION AND COMMAND SPECIFICATIONS

For example, suppose you create a file called AFIL.CMD containing the following:

```
IMG1,IMG1=IN1
IN2,IN3
/
PRI=100
COMMON=JRNAL:RO
//
```

Later, you can type:

```
>TKB
TKB>@AFIL
TKB>
```

or simply:

```
>TKB @AFIL
```

When TKB encounters the at sign (@), it directs its search for commands to the file named AFIL.CMD. The example above is equivalent to the keyboard sequence

```
>TKB
TKB>IMG1,IMG1=IN1
TKB>IN2,IN3
TKB>/
Enter Options:
TKB>PRI=100
TKB>COMMON=JRNAL:RO
TKB>//
```

When TKB encounters two terminating slash characters (//) in the indirect command file, it terminates indirect command file processing, builds the task, and exits to MCR.

When TKB encounters a single slash (/) in an indirect command file and the slash is the last character in the file, TKB directs its search for commands to the terminal. For example, suppose the file AFIL.CMD in the last example is changed to read:

```
IMG1,IMG1=IN1
IN2,IN3
/
```

Later, you can type:

```
>TKB
TKB>@AFIL
```

In this case, TKB goes to the terminal and prompts:

```
Enter Options:
TKB>
```

From this point, you input options to TKB directly from the keyboard. If you then conclude option input from the keyboard with double slashes (//), TKB suspends command processing, as described above, and exits to MCR following the task build. If you conclude option input with a single slash (/), TKB prompts for new command input following the task build of IMG1.TSK, as follows:

```
TKB>
```


INTRODUCTION AND COMMAND SPECIFICATIONS

Using the single slash (/) following option input in indirect command files is a convenient way to return control to your terminal between successive task builds. For example, suppose you create two indirect command files. The first, AFIL.CMD, contains:

```
IMG1,IMG1=IN1
IN2,IN3
/
PRI=100
COMMON=JRNAL
/
```

The second, AFIL1.CMD, contains:

```
IMG2,IMG2=IN4
IN5,IN6
/
PRI=100
//
```

Then, the terminal sequence to build these two tasks is:

```
>TKB
TKB>@AFIL
TKB>@AFIL1
>
```

NOTE

For interaction with a TKB indirect command file as described above, you must use the multiline format when you specify the indirect command file.

TKB permits two levels of indirection in file references. That is, the indirect command file referenced in a terminal sequence can contain a reference to another indirect command file. For example, if the file BFIL.CMD contains all the standard options that are used by a particular group of users at an installation, you can modify AFIL to include an indirect command file reference to BFIL.CMD as a separate line in the option sequence.

The contents of AFIL.CMD would then be:

```
IMG1,IMG1=IN1
IN2,IN3
/
PRI=100
COMMON=JRNAL:RO
@BFIL
//
```

To build these files, you type:

```
>TKB
TKB> @AFIL
```

Suppose the contents of BFIL.CMD are:

```
STACK=100
UNITS=5:ASG=DT1:5
```

INTRODUCTION AND COMMAND SPECIFICATIONS

Then the terminal equivalent of building these files is:

```
>TKB
TKB>IMG1,IMG1=IN1
TKB>IN2,IN3
TKB>/
Enter Options:
TKB>PRI=100
TKB>COMMON=JRNAL:RO
TKB>STACK=100
TKB>UNITS=5!ASG=DT1:5
TKB>//
```

The indirect command file reference must appear on a separate line. For example, if you modify AFIL.CMD by adding the @BFIL reference on the same line as the COMMON=JRNAL:RO option, the substitution would not take place and TKB would report an error.

1.6 COMMENTS IN LINES

You can include comments at any point in the command sequence, except in lines that contain file specifications. You begin a comment with a semicolon (;) and terminate it with a carriage return. All text between these delimiters is a comment.

For example, in the indirect command file AFIL.CMD, described in Section 1.5, you can add comments to provide more information about the purpose and the status of the task.

```
;
; TASK 33A
;
; DATA FROM GROUP E-46 WEEKLY
;
IMG1,IMG1=
;
; PROCESSING ROUTINES
;
IN1
;
; STATISTICAL TABLES
;
IN2
;
; ADDITIONAL CONTROLS
;
IN3
/
PRI=100
;
COMMON=JRNAL:RO ; RATE TABLES
;
; TASK STILL IN DEVELOPMENT
;
//
```

1.7 FILE SPECIFICATIONS

TKB adheres to the standard RSX-11M/M-PLUS conventions for file specifications. For any file, you can specify the device, the User File Directory (UFD), the file name, the file type, the file version number, and any number of switches.

INTRODUCTION AND COMMAND SPECIFICATIONS

The file specification has the form

```
device:[group,member]filename.type;version/sw1/sw2.../swn
```

When you specify files by name only, TKB applies the default switch settings for device, group, member, type, and version.

For example:

```
>TKB
TKB>IMG1,IMG1=IN1
TKB>IN2,IN3
TKB>//
```

If the current User Identification Code (UIC) of the terminal that TKB is running on is [200,200], the task image file specification of the example is assumed to be:

```
SY0:[200,200]IMG1.TSK;1
```

That is, TKB creates the task image file on the system device (SY0:) under UFD [200,200]. The default type for a task image file is .TSK and, if the name IMG1.TSK is new, the version number is 1. The default settings for all the task image switches also apply. Switch defaults are described in detail in Chapter 6.

For example:

```
>TKB
TKB>[[20,23]]IMG1/CP/DA,IMG1/CR=IN1
TKB>IN2;3,IN3
TKB>//
```

This sequence of commands instructs TKB to create a task image file IMG1.TSK;1 and a memory allocation (map) file IMG1.MAP;1 (actually, it produces IMG1.TSK and IMG1.MAP with versions one higher than the current versions) under UFD [20,23] on the device SY:. The task image is checkpointable and contains the standard debugging aid (ODT). TKB outputs the map to the line printer with a global cross-reference listing appended to it. TKB builds the task from the latest versions of IN1.OBJ and IN3.OBJ, and the specific version of IN2.OBJ. The input files are all found on the system device.

The system device (SY:) is always the default device unless you specify otherwise. If you specify another device on either side of the equal sign, that device becomes the default device for the files on that side of the equal sign. For example:

```
>TKB
TKB>[[20,23]]IMG1,IMG1,IMG1=DB1:IMG1,IN1,IN2
```

This command line produces a task image file, map file, and listing file in UFD [20,23] on device SY:. All the object files are in UFD [20,23] on device DB1. In cases where files are scattered among several devices, the devices must be specified in the command line.

For some files, a device specification is sufficient. In the example above, the map file could be fully specified by the device LP:. The map listing is produced on the line printer, but is not retained as a file.

This example also used switches /CP, /CR, and /DA. The code, syntax, and meaning for each switch are given in Chapter 6.

INTRODUCTION AND COMMAND SPECIFICATIONS

1.8 SUMMARY OF SYNTAX RULES

The syntax rules for issuing commands to TKB are as follows:

- A task-build command can take any one of four forms. The first form is a single line:

```
>TKB task-command-line
```

The second form has additional lines for input file names:

```
>TKB
TKB>task-command-line
TKB>input-line
.
.
TKB>terminating-symbol
```

The third form allows you to specify options:

```
>TKB
TKB>task-command-line
TKB>/
Enter Options:
TKB>option-line
.
.
TKB>terminating-symbol
```

The fourth form has both input lines and option lines:

```
>TKB
TKB>task-command-line
TKB>input-line
.
.
TKB>/
Enter Options:
TKB>option-line
.
.
TKB>terminating-symbol
```

The terminating symbol can be:

```
/ if you intend to build more than one task
```

```
// if you want TKB to return control to MCR
```

- A task command line has one of the three forms:

```
output-file-list=input-file,...
```

```
=input-file,...
```

```
@indirect-command-file
```

The third form is an indirect command file specification, as described in Section 1.5.

INTRODUCTION AND COMMAND SPECIFICATIONS

- An output file list has one of the three forms:

task-image-file,map-file,symbol-definition-file

task-image-file,map-file

task-image-file

The task-image-file is the file specification for the task image file; map-file is the file specification for the memory allocation (map) file; and symbol-definition-file is the file specification for the symbol definition file. Any of the specifications can be omitted, so that, for example, the following form is permitted:

task-image-file,,symbol-definition-file

- An input line has one of two forms:

input-file,...

@indirect-command-file

Both input-file and indirect-command-file are file specifications.

- An option line has one of two forms:

option!...

@indirect-command-file

The indirect-command-file is a file specification.

- An option has the form:

keyword=argument-list,...

The argument-list is:

arg:...

The syntax for each option is given in Chapter 6.

- A file specification conforms to standard RSX-1M/M-PLUS conventions. It has the form:

device:[group,member]filename.type;version/sw1/sw2.../swn

device:

The name of the physical device on which the volume containing the desired file is mounted. The name consists of two ASCII characters followed by an optional 1- or 2-digit octal unit number and a colon; for example, LP: or DT1:.

group

The group number, in the range of 1 through 377(8).

INTRODUCTION AND COMMAND SPECIFICATIONS

member

The member number, in the range 1 through 377(8).

filename

The name of the desired file. The file name can contain up to 9 alphanumeric characters.

type

The 3-character file type identification. Files having the same name but a different function are distinguished from one another by the file type; for example, CALC.TSK and CALC.OBJ.

version

The version number, in octal, of the file. Various versions of the same file are distinguished from one another by this number; for example, CALC.OBJ;1 and CALC.OBJ;2.

All components of a file specification are optional. The combination of the group number and the member number is the User File Directory (UFD) that contains the file name.

CHAPTER 2

TASK BUILDER FUNCTIONS

The process of building a task involves three distinct Task Builder (TKB) functions:

1. Linking object modules
2. Assigning addresses to the task image
3. Building data structures into the task

First, TKB is a linker. It collects and links the relocatable object modules that you specify to it into a single task image, and resolves references to global symbols across the module boundaries.

Second, TKB assigns addresses to the task image. On mapped systems, TKB assigns addresses for a task beginning at 0. The Executive then relocates the addresses at run time. On unmapped systems, TKB assigns addresses for a task beginning at the base address of the partition in which the task is to run. The addresses of tasks that run on unmapped systems are not relocated at run time.

NOTE

Unless otherwise indicated, references to tasks that run on mapped systems assume that the tasks are nonprivileged and residing within system-controlled partitions.

Third, TKB builds data structures into the task image that are required by the INSTALL processor to install the task and by the Executive to run it.

This chapter describes the three TKB functions in detail. It also describes the concepts of mapped and unmapped systems. In addition, this chapter introduces regions, supervisor-mode libraries, overlays, privileged tasks, I- and D-space tasks, and many of the mapping concepts necessary for an understanding of task mapping and Task Builder functions.

2.1 LINKING OBJECT MODULES

TKB links object modules within the context of program sections and resolves references to global symbols across module boundaries.

When the language translators convert symbolic source code within a module to object code, they assign provisional 16-bit addresses to the code. A single assembly or compilation produces a single object

TASK BUILDER FUNCTIONS

module. In its simplest form, each module begins at 0 and extends upward to the highest address in the module. Three object modules produced at separate times might have the address limits shown in Figure 2-1.




Figure 2-1 Relocatable Object Modules

If these modules represent the separate modules of a single program, TKB links them together and modifies the provisional addresses to one of the following:

- For a mapped system, a single sequence of addresses beginning at 0 and extending upward to the sum of the lengths of all the modules (-1 byte)
- For an unmapped system, a single sequence of addresses beginning at a base address assigned at task-build time and extending upward to the sum of the lengths of all the modules (-1 byte)

For example, Figure 2-2 shows the three modules linked for a mapped system and the modules linked for an unmapped system.


2.1.1 Allocating Program Sections

The language translators process source code and TKB links object modules within the context of program sections. A program section is a block of code or data that consists of three elements:

- A name
- A set of attributes
- A length

A program section is the basic unit used by TKB to determine the placement of code and data in a task image. The language translators maintain a separate location counter for each program section in a program. The name of each program section, its attributes, and its length are conveyed to TKB through the object module.

TASK BUILDER FUNCTIONS



ZK-378-81

Figure 2-2 Modules Linked for Mapped and Unmapped Systems

You can create as many program sections within a module as you wish by explicitly declaring them (with the COMMON statement in FORTRAN or the .PSECT directive in MACRO-11, for example) or by allowing the language translator to create them. If you do not explicitly create a program section in your source code, the language translator you are working with will create a "blank" program section within each module translated. This program section will appear on your listings and maps as .BLK.. For more information on explicitly declared program sections, see your language reference manual.

A program section's name is the name by which the language translator and TKB reference it. When processing files, both the language translator and TKB create internal tables that contain program section names, attributes, and lengths. A named program section can be declared more than once. However, all occurrences of that named program section must have identical attributes if the section occurs more than once in the same module or if the section is a global program section. Identically named program sections within the same module and global program sections with differing attributes cause TKB to declare the program section as having multiple attributes, which is an error. However, identically named program sections with differing attributes may appear in different trees of an overlaid task if the program sections have the local (LCL) attribute.

TASK BUILDER FUNCTIONS

Program section attributes define a program section's contents, its placement in a task image, and, in some cases, the allowed mode of access (read/write or read-only).

A program section's length determines how much address space TKB must reserve for it.

When a program consists of more than one module, it is not unusual for program sections of the same name to exist in more than one of the modules. Therefore, as TKB scans the object modules, it collects scattered occurrences of program sections of the same name and combines them into a single area of your task image file. The attributes listed in Table 2-1 control the way TKB collects and places each program section in the task image.

Table 2-1
Program Section Attributes

Attribute	Value	Meaning
access-code	RW	Read/write: data can be read from, and written into, the program section.
	RO	Read-only: data can be read from, but cannot be written into, the program section.
allocation-code	CON	Concatenate: all references to a given program section name are concatenated; the total allocation is the sum of the individual allocations.
	OVR	Overlay: all references to a given program section name overlay each other; the total allocation is the length of the longest individual allocation.
relocation-code	REL	Relocatable: the base address of the program section is relocated relative to the base address of the task.
	ABS	Absolute: the base address of the program section is not relocated; it is always 0.
save	SAV	The program section has the SAVE attribute, and TKB forces the program section into the root.
scope-code	GBL	Global: the program section name is recognized across overlay segment boundaries; TKB allocates storage for the program section from references outside the defining overlay segment.
	LCL	Local: the program section name is recognized only within the defining overlay segment; TKB allocates storage for the program section from references within the defining overlay segment only.

(continued on next page)

TASK BUILDER FUNCTIONS

Table 2-1 (Cont.)
Program Section Attributes

Attribute	Value	Meaning
type-code	D	Data: the program section contains data.
	I	Instruction: the program section contains either instructions, or data and instructions.

2.1.1.1 Access-code and Allocation-code - TKB uses a program section's access-code and allocation-code to determine its placement and size in a task image. If you specify /SG in the command sequence, TKB divides address space into read/write and read-only areas, and places the program sections in the appropriate area according to access-code. However, the default is to order the program sections alphabetically.

TKB uses a program section's allocation-code to determine its starting address and length. If a program section's allocation-code indicates that TKB is to overlay it (OVR), TKB places each allocation to the program section from each module at the same address within the task image. TKB determines the total size of the program section from the length of the longest allocation to it.

If a program section's allocation-code indicates that TKB is to concatenate it (CON), TKB places the allocation from the modules one after the other in the task image, and determines the total allocation from the sum of the lengths of each allocation.

TKB always allocates address space for a program section beginning on a word boundary. If the program section has the D (data) and CON (concatenate) attributes, TKB appends to the last byte of the previous allocation all storage contributed by subsequent modules. It does this regardless of whether that byte is on a word or nonword boundary. For a program section with the I (instruction) and CON attributes, however, TKB allocates address space contributed by subsequent modules beginning with the nearest following word boundary.

For example, suppose three modules, IN1, IN2, and IN3, are to be task built. Table 2-2 lists these modules with the program sections that each contains and their access codes and allocation codes.

In this example, the program section named B, with the attribute CON (concatenate), occurs twice. Thus, the total allocation for B is the sum of the lengths of each occurrence; that is, $100 + 120 = 220$. The program section named A also occurs twice. However, it has the OVR (overlay) attribute; so its total allocation is the largest of the two sizes, or 300. Table 2-3 lists the individual program section allocations.

TASK BUILDER FUNCTIONS

Table 2-2
Program Sections for Modules IN1, IN2, and IN3

File Name	Program Section Name	Access Code	Allocation Code	Size (Octal)
IN1	B	RW	CON	100
	A	RW	OVR	300
	C	RO	CON	150
IN2	A	RW	OVR	250
	B	RW	CON	120
IN3	C	RO	CON	50


Table 2-3
Individual Program Section Allocations

Program Section Name	Total Allocation
B	220
A	300
C	220

TKB then groups the program sections according to their access codes and alphabetizes each group, as shown in Figure 2-3.

NOTE

The example shown in Figure 2-3 represents the Task Builder's allocation of program sections if the /SG or /MU switches are used. For more information, see the description of the /MU, /SQ, and /SG switches in Chapter 10.



ZK-379-81

Figure 2-3 Allocation of Task Memory

TASK BUILDER FUNCTIONS

The save attribute (SAV) is useful in cases where the information in a program section must be kept available to all task segments. The SAV attribute of a program section causes TKB to force the program section into the root of an overlaid task. Therefore, the named common block in the FORTRAN SAVE statement or the named program section in the MACRO-11 .PSECT directive specified with the SAV attribute are in the root of the task.

2.1.1.2 Type-Code and Scope-Code - The scope-code is meaningful only when you define an overlay structure for a task. The scope-code is described in Chapters 3 and 4 within the context of the descriptions of overlays. {The type-code is meaningful in the context of program sections within an I- and D-space task, as described in Chapter 7.}

2.1.2 Resolving Global Symbols

TKB resolves references to global symbols across module boundaries and any references (explicit or implicit) to the system library. When the language translators process a text file, they assume that references to global symbols within the file are defined in other, separately assembled or compiled modules. As TKB links the relocatable object modules, it creates an internal table of the global symbols it encounters within each module. If, after TKB examines and links all the object modules, references remain to symbols that have not been defined, TKB assumes that it will find the definition for the symbols within the default system object module library (LB:[1,1]SYSLIB.OLB). If undefined symbols still remain after SYSLIB is examined, TKB flags the symbols as undefined. If you have not specified an output map in your TKB command sequence, TKB reports the names of the undefined symbols to you on your terminal. If you have specified an output map, TKB outputs to your terminal only the fact that the task contains undefined symbols. The names of the symbols appear on your map listing.

When creating the task image file, TKB resolves global references, as shown in the following example. Table 2-4 lists the three files IN1, IN2, and IN3, showing the program sections within each file, the global symbol definitions within each program section, and the references to global symbols in each program section.

Table 2-4
Resolution of Global Symbols for IN1, IN2, and IN3

File Name	Program Section Name	Global Definition	Global Reference
IN1	B	B1	A
		B2	L1
	A		C1
	C		XXX
IN2	A	A	
	B	B1	B2
IN3	C		B1

TASK BUILDER FUNCTIONS

In processing the first file, IN1, TKB finds definitions for B1 and B2 and references to A, L1, C1, and XXX. Because no definition exists for these references, TKB defers the resolution of these global symbols. In processing the next file, IN2, TKB finds a definition for A, which resolves the previous reference, and a reference to B2, which can be immediately resolved.

When all the object files have been processed, TKB has three unresolved global reference: C1, L1, and XXX. Assume that a search of the system library LB:[1,1]SYSLIB.OLB resolves L1 and XXX, and TKB includes the defining modules in the task's image. Assume also that TKB cannot resolve the global symbol C1. TKB lists it as an undefined global symbol.

The relocatable global symbol B1 is defined twice. TKB lists it as a multiply defined global symbol. TKB uses the first definition of that multiply defined symbol.

Finally, an absolute global symbol (for example, symbol=100) can be defined more than once without being listed as multiply defined, as long as each occurrence of the symbol has the same value.

2.2 THE TASK STRUCTURE

TKB builds the data structures required by other system programs and incorporates them into the task image. The Executive (which is responsible for the allocation of system resources) must have access to the data for all tasks on the system. It must know, for example, a task's size and priority, and it must have information about the way each task expects to use the system. It is the Task Builder's responsibility to allocate space in the task image for the data structures required by the Executive. For example, TKB allocates space for the task header and initializes it.

The disk image file created by TKB contains the linked task and all of the information required by the system programs to install and run it. In its simplest form, the disk image file consists of three physically contiguous parts:


- The label block group
- The task header
- The task memory image

Figure 2-4 illustrates the basic simplified structure of this file.

The label block group contains data produced by TKB and used by INSTALL command processing. It contains information about the task, such as the task's name, the partition in which it runs, its size and priority, and the logical units assigned to it. When you install the task, INSTALL command processing (hereinafter called INSTALL) uses this information to create a Task Control Block (TCB) entry for the task in the System Task Directory (STD) and to initialize the task's header information.

The task's header contains information that the Executive uses when it runs the task. The header also provides a storage area for saving the task's essential data when the task is checkpointed. TKB creates and partially initializes the header; INSTALL initializes the rest of the header.

TASK BUILDER FUNCTIONS




ZK-380-81

Figure 2-4 Disk Image of the Task

The task memory contains the linked modules of the program and, therefore, the code and data. It also contains the task's stack. The stack is an area of task memory that a task can use for temporary storage and subroutine linkage. It can be referenced through general register 6, the stack pointer (SP). The label block group, the task's header, and the task memory are described in detail in Appendix B.

The task's memory image is the part of your task that the system reads into physical memory at run time. The label block group is not required in physical memory. Therefore, in its simplest form, the task's memory image consists of only two parts: the task header and task memory. Figure 2-5 shows the memory image.



ZK-381-81

Figure 2-5 Memory Image

TASK BUILDER FUNCTIONS

2.3 OVERLAYS

This section is an introduction to overlaid tasks. Details about overlaid tasks can be found in Chapters 3 and 4.

Using overlays can save memory space by reducing the size of the executing portion of the task or the physical memory required by the task. Parts of an overlaid task reside on disk, thereby saving memory space.

An overlaid task is a task designed to have discrete parts. The parts of a task designed this way can execute relatively independently of other parts. Parts of an overlaid task reside on disk until they are needed for their required function. The common part of the task, which stays in memory, is the root. The root calls the other parts of the task, which are referred to as segments, from disk into memory.

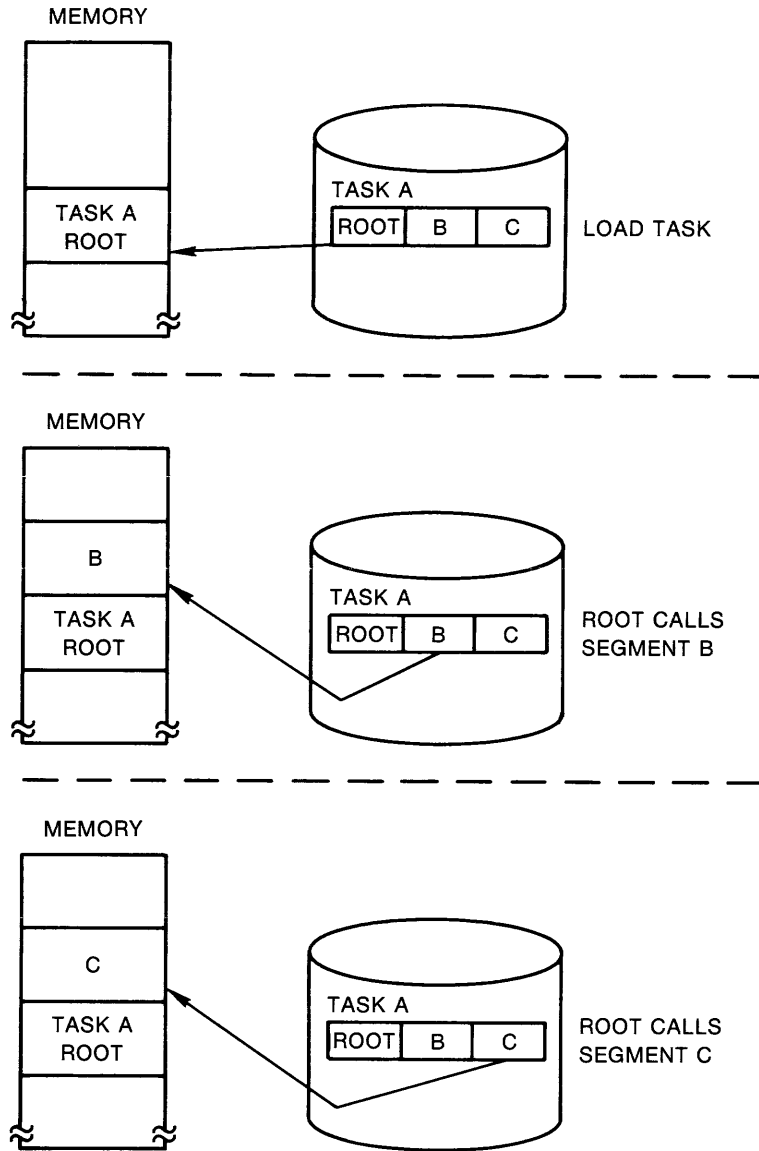
The RSX-11M/M-PLUS systems have two types of overlaid tasks. One type of overlaid task reads in segments from disk over other segments already in memory. A task of this type is called a disk-resident overlaid task. In this task, segments reside on disk until they are needed. The segments in disk-resident overlays that share the same memory address space of the task with other segments must be logically independent of those segments. The independence is necessary because the other segments are on disk and cannot be referenced. For example, Task A, an overlaid task root, can call either of two segments: segment B or segment C. The root of Task A initially calls segment B. Segments B and C occupy the same memory space. Segment B cannot call segment C and segment C cannot call segment B. However, if segment B returns control of the task to the root of task A, the root can then call segment C. Segment C would then be read into memory over segment B. Figure 2-6 illustrates this sequence.

Because segments of a disk-resident overlaid task can occupy the same memory space, a disk-overlaid task can occupy less memory than it would if it were not overlaid. However, more disk I/O transfers (and, therefore, more time) are needed for this type of task.

Another type of overlaid task is the memory-resident overlaid task. In this task, the segments reside on disk until they are needed. At that time, the needed segment is read into a sequentially adjacent area of memory and resides there until the task ends. For example, a memory-resident overlaid Task A has two segments: segment B and segment C. If the root of task A calls segment B, segment B is read into memory adjacent to the root. When the root regains control and then calls segment C, segment C is read into memory adjacent to segment B. Figure 2-7 illustrates this sequence.

Memory-resident overlaid tasks execute faster than disk-resident overlaid tasks. The increase in speed occurs because fewer disk I/O transfers are needed during task execution.

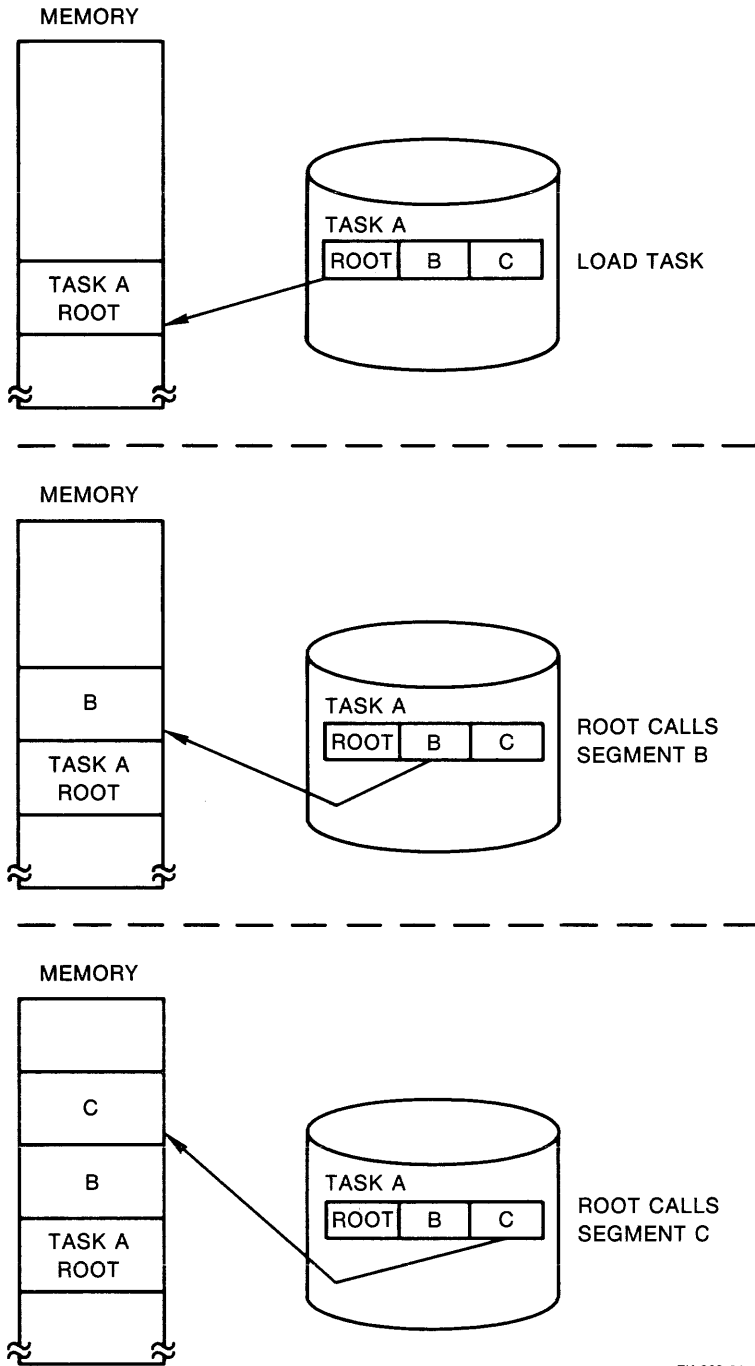
TASK BUILDER FUNCTIONS



ZK-382-81

Figure 2-6 Simple 2-Segment, Disk-Resident Overlay Calling Sequence

TASK BUILDER FUNCTIONS



ZK-383-81

Figure 2-7 Simple 2-Segment, Memory-Resident Overlay Calling Sequence

TASK BUILDER FUNCTIONS

2.4 ADDRESSING CONCEPTS

The primary addressing mechanism of the PDP-11 is the 16-bit computer word. The maximum physical address space that the PDP-11 can reference at any one time is a function of the length of this word. Because of the 16-bit word size, a task can have an address no larger than 177777(octal) (32K words) within the task image for nonprivileged tasks on an unmapped system. In practice, the task size may be limited to a few words less than 32K because of system design.

2.4.1 Physical, Virtual, and Logical Addresses

Physical, virtual, and logical addresses, and virtual and logical address space, are concepts that provide a basis for understanding the functions of task addressing and the use of task windows.

- Physical addresses - A single, physical location in memory is called the physical address.

Memory is divided into parts called bytes. They are numbered according to their position in memory. Therefore, the lowest byte is 0 and the highest byte is whatever the upper limit of memory may be for a particular system; for example, 32K, 64K, and so forth. The assigned number is called the physical address.

A task contains addresses (for example, 0 through 2200). TKB relocates the task's addresses in an unmapped system by a number represented by the base address of the partition in which it is installed. After installation, the task's addresses refer to physical addresses of memory, which always correspond to the same physical memory in an unmapped system.

Therefore, the task addresses have an actual one-to-one relationship to physical memory. The same relationship exists any time the task is in memory. The memory (physical) addresses will not be from 0 through 2200. For example, after the task is installed in the partition, the task's address of 0 may become physical address 17000 because the Task Builder added in the offset, which is equal to the partition base address.

In a mapped system, the task's addresses remain the same but the physical memory addresses may change due to Executive processes (checkpointing, swapping, and so forth.). Therefore, the task addresses do not always correspond to the same physical memory. If the task uses memory management directives, the memory addressing can be changed by the task to include any part of physical memory that it is allowed to access.

- Virtual addresses - A task's virtual addresses are the addresses within the task.

The PDP-11's 16-bit word length (a mapped system) imposes the address range of 32K words on the virtual addresses. Therefore, these task addresses could include addresses 0 through 177777(octal) depending on the length of the task. These task addresses are not the same as the actual addresses of the memory in which the task resides.

TASK BUILDER FUNCTIONS

- Virtual address space - A task's virtual address space is that space encompassed by the range of virtual addresses that the task uses.

With the Create Address Window (CRAW\$) memory management directive, a task can divide its virtual address space into segments called virtual address windows. By using address windows, you can manipulate the mapping of virtual addresses to different areas of physical memory.

- Logical addresses - A task's logical addresses are the actual physical memory addresses that the task can access.
- Logical address space - The task's logical address space is the total amount of physical memory to which the task has access rights.

The physical memory represented by the logical addresses may or may not be continuous. The items in physical memory that logical address space includes are the task itself, and static and dynamic regions.

2.4.2 Unmapped Systems

In an unmapped system, the task's virtual address space and its logical address space coincide exactly, as shown in Figure 2-8.

In an unmapped system, the task's address space is limited to 32K words. All of the machine's physical memory and all of its device registers are accessible to all tasks running on the system. The top 4K words of address space are reserved for the UNIBUS addresses that correspond to the peripheral device registers (the I/O page), and a segment of low memory is occupied by the Executive. Therefore, in an unmapped system, the largest task size is 32K words minus the I/O page and the size of the Executive. Figure 2-9 shows the memory layout for an unmapped system.

Unmapped systems contain only user-controlled partitions. When TKB links the relocatable object modules of a task that is to run on an unmapped system, it requires that you specify the partition in which the task is to run, and the partition's base address and length. TKB sets the base address of the task to the base address of the partition. This means that the task's location in physical memory is bound to the partition and does not change. Because all of physical memory in an unmapped system is directly addressable, and the task's location within memory does not change, the addresses that TKB assigns coincide exactly with the physical addresses of the machine and, therefore, do not need to be relocated at run time.

2.4.3 Mapped Systems

A mapped system is one in which the processor contains a KT-11 memory management unit. The processor handbook for your machine contains a complete description of the memory management unit.

Mapped processors have up to three modes of operation: kernel, supervisor, and user (the PDP-11/34 does not have supervisor mode). The information in this section is relevant to user mode only.

TASK BUILDER FUNCTIONS




Figure 2-8 Virtual and Logical Address Space Coincidence in an Unmapped System

In a mapped system, the relationship between virtual address space and physical address space is different from that of an unmapped system. The primary addressing mechanism for a mapped system is still the 16-bit word, and virtual address space is still 32K words. However, a mapped system has a much greater physical memory capacity and, therefore, physical memory and virtual address space do not coincide.

To address all of physical memory in a mapped system, a machine must have an effective word length of 18 or 22 bits, depending on the model of the machine. When TKB links the relocatable object modules of a task that is to run on a mapped system, it assigns 16-bit addresses to the task image. The memory management unit's function (under control of the Executive) is to convert the task's 16-bit addresses to effective 18- or 22-bit physical addresses. The mechanical job of task relocation is performed by the Executive and the memory management unit at task run time. Figure 2-10 illustrates the relationship between physical memory and virtual address space in a mapped system.

TASK BUILDER FUNCTIONS




Figure 2-9 Memory Layout for Unmapped System

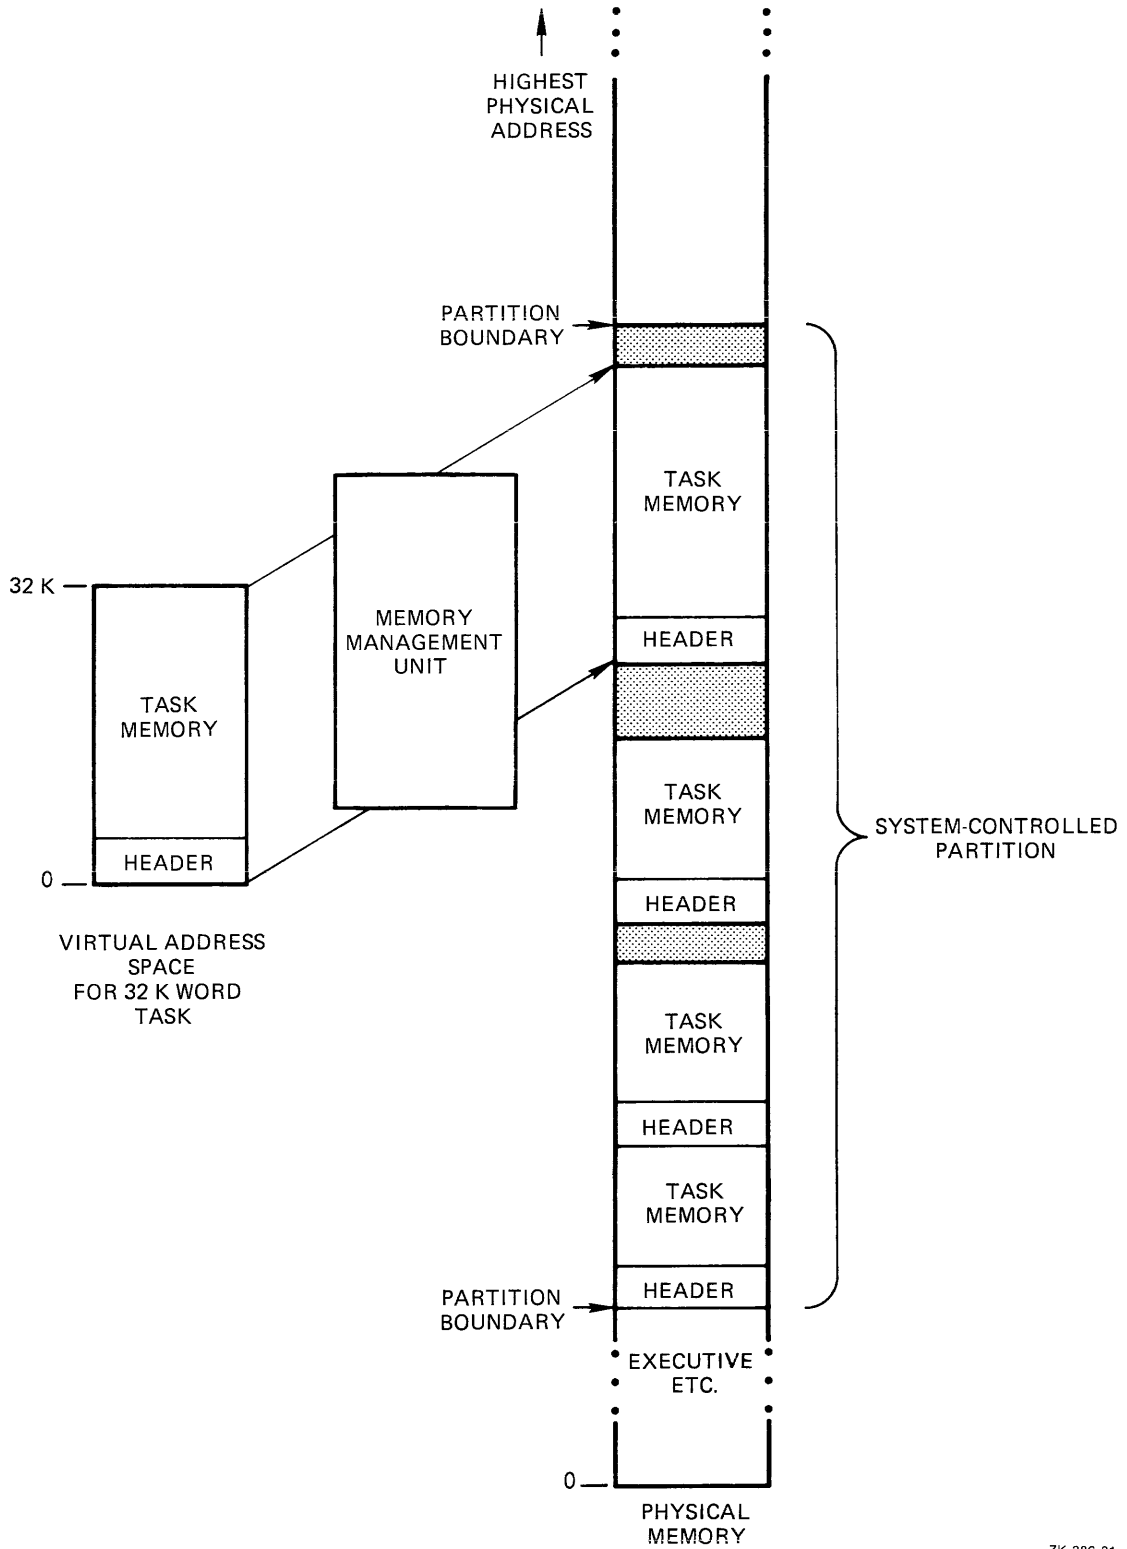
The memory management unit divides a machine's 32K words of virtual address space into eight 4K-word segments or pages. Each page has two registers associated with it:

- A 16-bit Page Description Register (PDR), which contains control and access information about the page with which it is associated
- A 16-bit Page Address Register (PAR), which is an address relocation register

The PDRs and PARs are always used as a pair. Each pair is called an Active Page Register (APR). Figure 2-11 shows how the memory management unit divides the 32K words of virtual address space.

The Executive allocates only as many APRs as are necessary to map a given task into physical memory. Therefore, a 4K-word task requires one APR; a 6K-word task requires two. Figure 2-12 illustrates this mapping.

TASK BUILDER FUNCTIONS



ZK-386-81

Figure 2-10 Task Relocation in a Mapped System

TASK BUILDER FUNCTIONS

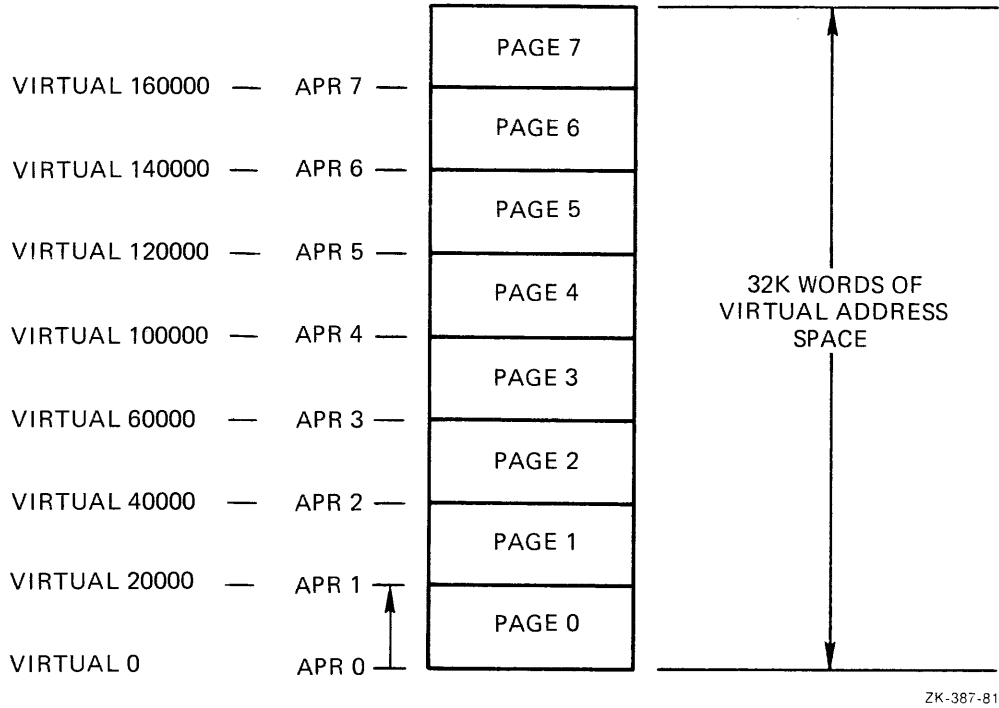


Figure 2-11 Memory Management Unit's Division of Virtual Address Space

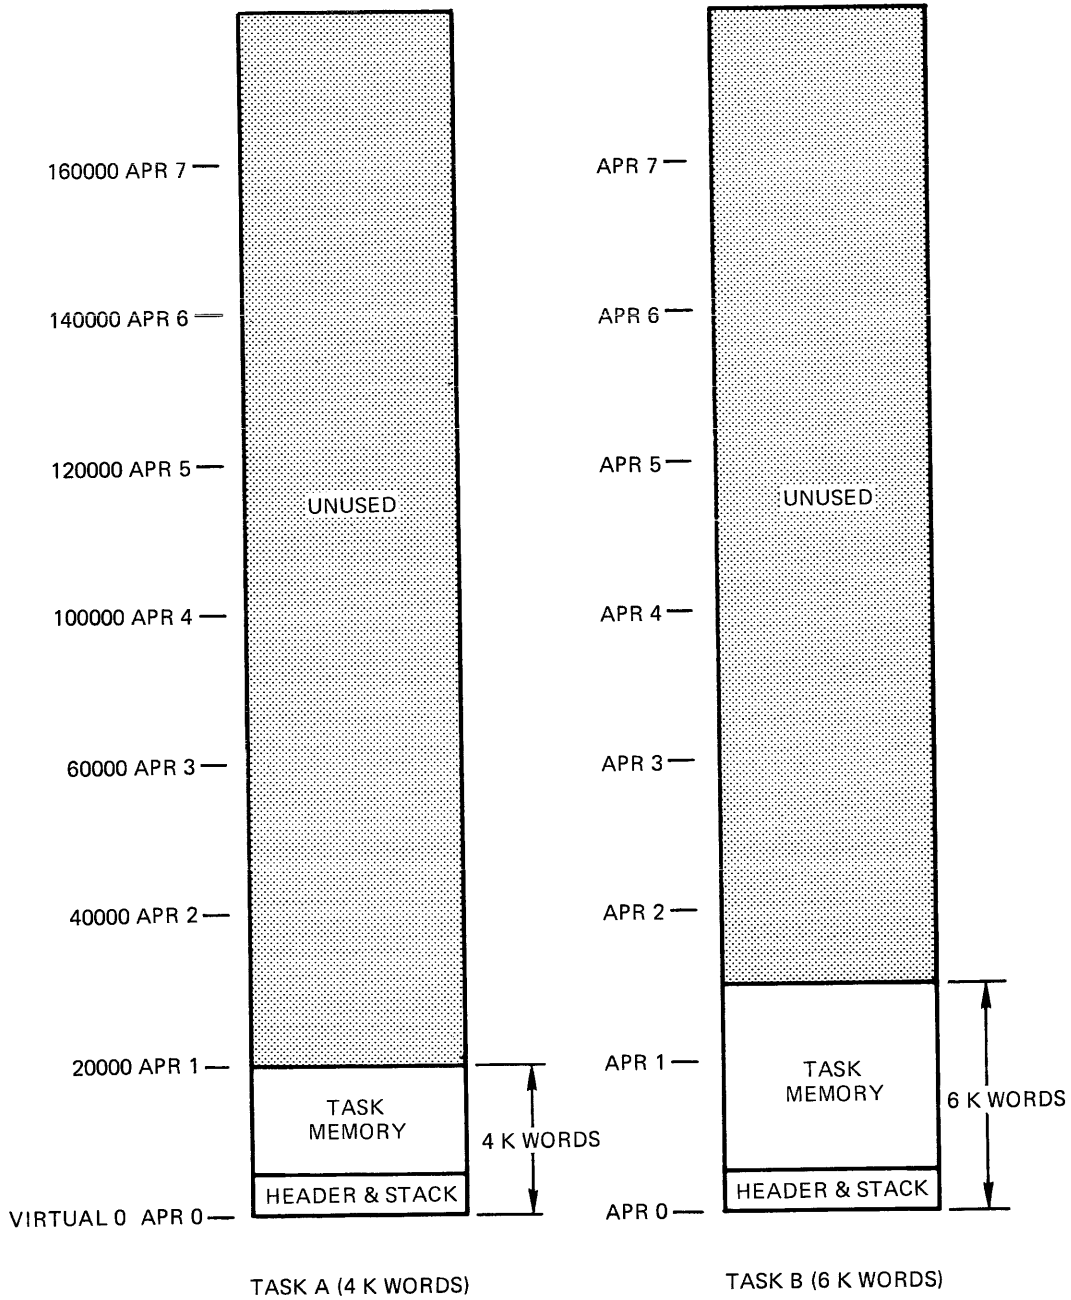
Finally, the layout of the virtual address space for a task that is to run in a mapped system is different in most cases from that of a task that is run in an unmapped system. Unless a task is privileged, the I/O page and the Executive are not normally part of a task's virtual address space and, unlike in an unmapped system, a task is inhibited by the system from accessing any portion of physical memory that it does not specifically own. Because the I/O page and the Executive are not part of a task's virtual address space, a task can be approximately 32,767 words long (32K minus 32 words needed by the loader) on a mapped system. TKB can build a task of 32K minus 1 word in size. However, overlaid tasks, and tasks that become extended, may use the entire 32K-word space.

2.4.4 Regions

This section briefly describes regions and their relationship to and use by tasks. Regions and their use are more thoroughly described in Chapter 5.

A region is a defined area of memory that can contain code or data. It can also be a blank area reserved for use by one or more tasks. The region is named and built like a task except that the /HD header switch is negated (/HD) because the region is not a task and does not need a task header. Tasks can also create regions dynamically as they execute. Dynamic regions are useful because they increase the task's logical address space while saving its virtual address space. Regions also allow tasks to share code and data with other tasks.

TASK BUILDER FUNCTIONS



ZK-388-81

Figure 2-12 Mapping for 4K-Word and 6K-Word Tasks

Regions are named according to their use or the way in which they were built. These regions are:

- Task Region -- A continuous block of memory in which the task runs.
- Common Shared Region -- On unmapped systems, a shared region defined by an operator at run time or built into the system during system generation; for example, a global common area.

TASK BUILDER FUNCTIONS

Resident commons are usually called shared regions because they are used as an area in which tasks share common data. Shared regions can be absolute or position independent. Shared regions and their use are described in Chapter 5.

- Library Shared Region -- A shared region containing common code or routines shared by tasks, and in this way saving virtual address space in the tasks.
- Dynamic Region -- A region created dynamically at run time by the Create Region (CRRG\$) memory management directive in the task. This directive and associated directives are described in the RSX-11M/M-PLUS Executive Reference Manual.

By convention, a shared region that contains code is a library and a shared region that contains data is a common.

Tasks must map to a region by using task windows which must be defined and numbered in the task when the task is built. Usually, a task uses one window for each region to which mapping must occur. Task windows are described in the next section, Task Mapping and Windows.

Figure 2-14 shows a sample collection of regions that could make up a task's logical address space. A task's logical address space can expand and contract dynamically as the task issues the appropriate memory management directives. The header and root segment are always part of the region. Therefore, the task header and root segment always use window 0 (UAPR 0) and region 0. Because a region occupies a continuous area of memory, each region is shown as a separate block.

2.5 TASK MAPPING AND WINDOWS

As mentioned earlier, tasks that run on mapped systems must be relocated at run time. When you build a task that is to run on a mapped system, TKB creates and places in the header of the task one or more 8-word data structures called window blocks. When you install a task, INSTALL initializes the window block(s). Once initialized, a window block describes a range of continuous virtual addresses called a window.

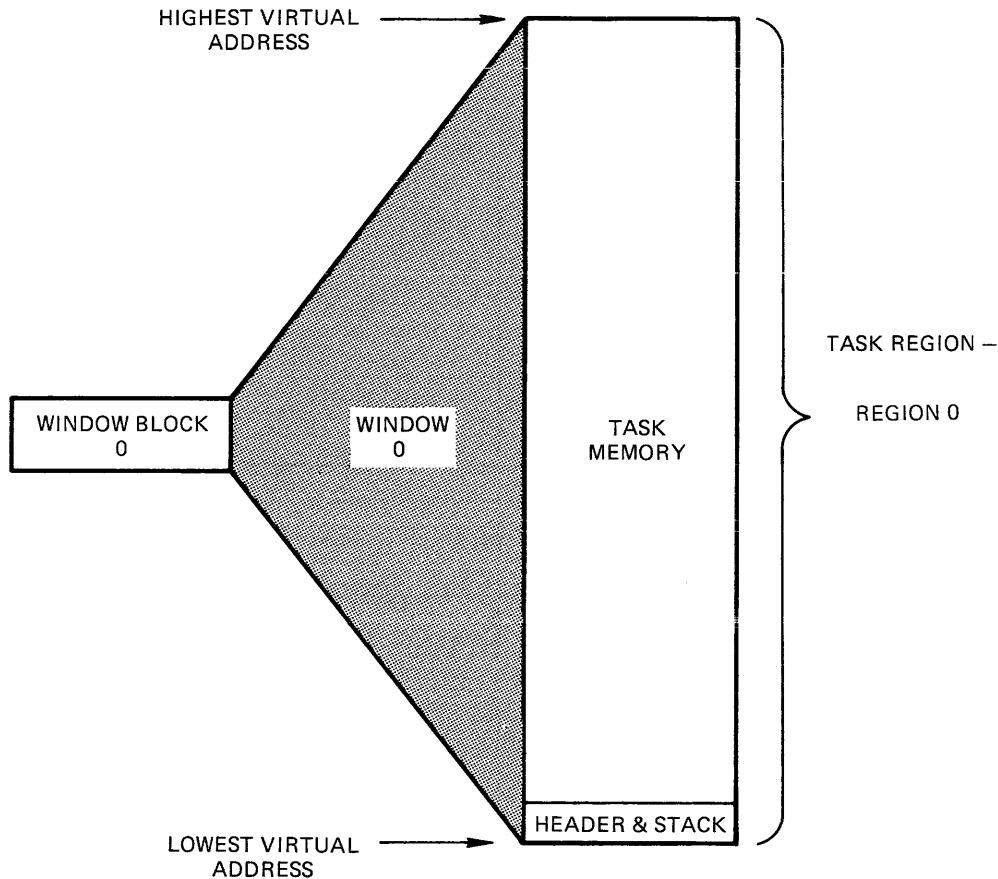
2.5.1 Task Windows

A window can be as small as 32 words or as large as 32K words. When a task consists of one continuous range of addresses (a single region task) only one window block is required to describe the entire task from the beginning of its header to the highest virtual address in the task. When a task consists of two or more regions (such as a task that references a shared region as described in Chapter 5), each region must have at least one window block associated with it that describes all or a portion of the region.

When the Executive maps a task into physical memory, it extracts the information it requires to set up the APRs of the memory management unit from the task's window block.

TASK BUILDER FUNCTIONS

In an RSX-11M system, regardless of the number of regions associated with a task, the region that contains the task's header and root is always described by window 0. Windows 0 and 1 describe the root of an I- and D-space task in an RSX-11M-PLUS system. Window 0 describes the I-space root and window 1 describes the D-space root and task header. Furthermore, this region is referred to as the task region and is identified as region 0. Figure 2-13 illustrates window block 0 for a system without I- and D-space. Windows for an I- and D-space task are described in Chapter 7.



ZK-389-81

Figure 2-13 Window Block 0

When you run your task, the Executive determines where in physical memory the task is to reside. The Executive then loads the Page Address Register portion of the APRs with a relocation constant that, when combined with the addresses of the task, yields the 18- or 22-bit physical address range of the task.

Referring to Figure 2-14, which illustrates a mapped system without I- and D-space, you can observe that a large 32K user task contains three distinct areas of continuous space called "windows." The term "task window" is a construct that maps a continuous portion of the task's virtual address space to a continuous portion of a region in the task's logical address space. Windows must have a specified size and starting address. The window size can be from 32 words to 32K minus

TASK BUILDER FUNCTIONS

32 words, and windows must start on a 4K address boundary. Figure 2-14 shows three windows that are not continuous in the task's virtual address space. However, the space within each window is continuous. In this task, the size of window 0 is 11K; the size of window 1 is 11K; and the size of window 2 is 8K. The concept of windows exists for the following specific reason.

By using the concept of windows and the memory management directives, a nonprivileged task can access a larger logical memory space than that implied by the 32K virtual addressing range and normally accessible by the 16-bit address. A task can, in fact, only access 32K of memory at one time. However, a nonprivileged task can change its access to logical addresses (real, physical memory). The area that your program accesses can be changed by the program during program execution. The process of accessing different logical areas of memory is called "mapping."

By referring to Figure 2-14, you can see that window 1 in the task is mapped to region 1 in physical memory. The task can change the window 1 mapping to region 0 in physical memory. In effect, then, though a task is limited to a range of 32K virtual addresses, a task can access all the physical memory available to it (determined by the way that you set up the mapping) by changing the mapping of its windows to different logical addresses. Figure 2-14 provides a visual description of the concept of mapping to different logical addresses.

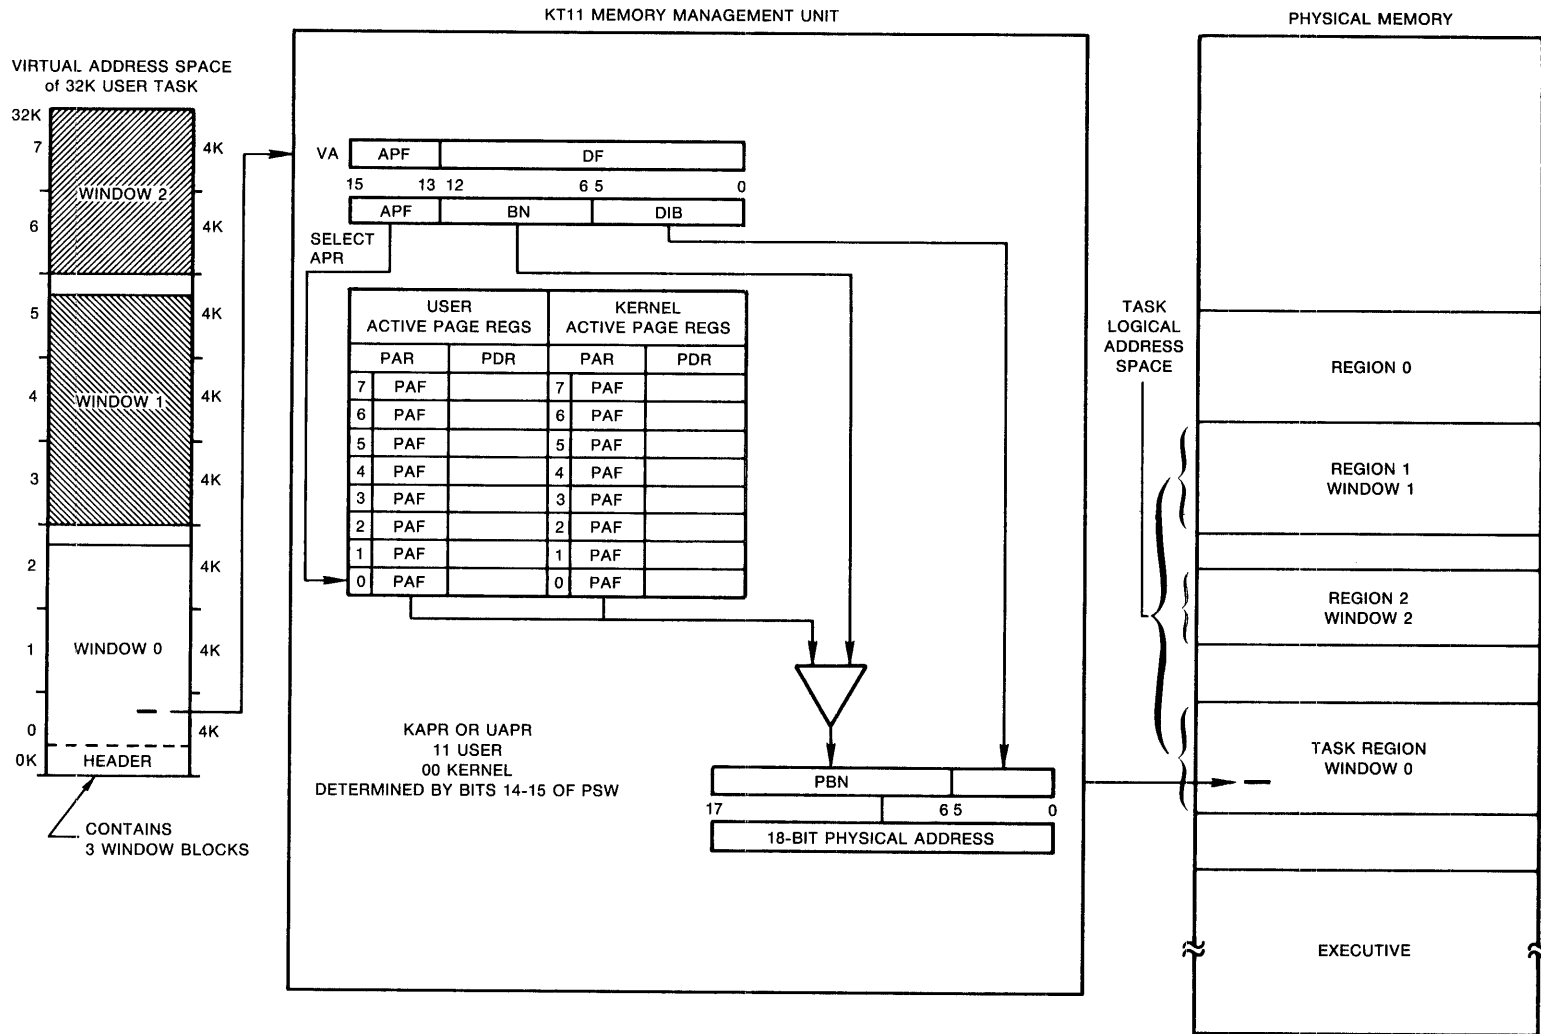
The discussion now proceeds to setting up the task's windows. This is done by defining task window blocks to TKB.

To manipulate virtual address mapping to various logical areas, you must first divide a task's 32K of virtual address space into segments. These segments are task (virtual address) windows. Each window encompasses a continuous range of virtual addresses. The first address of the window address range must be a multiple of 4K (the first address must begin on a 4K boundary) because of the way that the KT-11 memory management unit uses APRs.

On an RSX-11M system, you can specify up to seven windows. Task mapping for the task's code requires the use of window 0. Therefore, there is a total of eight windows. However, window 0 is not available to nonprivileged tasks. The size of each window can range from a minimum of 32 words to a maximum of 32K minus 32 words.

RSX-11M-PLUS tasks that use I- and D-space or supervisor-mode libraries have a total of 16 windows. You can specify up to 14 windows in this type of task. Windows 0 and 1 are not available to nonprivileged tasks in this kind of system.

A task that includes directives that dynamically manipulate address windows must have task window blocks set up in the task header as well as Window Definition Blocks in the code for use by the Create Address Window directive. The Executive uses task window blocks to identify and describe each currently existing window. When linking the task, the programmer specifies the number of extra window blocks needed by the task. The number of blocks should equal the maximum number of windows that will exist concurrently while the task is running.



ZK-390-81

Figure 2-14 Virtual to Logical Address Space Translation

TASK BUILDER FUNCTIONS

In RSX-11M or RSX-11M-PLUS without I- and D-space, a window's identification is a number from 0 to 7, which is an index to the window's corresponding window block. The address window identified by 0 is the window that always maps the task's header and root segment. TKB creates window 0, which the Executive uses to map the task. No directive may specify window 0; a directive that does so is rejected.

In an RSX-11M-PLUS system using an I- and D-space task, a window's identification is a number from 0 to 15, which is an index to the window's corresponding window block. The address windows identified by 0 and 1 are the windows that always map the task's header and root. TKB creates windows 0 and 1, which the Executive uses to map the task. No directive may specify windows 0 or 1; a directive that does so is rejected.

When a task uses memory management directives, the Executive views the relationship between the task's virtual and logical address space in terms of windows and regions. Unless a virtual address is part of an existing address window, the address does not point anywhere. This is a point to watch when setting up windows with the Create Address Window directive (CRAW\$). Similarly, a window can be mapped only to an area that is all or part of an existing region within the task's logical address space.

Once a task has defined the necessary windows and regions, the task can issue memory management directives to perform operations such as the following:

- Map a window to all or part of a region.
- Unmap a window from one region in order to map it to another region.
- Unmap a window from one part of a region in order to map it to another part of the same region.

2.6 RSX-11M-PLUS SUPERVISOR MODE

Three modes of operation are possible in the PDP-11: user mode, supervisor mode, and kernel mode. Each mode has associated with it 16 APRs for mapping memory: 8 I-space APRs and 8 D-space APRs. A task in the RSX-11M-PLUS system can use supervisor-mode libraries and thereby double the task's virtual address space to 64K words. Supervisor-mode libraries are described in Chapter 8. This section briefly describes supervisor mode and the mapping that occurs when the task uses supervisor mode.

Supervisor-mode libraries are libraries of routines that are used only in supervisor mode. The task switches to supervisor mode when it calls a routine within the supervisor-mode library. By using a supervisor mode library as described in Chapter 8, you make the RSX-11M-PLUS system, for large systems, use the supervisor-Mode APRs.

2.6.1 Supervisor-Mode Mapping

Normally, a task has an address space of 32K-words by using eight user APRs. When a conventional task links to a supervisor-mode library and calls a routine in the library, the Executive copies the user-mode I-space APRs into the supervisor-mode D-space APRs and maps the supervisor-mode library with supervisor I-space APRs. Therefore, while in supervisor mode and within the library, the task can access 32K-words of its own space with D-space APRs and 32K-words of library

TASK BUILDER FUNCTIONS

routines with I-space APRs. The amount of possible logical address space totals to 64K-words.

When an I- and D-space task links to a supervisor-mode library, the Executive copies the user-mode D-space APRs into the supervisor D-space APRs. Therefore, the supervisor-mode routines can access user data space and access supervisor-mode instruction space with supervisor I-space APRs. Figure 8-2 illustrates this mapping. The mapping just described is the default mapping for an I- and D-space task. You can explicitly create supervisor-mode D-space mapping to override the user-mode D-space overmapping that occurs by using the MSDSS\$ Executive directive. Chapter 8 discusses the use of MSDSS\$.

The mapping of a conventional task in a system that contains a supervisor-mode library is shown in Figure 2-15.

The mapping of a conventional task in a system while using a supervisor-mode library is shown in Figure 2-16.

2.7 PRIVILEGED TASKS


RSX-11M/M-PLUS systems have two classes of tasks: privileged and nonprivileged. However, the term "privileged" has meaning in mapped systems only, because in mapped systems certain areas of memory are protected from nonprivileged tasks. In an unmapped system, any task has the ability to access all of physical memory if so programmed. Therefore, the distinction between these two classes of tasks is primarily one of their mapping to memory in a mapped system.

Privileged tasks in a mapped system can access system data areas and the Executive. Altering system data areas or the Executive can cause obscure and difficult problems. Therefore, privileged tasks must be programmed and used with all caution.

You can specify a task as privileged by using the /PR switch in the TKB command line. The /PR:0 switch allows a task to perform certain privileged operations; but, the /PR:0 task cannot access the Executive or system data structures. The /PR:4 switch allows the task to directly map the I/O page, Executive routines, and system data structures. The /PR:4 switch is used for a privileged task in a system that has an Executive of 16K or less. The /PR:5 switch allows a task to directly map to the I/O page, Executive routines, and system data structures. The /PR:5 switch is used for a privileged task in a system that has an Executive of 20K or less.

Chapter 6 describes privileged tasks and their mapping in detail.


TASK BUILDER FUNCTIONS



ZK-391-81

Figure 2-15 Mapping for a Conventional User Task and a System Containing a Supervisor-Mode Library in an RSX-11M-PLUS System

TASK BUILDER FUNCTIONS



ZK-392-81

Figure 2-16 Mapping for a Conventional User Task Using a Supervisor-Mode Library in an RSX-11M-PLUS System

TASK BUILDER FUNCTIONS

2.8 MULTIUSER TASKS (RSX-11M-PLUS ONLY)

The following section is an introduction to multiuser tasks, which are fully described in Chapter 9.

TKB allows you to build multiuser tasks. A multiuser task is that which has one portion of its code and data designated as read-only and another portion designated as read/write. You specify the read-only portions of your task with program sections that have the read-only access code. When you then build your task with the /MU switch, TKB places the read-only portions in a region that has a high virtual address and the read/write portion in a region that has a low virtual address. Any other requests to run the task, if the task is already running, results in a copy of the read/write portion of the task in physical memory for the other user. There is always only one copy of the read-only code regardless of the number of tasks that may be running.

The /MU switch is described in Chapter 10.

2.9 USER-MODE I- AND D-SPACE TASKS (RSX-11M-PLUS)

User tasks that use both I- and D-space differ from conventional tasks because I- and D-space tasks have specifically defined locations within the task for both instructions and data. Because of this separation, the I- and D-space task image is structurally different.

Additionally, the separate instruction areas are mapped through separate APRs in the memory management unit. Hence, up to 8 user-mode instruction APRs map the task's instructions, and up to 8 user-mode data APRs map the task's data.

Also, overlaid I- and D-space tasks are more complex because each overlaid part (segment) of such a task may reside in both instruction space and data space.

I- and D-space tasks differ from conventional tasks in the following major ways:

- PSECTs with the "I" attribute contain only instructions and PSECTs with the "D" attribute contain only data.
- Two sets of APRs map the I- and D-space task: the I-space APRs and the D-space APRs.
- I- and D-space tasks can use up to 64K words of virtual space instead of 32K words because of their use of the two sets of APRs. With supervisor-mode libraries, an I- and D-space task can use up to 96K words of virtual space.

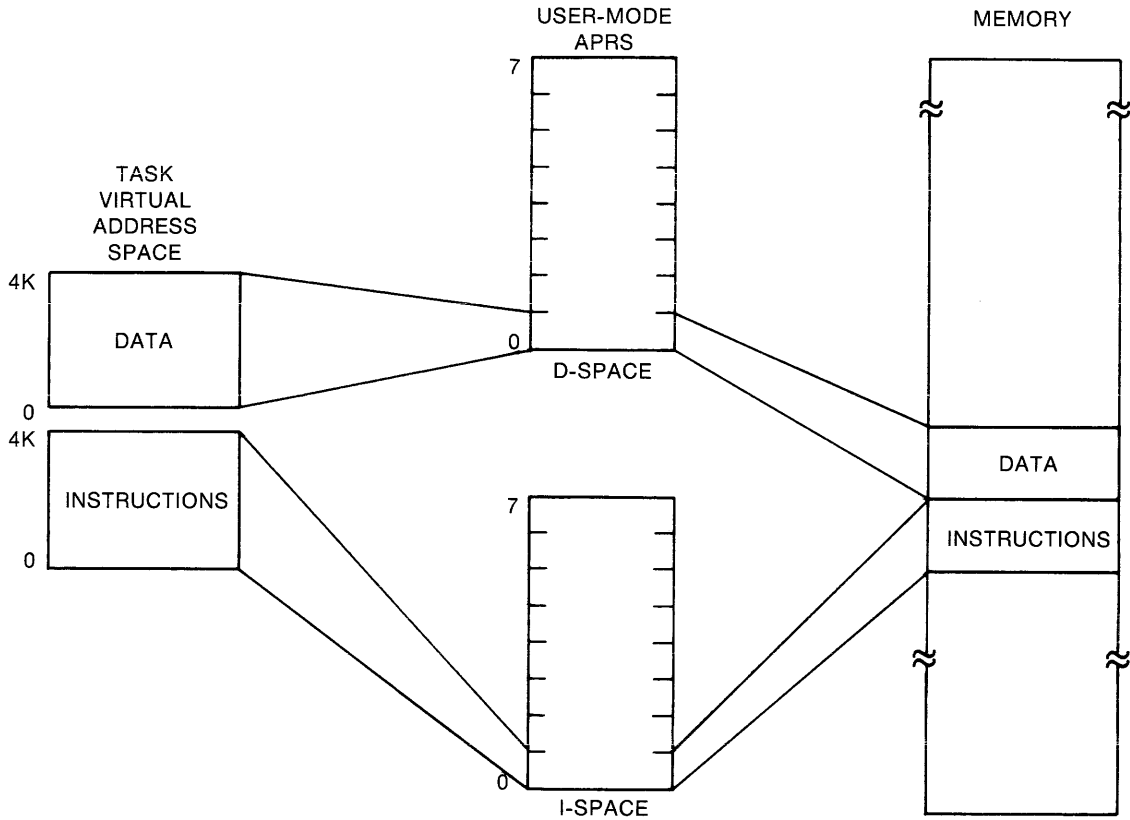
The following have data contiguously adjacent in memory, and instructions contiguously adjacent in memory:

- Non-overlaid I- and D-space tasks
- Segments in an overlaid I- and D-space task that contain both I- and D-space

TASK BUILDER FUNCTIONS

In these tasks or task segments, I-space PSECTs are segregated from D-space PSECTs in memory; they cannot be intermixed.

Figure 2-17 shows a user-mode I- and D-space task with data separated from the instructions and mapped to memory through two sets of APRs.



ZK-1049-82

Figure 2-17: Simplified APR Mapping for an I- and D-space Task

I- and D-space tasks are more fully described in Chapter 7. The task images for both conventional and I- and D-space tasks are described in Appendix B.

CHAPTER 3

OVERLAY CAPABILITY

TKB provides you with the means to reduce the memory and/or virtual address space requirements of your task by using tree-like overlay structures created with the Overlay Description Language (ODL). You can divide your conventional task into pieces called segments, which are loadable with one disk access. In an I- and D-space task, an overlaid segment that contains I- and D-space PSECTS requires a maximum of two disk accesses to load the segment. The segments are the discrete parts of the overlay structure that form the tree. You can specify two kinds of overlay segments: those that reside on disk, and those that reside permanently in memory after being loaded from disk.

3.1 OVERLAY STRUCTURES

To create an overlay structure, you divide a task into a series of segments consisting of:

- A single root segment, which is always in memory
- Any number of overlay segments, you must consider which either 1) reside on disk and share virtual address space and physical memory with one another (disk-resident overlays); or 2) reside in memory and share only virtual address space with one another (memory-resident overlays)¹

Segments consist of one or more object modules, which in turn consist of one or more program sections. Segments that overlay each other must be logically independent; that is, the components of one segment cannot reference the components of another segment with which it shares virtual address space. In addition to the logical independence of the overlay segments, you must consider the general flow of control within the task when creating overlay segments.

You must also consider the kind of overlay segment to create at a given position in the structure, and how to construct it. Dividing a task into disk-resident overlays saves physical space, but introduces the overhead activity of loading these segments each time they are needed -- but are not present -- in memory. Memory-resident overlays, on the other hand, are loaded from disk only the first time they are referenced. Thereafter, they remain in memory and are referenced by remapping.

1. Note that memory-resident overlays can be used only if the hardware has a memory management unit, and if support for the memory management directives has been included in the system on which the task is to run.

OVERLAY CAPABILITY

Several large classes of tasks can be handled effectively when built as overlay structures. For example, a task that moves sequentially through a set of modules is well suited to use as an overlay structure. A task that selects one of a set of modules according to the value of an item of input data is also well suited to use as an overlay structure.

Tasks that have separate I- and D-space may also use overlays where the root has instructions and data separately defined by PSECTs, and each individual segment of the task also has instructions and data separately defined. Chapter 7 contains more information about I- and D-space tasks.

3.1.1 Disk-Resident Overlay Structures

Disk-resident overlays conserve virtual address space and physical memory by sharing them with other overlays. Segments that are logically independent need not be present in memory at the same time. They, therefore, can occupy a common physical area in memory (and, therefore, common virtual address space) whenever either needs to be used.

The use of disk-resident overlays is shown in this section by an example, task TK1, which consists of four input files. Each input file consists of a single module with the same name as the file. The task is built by the command

```
>TKB TK1=OVLAY.ODL/MP
```

and the file OVLAY.ODL contains the modules CNTRL, A, B, C in an overlay description for the task being built. The /MP switch specifies that the input file is an Overlay Description Language (ODL) file.

In this example, the modules A, B, and C are logically independent; that is:

A does not call B or C and does not use the data of B or C.

B does not call A or C and does not use the data of A or C.

C does not call A or B and does not use the data of A or B.

A disk-resident overlay structure can be defined in which A, B, and C are overlay segments that occupy the same storage area in physical memory. The flow of control for the task is as follows:

CNTRL calls A and A returns to CNTRL.

CNTRL calls B and B returns to CNTRL.

CNTRL calls C and C returns to CNTRL.

CNTRL calls A and A returns to CNTRL.

In this example, the loading of overlays occurs only four times during the execution of the task. Therefore, the virtual address space and physical memory requirements of the task can be reduced without unduly increasing the overhead activity.

The effect of the use of an overlay structure on allocating virtual address space and physical memory for task TK1 is described in the following paragraphs.

OVERLAY CAPABILITY

The lengths of the modules are:

Module	Length (in Octal)
CNTRL	20000 bytes
A	30000 bytes
B	20000 bytes
C	14000 bytes

Figure 3-1 shows the virtual address space and physical memory required as a result of building TK1 as a single-segment task on a system with memory management hardware.

The virtual address space and physical memory requirement to build TK1 as a single-segment task is 104000(octal) bytes.

In contrast, Figure 3-2 shows the virtual address space and physical memory required as a result of building TK1 as a multisegment task and using the overlay capability.

The multisegment task requires 50000(octal) bytes.

NOTE

In addition to the storage required for modules A, B, and C, storage is required for overhead in handling the overlay structures. This overhead is not reflected in this example.

In using the overlay capability, the total amount of virtual address space and physical memory required for the task is determined by the sum of the length of the root segment and the length of the longest overlay segment. Overlay segments A and B in this example are much longer than overlay segment C. If A and B are divided into sets of logically independent modules, task storage requirements can be further reduced. Segment A can be divided into a control program (A0) and two overlays (A1 and A2). Segment A2 can then be divided into the main part (A2) and two overlays (A21 and A22). Similarly, segment B can be divided into a control module (B0) and two overlays (B1 and B2).

Figure 3-3 shows the virtual address space and physical memory required for the task produced by the additional overlays defined for A and B.

As a single-segment task, TK1 requires 104000(octal) bytes of virtual address space and physical memory. The first overlay structure reduces the requirement by 34000(octal) bytes. The second overlay structure further reduces the requirement by 14000(octal) bytes.

The vertical and horizontal lines in the diagrams of Figures 3-2 and 3-3 represent the state of virtual address space and physical memory at various times during the calling sequence of TK1. For example, in Figure 3-3 the leftmost vertical line in both diagrams shows virtual address space and physical memory, respectively, when CNTRL, A0, and A1 are loaded. The next vertical line shows virtual address space and physical memory when CNTRL, A0, A2, and A21 are loaded, and so on.

OVERLAY CAPABILITY

The horizontal lines in the diagrams of Figures 3-2 and 3-3 indicate segments that share virtual address space and physical memory. For example, in Figure 3-3, the uppermost horizontal line of the task region in both diagrams shows A1, A21, A22, B1, B2, and C, all of which can use the same virtual address space and physical memory. The next horizontal line shows A1, A2, B1, B2, and C, and so on.

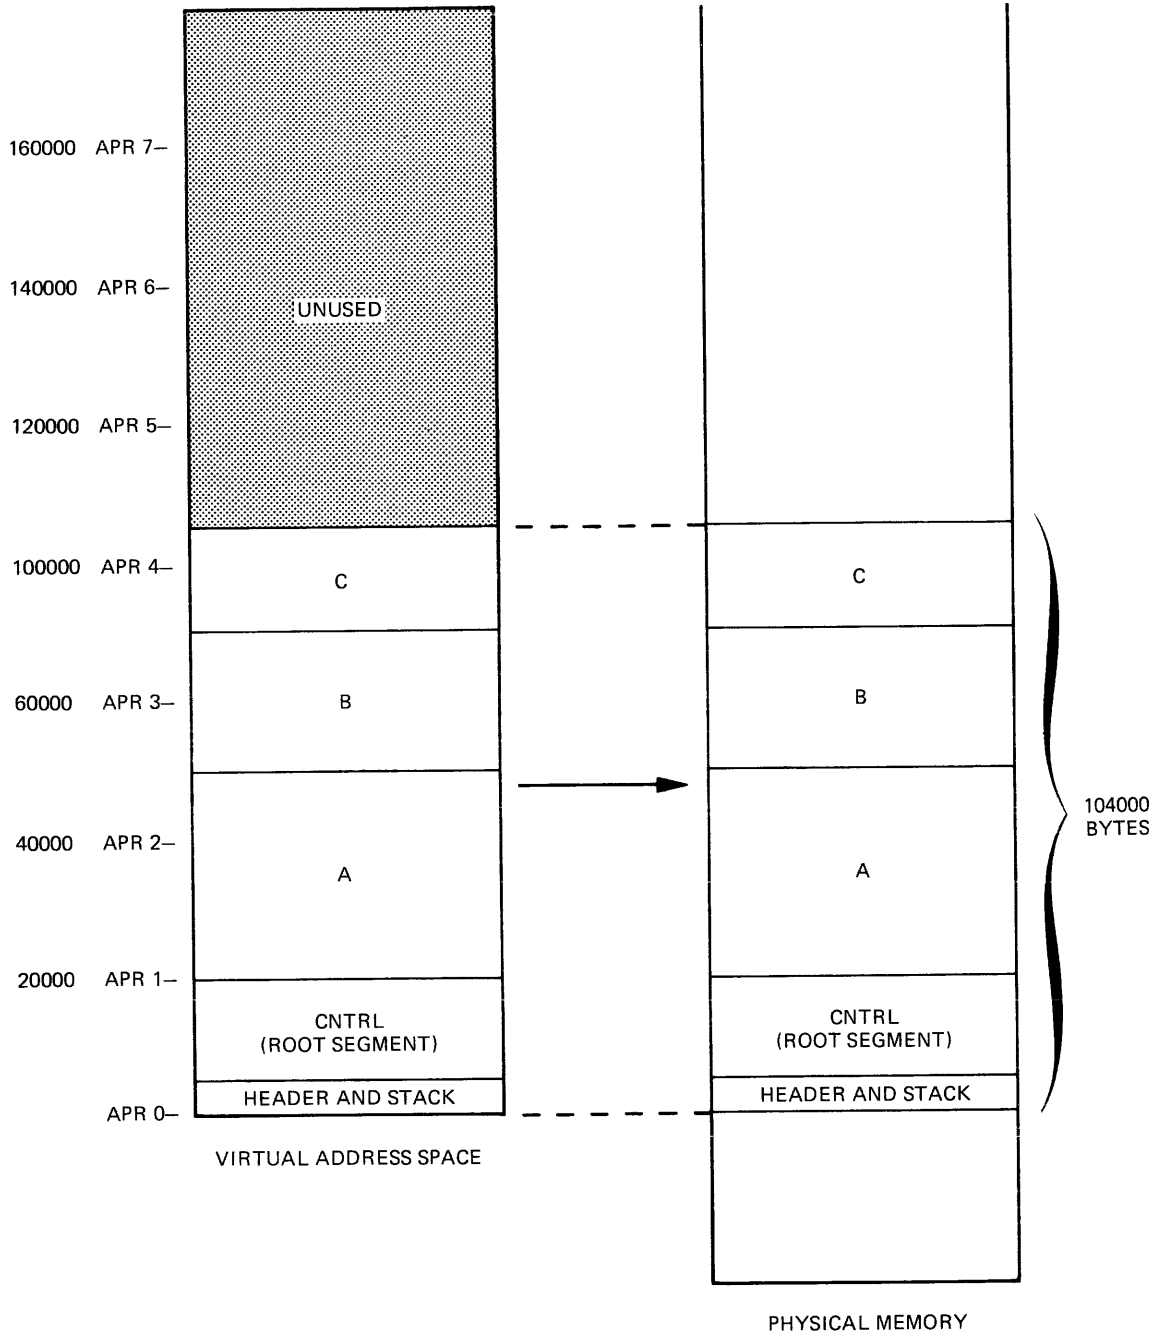
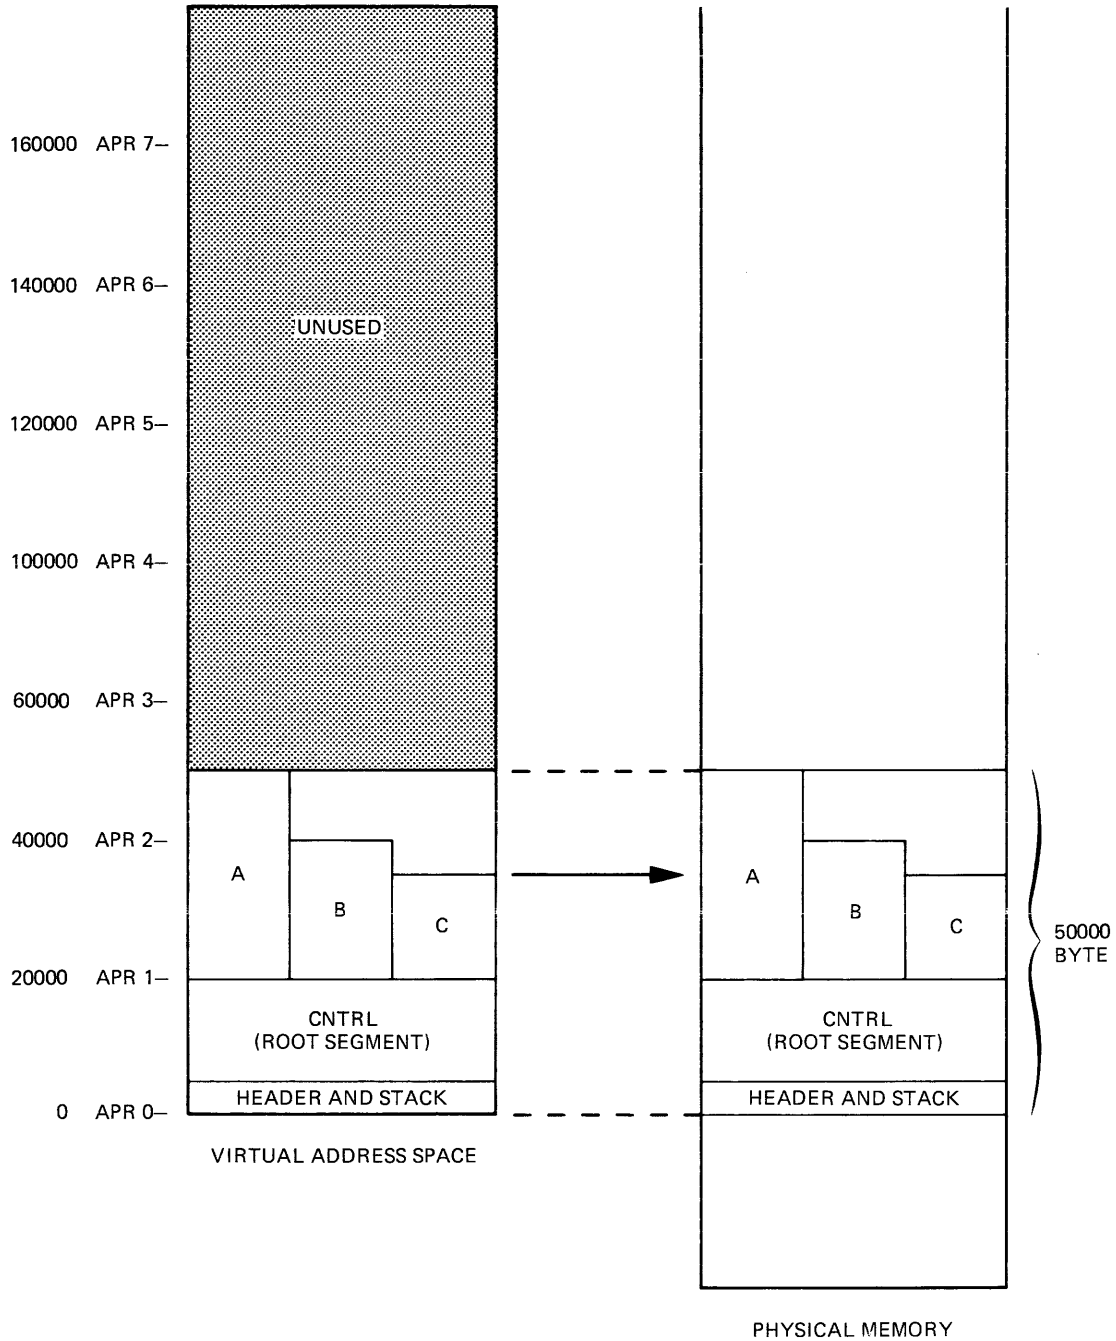


Figure 3-1 TK1 Built As a Single-Segment Task

ZK-393-81

OVERLAY CAPABILITY



ZK-394-81

Figure 3-2 TK1 Built As a Multisegment Task

3.1.2 Memory-Resident Overlay Structures (Not Supported on RSX-11S)

TKB provides for creating overlay segments that are loaded from disk only the first time they are referenced. Thereafter, they reside in memory. Memory-resident overlays share virtual address space just as disk-resident overlays do but, unlike disk-resident overlays, memory-resident overlays do not share physical memory. Instead, they reside in separate areas of physical memory, each segment aligned on a 32-word boundary. Memory-resident overlays save time for a running

OVERLAY CAPABILITY

task because they do not need to be copied from a secondary storage device each time they are to overlay other segments. "Loading" a memory-resident overlay reduces to mapping a set of shared virtual addresses to the unique physical area of memory containing the overlaying segment.

The use of memory-resident overlays is shown in this section by an example, task TK2, which consists of four input files. Each input file consists of a single module with the same name as the file. The task is built by the command

```
>TKB TK2=OVLAY2.ODL/MP
```

and the file OVLAY2.ODL contains the modules CNTRL, D, E, and F in an overlay description for the task being built. The /MP switch specifies that the input file is an Overlay Description Language (ODL) file.

In this example, the modules D, E, and F are logically independent; that is:

D does not call E or F and does not use the data of E or F.

E does not call D or F and does not use the data of D or F.

F does not call D or E and does not use the data of D or E.

A memory-resident overlay structure can be defined in which D, E, and F are overlay segments that occupy separate physical memory locations but the same virtual address space. The flow of control for the task is as follows:

CNTRL calls D and D returns to CNTRL.

CNTRL calls E and E returns to CNTRL.

CNTRL calls F and F returns to CNTRL.

The effect of the use of a memory-resident overlay structure on allocating virtual address space and physical memory for task TK2 is described in the following paragraphs.

The lengths of the modules are:

Module	Length (in Octal)
CNTRL	20000
D	10000
E	14000
F	12000

Figure 3-4 shows the virtual address space and physical memory requirements as a result of building TK2 as a single-segment task on a system with memory management hardware.

The virtual address space and physical memory requirements when TK2 is built as a single-segment task is 56000(octal) bytes.

If TK2 is built using the Task Builder's memory-resident overlay capability, the relationship of virtual address space to physical memory changes, as shown in Figure 3-5.

OVERLAY CAPABILITY


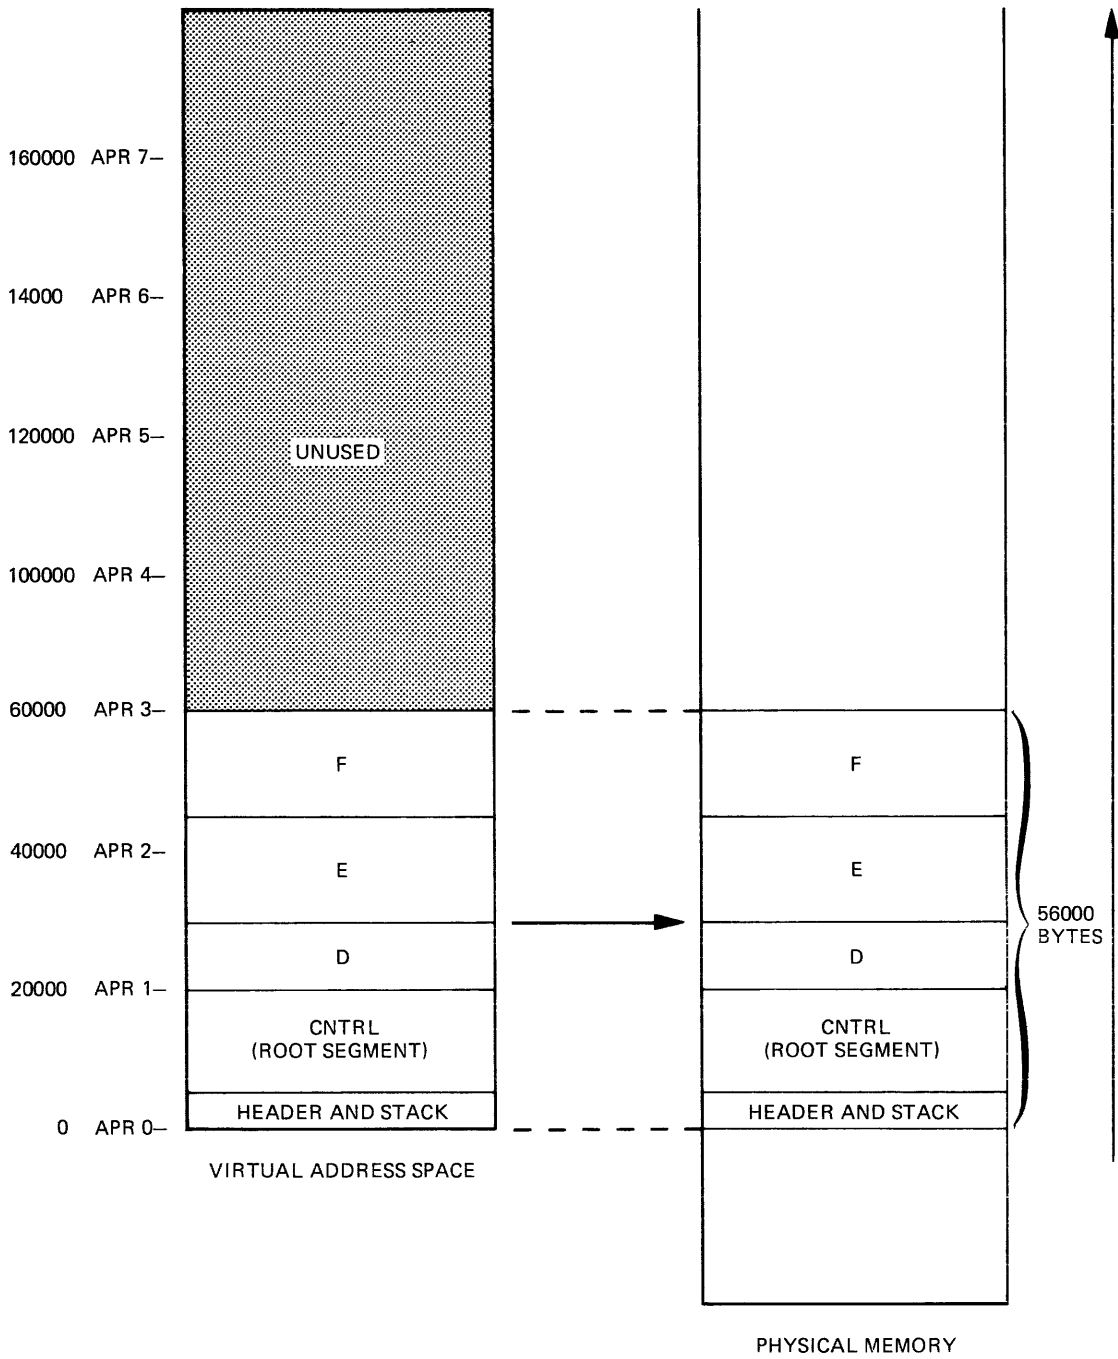


Figure 3-3 TK1 Built with Additional Overlay Defined

ZK-395-81

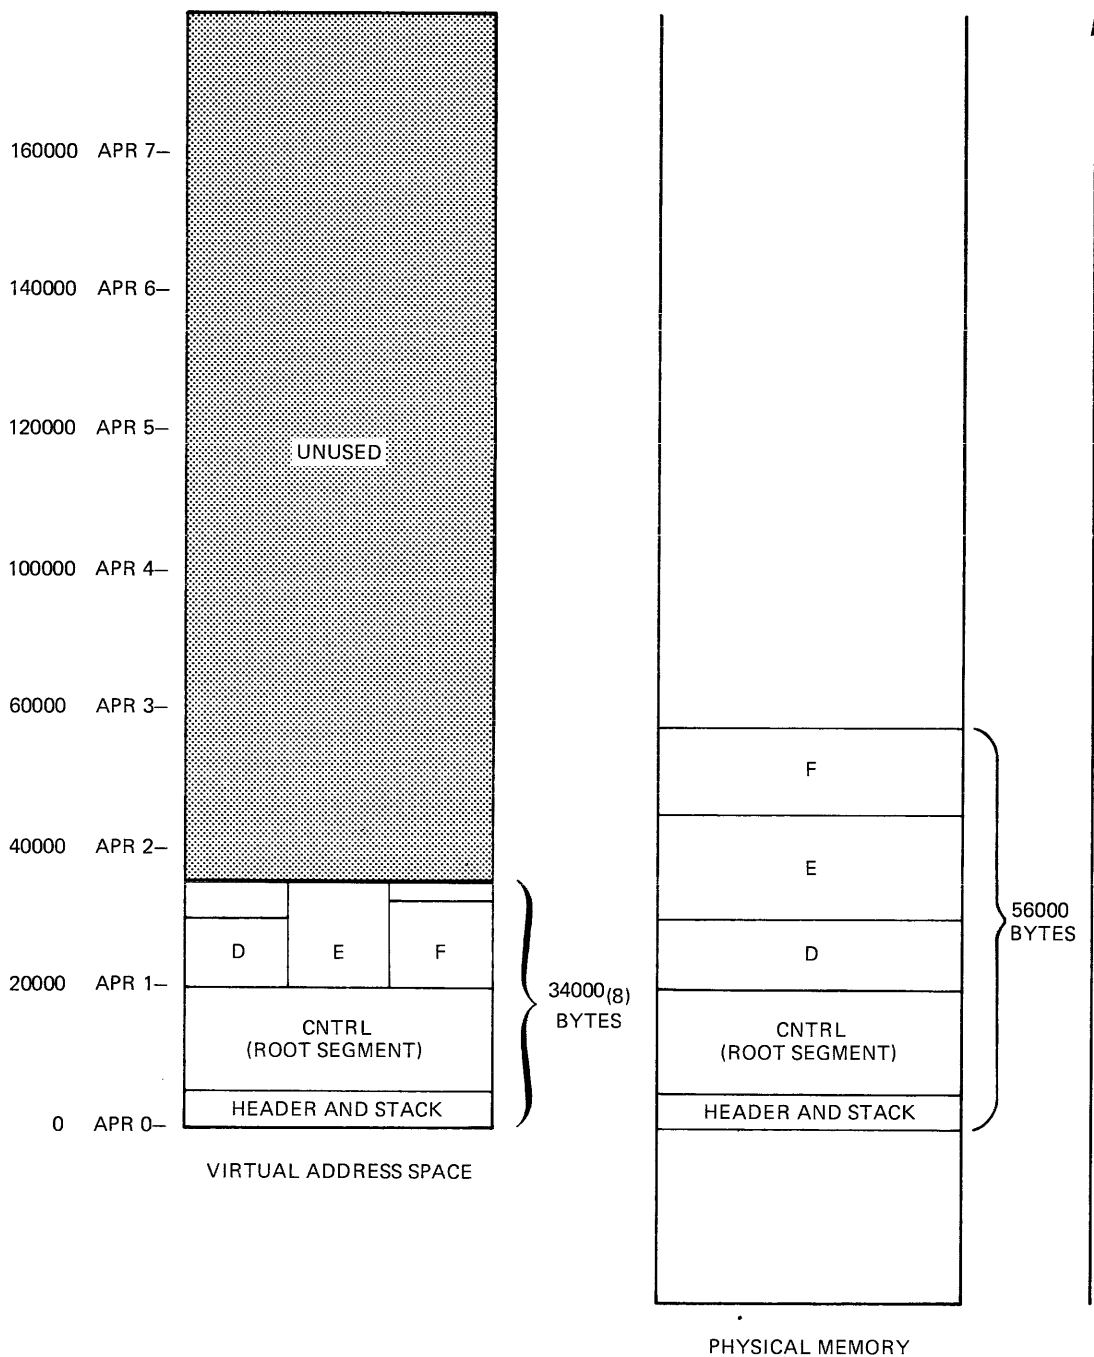
OVERLAY CAPABILITY



ZK-396-81

Figure 3-4 TK2 Built As a Single-Segment Task

OVERLAY CAPABILITY



ZK-397-81

Figure 3-5 TK2 Built As a Memory-Resident Overlay

The physical memory requirements for TK2 do not change (56000(octal) bytes), but the virtual address space requirements have been reduced to 34000(octal) bytes. This represents a savings in virtual address space of 22000(octal) bytes.

OVERLAY CAPABILITY

NOTE

In addition to the storage required for modules D, E, and F, storage is required for overhead in handling the overlay structures. This overhead is not reflected in this example.

In Figure 3-5, the vertical and horizontal lines in the virtual address space diagram represent the state of virtual address space at various times during the calling sequence of TK2. The leftmost vertical line shows virtual address space when CNTRL and D are loaded and mapped. The next vertical line shows virtual address space when CNTRL and E are loaded and mapped. The third vertical line shows virtual address space when CNTRL and F are loaded and mapped.

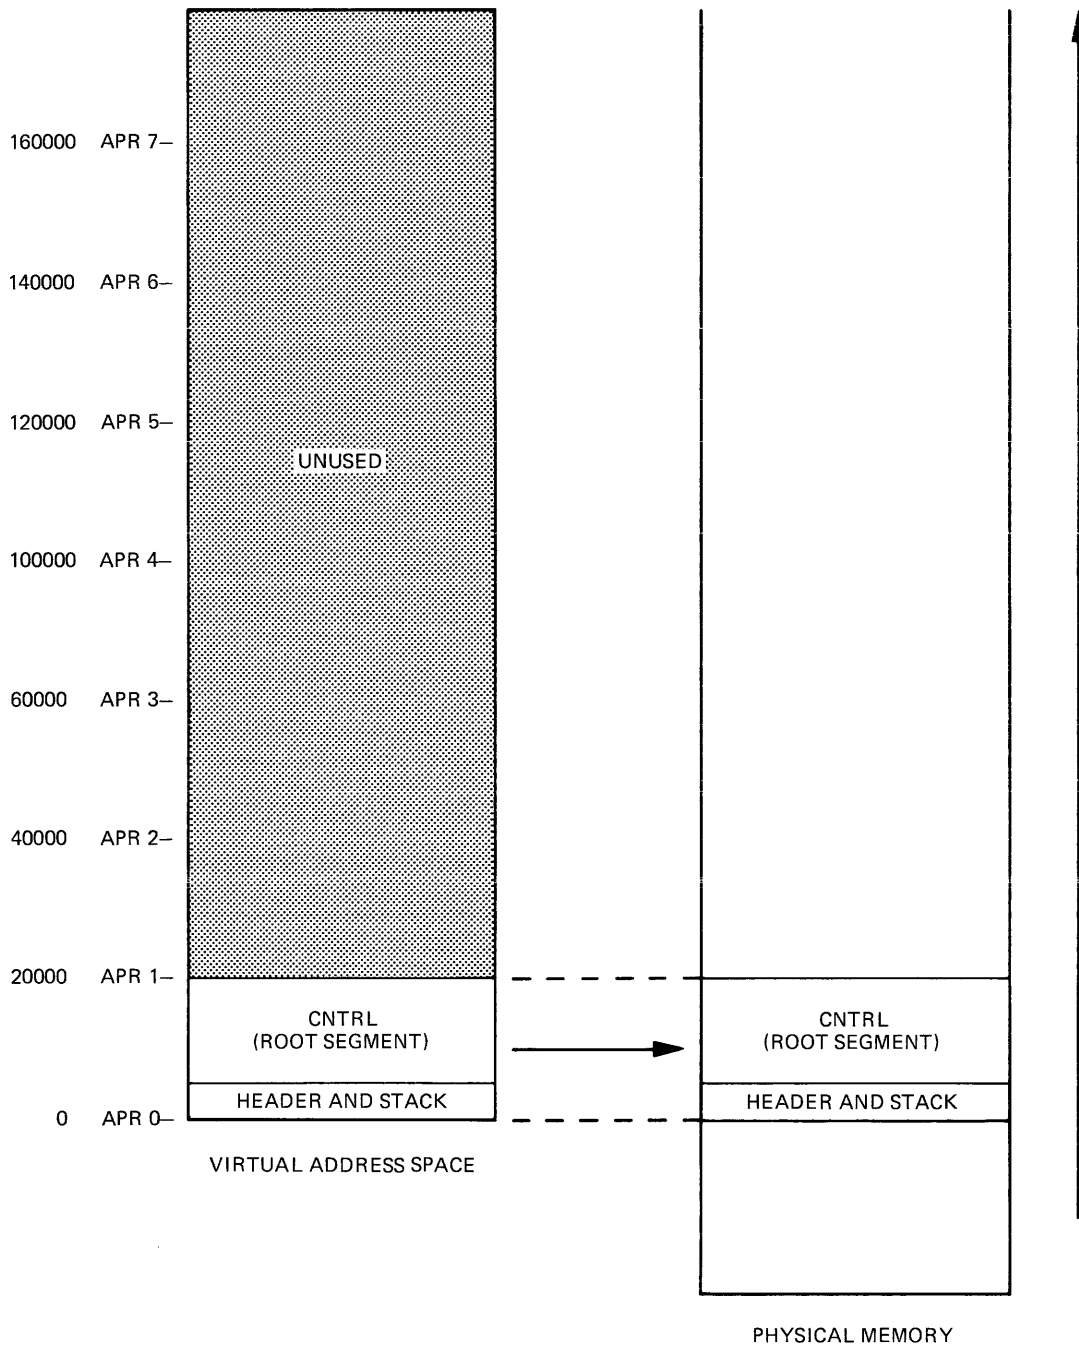
The uppermost horizontal line of the task region shows that segments D, E, and F share virtual address space.

When TK2 is activated, the Executive loads TK2's root segment into physical memory. The Executive loads segments D, E, and F into memory as they are called. Once all segments in the structure have been called, "loading" of the overlay segments reduces to the remapping of virtual address space to the physical locations in memory where the overlay segments permanently reside. Figures 3-6 and 3-7 illustrate the relationship between virtual address space and physical memory for task TK2 during four time periods:

- TIME 1 (Figure 3-6A) - TK2 is run and the system loads the root segment (CNTRL) into physical memory and maps to it.
- TIME 2 (Figure 3-6B) - CNTRL calls segment D. The system loads segment D into physical memory and maps to it. Segment D returns to CNTRL.
- TIME 3 (Figure 3-7A) - CNTRL calls segment E. The system loads segment E into physical memory, unmaps from segment D, and maps to segment E. Segment E returns to CNTRL.
- TIME 4 (Figure 3-7B) - CNTRL calls segment F. The system loads segment F into physical memory, unmaps from segment E, and remaps to segment F. Segment F returns to CNTRL.

OVERLAY CAPABILITY

Figure 3-6A Time 1




ZK-398-81

Figure 3-6A Relationship Between Virtual Address Space and Physical Memory -- Time 1

OVERLAY CAPABILITY

Figure 3-6B Time 2




ZK-399-81

Figure 3-6B Relationship Between Virtual Address Space and Physical Memory -- Time 2

OVERLAY CAPABILITY

Figure 3-7A Time 3




ZK-400-81

Figure 3-7A Relationship Between Virtual Address Space and Physical Memory -- Time 3

OVERLAY CAPABILITY

Figure 3-7B Time 4



ZK-401-81

Figure 3-7B Relationship Between Virtual Address Space and Physical Memory -- Time 4

It is important to be careful in choosing whether to have memory-resident overlays in a structure. Carelessly using these segments can result in inefficient allocation of virtual address space, because TKB allocates virtual address space in blocks of 4K words. Consequently, the length of each overlay segment should approach that limit if you are to minimize waste. (A segment that is one word longer than 4K words, for example, is allocated 8K words of virtual address space. All but one word of the second 4K words is unusable.)

OVERLAY CAPABILITY

You can also conserve physical memory by maintaining control over the contents of each segment. Including a module in several memory-resident segments that overlay one another causes physical memory to be reserved for each extra copy of that module. Common modules, including those from the system object module library (SYSLIB), should be placed in a segment that can be accessed from all referencing segments.

The primary criterion for choosing to have memory-resident overlays is the need to save virtual address space when disk-resident overlays are either undesirable (because they would slow down the system unacceptably), or impossible (because the segments are part of a resident library or other shared region that must permanently reside in memory).

Memory-resident overlays can help you use large systems to better advantage because of the time savings realized when a large amount of physical memory is available. Resident libraries, in particular, can benefit from the virtual address space saved when they are divided into memory-resident segments.

3.2 OVERLAY TREE

The arrangement of overlay segments within the virtual address space of a task can be represented schematically as a tree-like structure. Each branch of the tree represents a segment. Parallel branches denote segments that overlay one another and therefore have the same virtual address; these segments must be logically independent. Branches connected end to end represent segments that do not share virtual address space with each other; these segments need not be logically independent.

TKB provides an Overlay Description Language (ODL) for representing an overlay structure consisting of one or more trees (the ODL is described in Section 3.4).

The single overlay tree shown in Figure 3-8 represents the allocation of virtual address space for TK1 (see Section 3.1.1).

The tree has a root (CNTRL) and three main branches (A0, B0, and C). It also has six leaves (A1, A21, A22, B1, B2, and C).

The tree has as many paths as it has leaves. The path down is defined from the leaf to the root. For example:


A21-A2-A0-CNTRL

The path up is defined from the root to the leaf. For example:

CNTRL-B0-B1

Knowing the properties of the tree and its paths is important to understanding the overlay loading mechanism and the resolution of global symbols.

OVERLAY CAPABILITY



ZK-402-81

Figure 3-8 Overlay Tree for TK1

3.2.1 Loading Mechanism

Modules can call other modules that exist on the same path. The module CNTRL (Figure 3-8) is common to every path of the tree and, therefore, can call and be called by every module in the tree. The module A2 can call the modules A21, A22, A0, and CNTRL; but A2 cannot call A1, B1, B2, B0, or C.

When a module in one overlay segment calls a module in another overlay segment, the called segment must be in memory and mapped, or must be brought into memory. The methods for loading overlays are described in Chapter 4.

3.2.2 Resolution of Global Symbols in a Multisegment Task

In resolving global symbols for a multisegment task, TKB performs the same activities that it does for a single-segment task. The rules defined in Chapter 2 for resolving global symbols in a single-segment task apply also in this case, but the scope of the global symbols is altered by the overlay structure.

In a single-segment task, any module can refer to any global definition. In a multisegment task, however, a module can only refer to a global symbol that is defined on a path that passes through the called segment.

The following points, illustrated in the tree diagram in Figure 3-9, describe the two distinct cases of multiply defined symbols and ambiguously defined symbols.

In a single-segment task, if you define two global symbols with the same name, the symbols are multiply defined and an error message is produced.

In a multisegment task, you can define two global symbols with the same name if they are on separate paths, and not referenced from a segment that is common to both.

If you define a global symbol more than once on separate paths, but they are referenced from a segment that is common to both, the symbol is ambiguously defined. If you define a global symbol more than once on a single path, it is multiply defined.

OVERLAY CAPABILITY

TKB's procedure for resolving global symbols is summarized as follows:

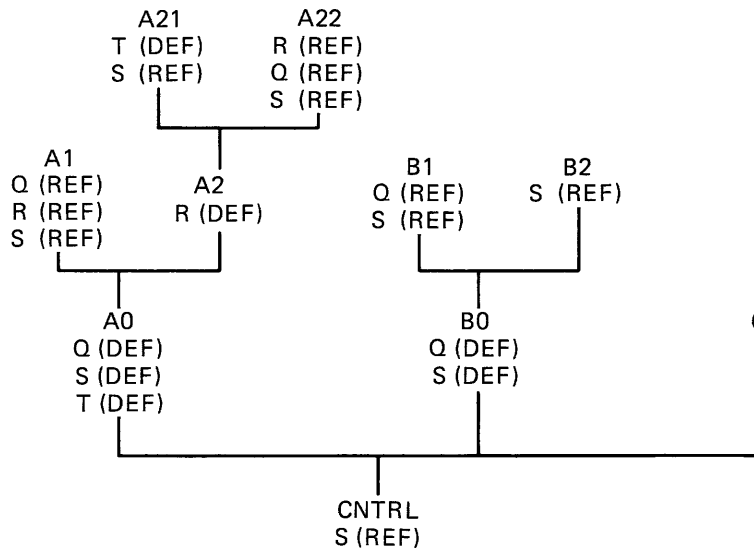
1. TKB selects an overlay segment for processing.
2. TKB scans each module in the segment for global definitions and references.
3. If the symbol is a definition, TKB searches all segments on paths that pass through the segment being processed, and looks for references that must be resolved.
4. If the symbol is a reference, TKB performs the tree search as described in step 3, looking for an existing definition.
5. If the symbol is new, TKB enters it in a list of global symbols associated with the segment.

Overlay segments are selected for processing in an order corresponding to their distance from the root. That is, TKB processes the segment farthest from the root first, before processing an adjoining segment.

When TKB processes a segment, its search for global symbols proceeds as follows:

1. The segment being processed
2. All segments toward the root
3. All segments away from the root
4. All co-trees (see Section 3.5)

Figure 3-9 illustrates the resolution of global symbols in a multisegment task.



ZK-403-81

Figure 3-9 Resolution of Global Symbols in a Multisegment Task

OVERLAY CAPABILITY

The following notes discuss the resolution of references in Figure 3-9:

1. The global symbol Q is defined in both segment A0 and segment B0. The references to Q in segment A22 and in segment A1 are resolved by the definition in A0. The reference to Q in B1 is resolved by the definition in B0. The two definitions of Q are distinct in all respects and occupy different overlay paths.
2. The global symbol R is defined in segment A2. The reference to R in A22 is resolved by the definition in A2 because there is a path to the reference from the definition (CNTRL-A0-A2-A22). The reference to R in A1, however, is undefined because there is no definition for R on a path through A1.
3. The global symbol S is defined in both segment A0 and segment B0. References to S from segments A1, A21, or A22 are resolved by the definition in A0, and references to S in B1 and B2 are resolved by the definition in B0. However, the reference to S in CNTRL cannot be resolved because there are two definitions of S on separate paths through CNTRL. The global symbol S is ambiguously defined.
4. The global symbol T is defined in both segment A21 and segment A0. Since there is a single path through the two definitions (CNTRL-A0-A2-A21), the global symbol T is multiply defined.

3.2.3 Resolution of Global Symbols from the Default Library

The process of resolving global symbols may require two passes over the tree structure. The global symbols discussed in the previous section are included in user-specified input modules that TKB scans in the first pass. If any undefined symbols remain, TKB initiates a second pass over the structure in an attempt to resolve such symbols by searching the default object module library (normally LB0:[1,1]SYSLIB.OLB). TKB reports any undefined symbols remaining after its second pass.

When multiple tree structures (co-trees) are defined, as described in Section 3.5, any resolution of global symbols across tree structures during a second pass can result in multiple or ambiguous definitions. In addition, such references can cause overlay segments to be inadvertently displaced from memory by the overlay loading routines, thereby causing run-time failures. To eliminate these conditions, the tree search on the second pass is restricted to:

- The segment in which the undefined reference has occurred
- All segments in the current tree that are on a path through the segment
- The root segment

When the current segment is the main root, the tree search is extended to all segments. You can unconditionally extend the tree search to all segments by including the /FU (full) switch in the task image file specification. (Refer to Chapter 10 for a description of the /FU switch.)

OVERLAY CAPABILITY

3.2.4 Allocation of Program Sections in a Multisegment Task

One of a program section's attributes indicates whether the program section is local (LCL) to the segment in which it is defined or is global (GBL).

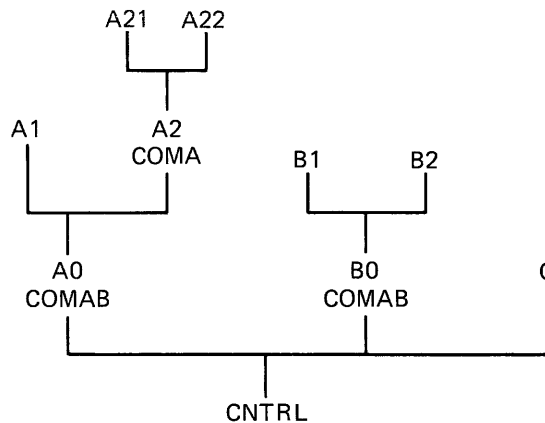
Local program sections with the same name can appear in any number of segments. TKB allocates virtual address space for each local program section in the segment in which it is declared. Global program sections that have the same name, however, must be resolved by TKB.

When a global program section is defined in several overlay segments along a common path, TKB allocates all virtual address space for the program section in the overlay segment closest to the root.

FORTRAN common blocks are translated into global program sections with the overlay (OVR) attribute. In Figure 3-10, the common block COMA is defined in modules A2 and A21. TKB allocates the virtual address space for COMA in A2 because that segment is closer to the root than the segment that contains A21.

If the segments A0 and B0 use the common block COMAB, however, TKB allocates the virtual address space for COMAB in both the segment that contains A0 and the segment that contains B0. A0 and B0 cannot communicate through COMAB. When the overlay segment containing B0 is loaded, any data stored in COMAB by A0 is lost.

You can specify the allocation of program sections explicitly. If A0 and B0 need to share the contents of COMAB, you can force the allocation of this program section into the root segment by the use of the .PSECT directive of the Task Builder's overlay description language, described in Section 3.4.



ZK-404-81

Figure 3-10 Resolution of Program Sections for TK1

3.3 OVERLAY DATA STRUCTURES AND RUN-TIME ROUTINES

When TKB constructs an overlaid task, it builds additional data structures and adds them to the task image. The data structures contain information about the overlay segments and describe the

OVERLAY CAPABILITY

relationship of each segment in the tree to the other segments in the tree. TKB also includes into the task image a number of system library routines (called overlay run-time routines). The overlay run-time routines use the data structures to facilitate the loading of the segments and to provide the necessary linkages from one segment to another at run time.

TKB links the majority of data structures and all of the overlay run-time routines into the root segment of the task. The number and type of data structures, and the functions the routines perform, depend on two considerations:

- Whether the task is built to use the Task Builder's autoloader or manual load facilities
- Whether the overlay segment is memory resident or disk resident

These considerations have a marked impact on the size and operation of the task. Chapter 4 describes the Task Builder's autoloader and manual load facilities and describes the methods for loading overlays. Appendix B describes the data structures and their contents in detail.

The contents of the root segment for a task with an overlay structure are discussed briefly in the following sections.

3.3.1 Overlaid Conventional Task Structures

Depending on the considerations just discussed, some or all of the following data structures are required by the overlay run-time routines:

- Segment tables
- Autoloader vectors
- Window descriptors
- Region descriptors

Figure 3-11 shows a typical overlay root segment structure.

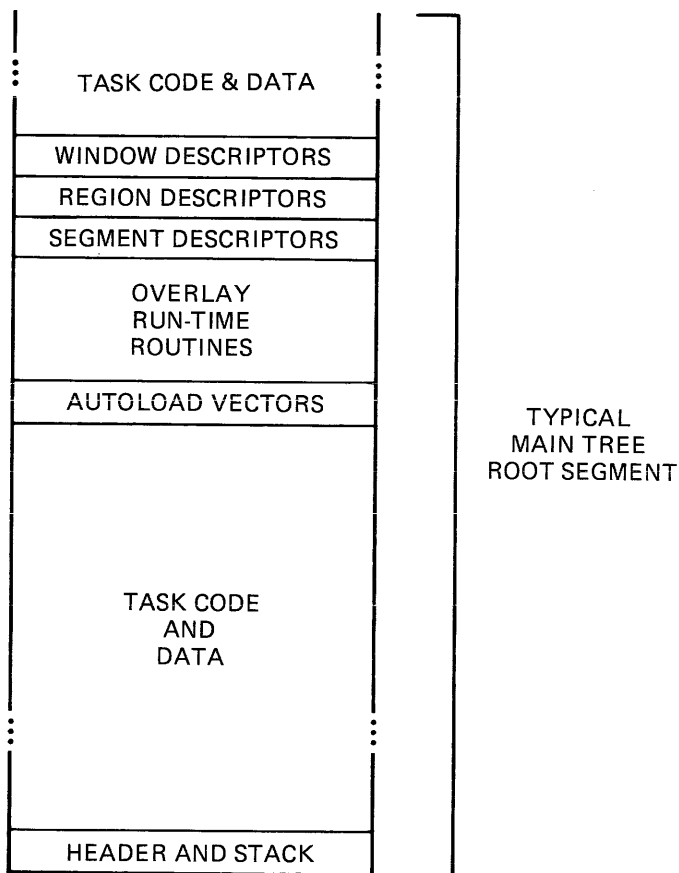
There is a segment descriptor for every segment in the task. The descriptor contains information about the load address, the length of the segment, and the tree linkage.

In an autoloader, overlaid task, autoloader vectors appear in the root segment and in every segment that calls modules in another segment located farther away from the root of the tree. All references to resident libraries are resolved through autoloader vectors in the root.

Window descriptors are allocated whenever a memory-resident overlay structure is defined for the task. The descriptor contains information required by the Create Address Window system directive (CRAW\$). One descriptor is allocated for each memory-resident overlay segment.

Region descriptors are allocated whenever a task is linked to a shared region containing memory-resident overlays. The descriptor contains information required by the Attach Region system directive (ATRG\$).

OVERLAY CAPABILITY



ZK-405-81

Figure 3-11 Typical Overlay Root Segment Structure

3.3.2 Overlaid I- and D-Space Task Structures

Overlaid I- and D-space tasks contain data structures similar to those in a conventional task. These data structures differ only in their virtual address space allocation (mapping), and some structures are mapped in two different address spaces.

Figure 3-12 shows a typical overlaid I- and D-space task with an up-tree segment.

The structures located in data space are:

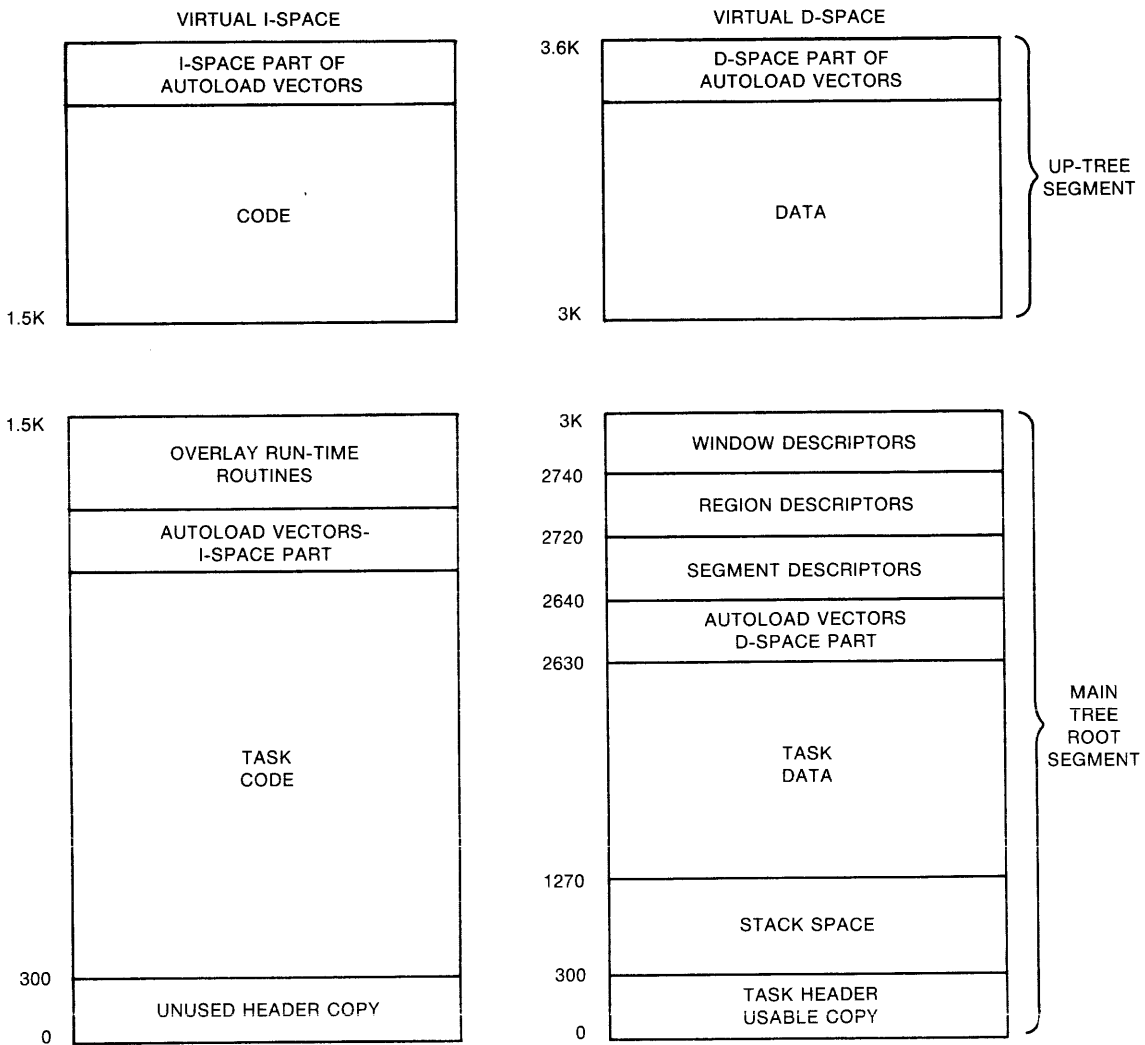
- Segment descriptors
- Window descriptors
- Region descriptors
- Autoload vectors (the data part)

Autoload vectors contain data and instructions. Therefore, TKB locates the instruction part of the autoload vector in I-space and the data part in D-space. Each segment of an autoloadable overlaid I- and D-space task may have an instruction part and a data part. Therefore, each I- and D-space segment in such a task would have its vectors separated into an instruction part and a data part.

OVERLAY CAPABILITY

The structures located in instruction space are:

- Autoload vectors (instruction part)
- Segment return point
- Overlay runtime system code (root segment only)



ZK-1050-82

Figure 3-12 Typical Overlaid I- and D-Space Task with Up-Tree Segment

OVERLAY CAPABILITY

3.4 OVERLAY DESCRIPTION LANGUAGE

TKB provides a language, called the Overlay Description Language (ODL), that allows you to describe the overlay structure of a task. An overlay description is a text file consisting of a series of ODL directives, one directive per line. Each line may have as many as 132 characters. You enter the name of this file in a TKB command line, and identify it as an ODL file by specifying the /MP switch (see Chapter 10) to the file name. For example, the following TKB command line specifies an ODL file:

```
>TKB TASK1=OVLAY/MP
```

If you specify an ODL file to TKB, it must be the only input file you specify.

A command line in an ODL file takes the form

```
label: directive argument-list ;comment
```

A label is required only for the .FCTR directive (see Section 3.4.2). Labels cannot be used with the other directives.

The ODL directives are listed below and described in Sections 3.4.1 through 3.4.6:

- .ROOT and .END
- .FCTR
- .NAME
- .PSECT
- @ (at sign; indirect command file specifier)

The ODL directives can act upon the following items: named input files, overlay segments, program sections, and lines in the ODL file itself. These items follow each directive on the same line as the directive, and form an argument-list. Operators, such as the hyphen, exclamation point, and comma, group the argument-list items (named task elements) or attach attributes to them.

If the named task element is a file, you can enter a complete file specification. Defaults for omitted parts of the file specification are as described in Chapters 1 and 10, except that the default device is SY0:, and the default UFD is taken from the terminal UIC.

In addition, the following restrictions apply to argument-lists:

- You can only use the dot character (.) in a file name.
- Comments cannot appear on a line ending with a file name.

3.4.1 .ROOT and .END Directives

The .ROOT directive defines the structure of the overlaid task. Because of this, .ROOT usually appears first in the overlay description. The .NAME directive may precede the .ROOT directive in certain circumstances discussed in Section 3.4.4. Each overlay description must end with one .END directive. The .ROOT directive tells TKB where to start building the tree, and the .END directive tells TKB where the input ends.

OVERLAY CAPABILITY

The arguments of the .ROOT directive use three operators to express concatenation, memory residency, and overlaying. These operators can be used also in the .FCTR directive.

- The hyphen (-) operator indicates the concatenation of virtual address space. For example, X-Y means that sufficient virtual address space will be allocated to contain module X and module Y simultaneously. TKB allocates segment X and segment Y in sequence to produce one segment.
- The exclamation point (!) operator indicates memory residency of overlays. (This operator is discussed in Section 3.4.3.)
- The comma (,) operator, appearing within parentheses, indicates the overlaying of virtual address space. For example, (Y,Z) means that virtual address space can contain either segment Y or segment Z. If no exclamation point (!) precedes the left parenthesis, segment Y and segment Z also share physical memory.


The comma (,) operator is also used to define multiple tree structures (as described in Section 3.5.1).

You use parentheses to delimit a group of segments that start at the same virtual address. The number of nested parenthetical groups cannot exceed 16.

For example:

```
.ROOT X-(Y,Z-(Z1,Z2))
.END
```

These directives describe the tree and its corresponding virtual address space shown in Figure 3-13:



ZK-406-81

Figure 3-13 Tree and Virtual Address Space Diagram

To create the overlay description for the task TK1 in Figure 3-3 (Section 3.1.1), you could create a file called TFIL.ODL that contains the directives:

```
.ROOT CNTRL-(A0-(A1,A2-(A21,A22)),B0-(B1,B2),C)
.END
```

To build the task with that overlay structure, you would type:

```
>TKB TK1=TFIL/MP
```

The /MP switch in the command string above tells TKB that there is only one input file (TFIL.ODL), and that this file contains the overlay description for the task.

OVERLAY CAPABILITY

3.4.2 .FCTR Directive

The .FCTR directive allows you to build large, complex trees and represent them clearly.

The .FCTR directive has a label at the beginning of the ODL line that is pointed to by a reference in a .ROOT or another .FCTR statement. The label must be unique with respect to module names and other labels. The .FCTR directive allows you to extend the tree description beyond a single line, enabling you to provide a clearer description of the overlay. (There can be only one .ROOT directive.)

For example, to simplify the tree given in the file TFIL (described in Section 3.4.1), you could use the .FCTR directive in the overlay description as follows:

```
                .ROOT CNTRL-(AFCTR,BFCTR,C)
AFCTR:         .FCTR A0-(A1,A2-(A21,A22))
BFCTR:         .FCTR B0-(B1,B2)
                .END
```

The label BFCTR is used in the .ROOT directive to designate the argument B0-(B1,B2) of the .FCTR directive. The resulting overlay description is easier to interpret than the original description. The tree consists of a root, CNTRL, and three main branches. Two of the main branches have sub-branches.

The .FCTR directive can be nested to a level of 16. For example, you could further modify TFIL as follows:

```
                .ROOT CNTRL-(AFCTR,BFCTR,C)
AFCTR:         .FCTR A0-(A1,A2FCTR)
A2FCTR:        .FCTR A2-(A21,A22)
BFCTR:         .FCTR B0-(B1,B2)
                .END
```

3.4.3 Arguments for the .FCTR and .ROOT Directives

The arguments for the .FCTR and .ROOT directives may have different forms or syntax. The examples in this chapter use forms such as A1, B1, X, and Y for clarity, but the actual arguments that you use may have somewhat different names. This section lists the forms that the arguments may take for these directives. If you use an argument that does not fall into one of the following five categories, TKB takes the argument as that of the name of an object module file; in other words, the file name that you use must have an extension of .OBJ.

3.4.3.1 Named Input File - You may use a named input file that has the object file format. For example,

```
CALC:  .FCTR [7,54]MULT.OBJ
```

The default is .OBJ.

OVERLAY CAPABILITY

3.4.3.2 Specific Library Modules - You may name and therefore use specific object modules from a library file. For example,

```
BAKER: .FCTR [300,3]COOKIE/LB:CHIP:OAT
```

where COOKIE.OLB is the library file and CHIP and OAT are the modules that you want to extract from the file. The default extension is .OLB and it need not be specified as part of the argument.

3.4.3.3 A Library to Resolve References Not Previously Resolved - You may specify a library as an argument in a .FCTR statement after extracting specific modules in a previous .FCTR statement. TKB uses the library to resolve symbols that may still be unresolved after extracting the modules. For example,

```
BAKER: .FCTR [300,3]COOKIE/LB:CHIP:OAT
LIB: .FCTR LB:[1,4]RECIPE/LB
```

3.4.3.4 A Section Name Used in a .PSECT Directive - You may use the name that you used as a program section name in the .PSECT directive as the argument in a .FCTR statement. For example,

```
.PSECT COM,GBL,D,RW,OVR
FSTCOM: .FCTR COM
```

3.4.3.5 A Segment Name Used in a .NAME Directive - You may use the name that you specified as the name of a segment in the .NAME directive. For example,

```
.NAME SEG1,GBL,DSK
OVLY: .FCTR SEG1-MOD1-MOD2
```

3.4.4 Exclamation Point Operator

The exclamation point operator allows you to specify memory-resident overlay segments (see Section 3.1.2). You specify memory residency by placing an exclamation point (!) immediately before the left parenthesis enclosing the segments to be affected. The overlay description for task TK2 in Figure 3-4 (Section 3.1.2) is as follows:

```
.ROOT CNTRL-!(D,E,F)
.END
```

In the example above, segments D, E, F are declared resident in separate areas of physical memory. The Task Builder determines the single starting virtual address for D, E, and F by rounding the octal length of segment CNTRL up to the next 4K boundary. The physical memory allocated to segments D, E, and F is determined by rounding the actual length of each segment to the next 32-word boundary (256-word boundary if the /CM switch is in effect), and adding this value to the total memory required by the task.

OVERLAY CAPABILITY

The exclamation point operator applies to that segment immediately to the right of the left parenthesis and those segments farther from the root on the same level with that segment. In other words, all parallel segments must be of the same residency type (disk resident or memory resident).

The exclamation point operator applies to segments at the same level from the root inside a pair of parentheses; segments nested in parenthesis within that level, but farther from the root, are not affected.

It is therefore possible to define an overlay structure that combines the space-saving attributes of disk-resident overlays with the speed of memory-resident overlays. For example:

```
.ROOT A-!(B1-(B2,B3),C)
.END
```

In this example, B1 and C are declared memory resident by the exclamation point operator. B2 and B3 are declared disk resident, however, because no exclamation point operator precedes the parentheses enclosing them.

Note that while a memory-resident overlay can call a disk-resident overlay, the converse is not legal; that is, you cannot use an exclamation point for segments emanating from a disk-resident segment. For example, you cannot build the following structure:

```
.ROOT A-(B1-!(B2,B3),C) ; this overlay description is illegal
.END
```

In this example, B1 is declared disk resident; so it is illegal to use the exclamation point to declare B2 and B3 memory resident.

3.4.5 .NAME Directive

The .NAME directive allows you to name a segment, and assign attributes to the segment. The name must be unique with respect to file names, program section names, .FCTR labels, and other segment names used in the overlay description. You use the .NAME directive prior to using the .ROOT or .FCTR directive. The Task Builder attaches attributes to a segment when it encounters the name in a .ROOT or .FCTR directive that defines the overlay segment. If you apply multiple names to a segment, the attributes of the last name given are in effect. This directive does the following:

- Names uniquely a segment that is loaded through the manual load facility (see Chapter 4)
- Permits a named data-only segment to be loaded through the autoloading mechanism

The format of the .NAME directive is:

```
.NAME segname[,attr][,attr]
```

segname

A 1- to 6-character name; this name can consist of the Radix-50 characters A-Z, 0-9, and \$ (the period (.) cannot be used).

OVERLAY CAPABILITY

attr

One of the following:

GBL The name is entered in the segment's global symbol table.

The GBL attribute makes it possible to load data-only overlay segments by means of the autoloading mechanism (see Chapter 4).

NODSK No disk space is allocated to the named segment.

If a data overlay segment has no initial values, but will have its contents established by the running task, no space for the named segment on disk need be reserved. If the code attempts to establish initial values for data in a segment for which no disk space is allocated (a segment with the NODSK attribute), TKB gives a fatal error.

NOGBL The name is not entered in the segment's global symbol table.

If the GBL attribute is not present, NOGBL is assumed.

DSK Disk storage is allocated to the named segment.

If the NODSK attribute is not present, DSK is assumed.

3.4.5.1 Example of The Use of The .NAME Directive - In the following modified ODL file for TK1 (Figure 3-3 of Section 3.1.1), you provide names for the three main branches, A0, B0, and C, by specifying the names in the .NAME directive and using them in the .ROOT directive. The default attributes NOGBL and DSK are in effect for BRNCH1 and BRNCH3, but BRNCH2 has the complementary attributes (GBL and NODSK) that cause TKB to enter the name BRNCH2 into the segment's global symbol table and suppress disk allocation for that segment. BRNCH2 contains uninitialized storage to be utilized at run time.

```
          .NAME BRNCH1
          .NAME BRNCH2,GBL,NODSK
          .NAME BRNCH3
          .ROOT CNTRL-!(BRNCH1-AFCTR,*BRNCH2-BFCTR,BRNCH3-C)
AFCTR:    .FCTR A0-(A1,A2-(A21,A22))
BFCTR:    .FCTR B0-*(B1,B2)
          .END
```

(The asterisk (*) is the autoloading indicator; it is discussed in Chapter 4.)

You can load the data overlay segment BRNCH2 by including the following statement in the program:

```
CALL BRNCH2
```

This action is immediately followed by an automatic return to the next instruction in the program.

OVERLAY CAPABILITY

You can also use segment names in making patches with the `ABSPAT` and `GBLPAT` options (see Chapter 11).

NOTE

In the absence of a unique `.NAME` specification, `TKB` establishes a segment name, using the first module name or library module name occurring in the segment.

3.4.6 .PSECT Directive

You can use the `.PSECT` directive to control the placement of a global program section in an overlay structure. The name of the program section (a 1- to 6-character name consisting of the Radix-50 characters A-Z, 0-9, and \$) and its attributes are given in the `.PSECT` directive. The attributes used in the `.PSECT` directive must match those in the actual program section in the module. Thus, you can use the name in a `.ROOT` or `.FCTR` statement to indicate to the Task Builder the segment to which the program section will be allocated. An example of the use of `.PSECT` is given in the modified version of task `TK1` (the original version is shown in Figure 3-3 in Section 3.1.1) shown below.

In this example, `TK1` has a disk-resident overlay structure. The example assumes that the programmer was careful about the logical independence of the modules in the overlay segment, but failed to take into account the requirement for logical independence in multiple executions of the same overlay segment.

The flow of task `TK1` can be summarized as follows. `CNTRL` calls each of the overlay segments, and the overlay segment returns to `CNTRL` in the order A, B, C, A. Module A is executed twice. The overlay segment containing A must be reloaded for the second execution.

Module A uses a common block named `DATA3`. The Task Builder allocates `DATA3` to the overlay segment containing A. The first execution of A stores some results in `DATA3`. The second execution of A requires these values. In this disk-resident overlay structure, however, the values calculated by the first execution of A are overlaid. When the segment containing A is read in for the second execution, the common block is in its initial state.

To permit the two executions of A to communicate, a `.PSECT` directive is used to force the allocation of `DATA3` into the root. The indirect command file for `TK1`, `TFIL.ODL`, is modified as follows:

```
                .PSECT DATA3,RW,GBL,REL,OVR
                .ROOT CNTRL-DATA3-(AFCTR,BFCTR,C)
AFCTR:         .FCTR A0-(A1,A2-(A21,A22))
BFCTR:         .FCTR B0-(B1,B2)
                .END
```

The attributes `RW`, `GBL`, `REL`, and `OVR` are described in Chapter 2.

OVERLAY CAPABILITY

3.4.7 Indirect Command Files

The Overlay Description Language processor can accept ODL text indirectly, that is, specified in an indirect command file. If an at sign (@) appears as the first character in an ODL line, the processor reads text from the file specified immediately after the at sign. The processor accepts the ODL text from the file as input at the point in the overlay description where the file is specified.

For example, suppose you create a file, called BIND.ODL, that contains the text:

```
B: .FCTR B1-(B2,B3)
```

A line beginning with @BIND can replace this text at the position where the text would have appeared:

	Indirect		Direct
	.ROOT A-(B,C)		.ROOT A-(B,C)
C:	.FCTR C1-(C2,C3)	C:	.FCTR C1-(C2,C3)
@BIND		B:	.FCTR B1-(B2,B3)
	.END		.END

The Task Builder allows two levels of indirection.

3.5 MULTIPLE-TREE STRUCTURES

You can define more than one tree within an overlay structure. These multiple tree structures consist of a main tree and one or more co-trees. The root segment of the main tree is loaded by the Executive when the task is made active, while segments within each co-tree are loaded through calls to the overlay run-time routines. Except for this distinction, all overlay trees have identical characteristics: a root segment that resides in memory, and two or more overlay segments.

The main property of a structure containing more than one tree is that storage is not shared among trees. Any segment in a tree can be referred to from another tree without displacing segments from the calling tree. Routines that are called from several main tree overlay segments, for example, can overlay one another in a co-tree. The same considerations in deciding whether to create memory-resident overlays or disk-resident overlays in a single-tree structure apply in building a structure containing co-trees.

3.5.1 Defining a Multiple-Tree Structure

Multiple-tree structures are specified within the Overlay Description Language by extending the function of the comma operator. As described in Section 3.4, this operator, when included within parentheses, defines a pair of segments that share storage. Including the comma operator outside all parentheses delimits overlay trees. The first overlay tree thus defined is the main tree. Subsequent trees are co-trees. For example:

```
.ROOT      X,Y
X:         .FCTR      X0-(X1,X2,X3)
Y:         .FCTR      Y0-(Y1,Y2)
          .END
```

OVERLAY CAPABILITY

In this example, two overlay trees are specified: 1) a main tree containing the root segment X0 and three overlay segments; and 2) a co-tree consisting of root segment Y0 and two overlay segments. The Executive loads segment X0 into memory when the task is activated. The task then loads the remaining segments through calls to the overlay run-time routines.

3.5.1.1 Defining Co-trees With a Null Root by Using .NAME - A co-tree must have a root segment to establish linkage with its own overlay segments. However, co-tree root segments need not contain code or data and, therefore, can be 0 length. You can create a segment of this type, called a null segment, by means of the .NAME directive. The previous example is modified, as shown below, to move file Y0.OBJ to the root and include a null segment.

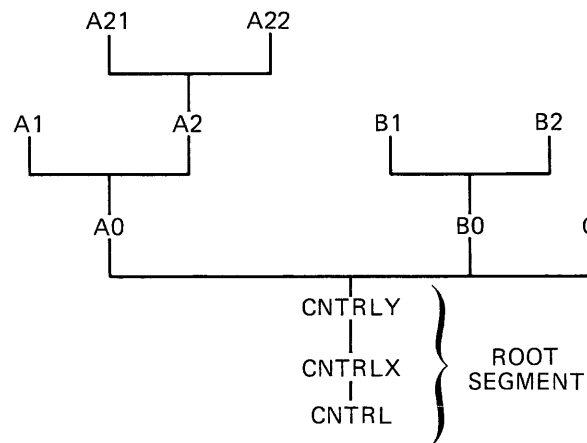
```
X:      .ROOT      X,Y
        .FCTR      X0-Y0-(X1,X2,X3)
        .NAME      YNUL
Y:      .FCTR      YNUL-(Y1,Y2)
        .END
```

The null segment YNUL is created by using the .NAME directive, and replaces the co-tree root that formerly contained Y0.OBJ.

3.5.2 Multiple-Tree Example

The following example illustrates the use of multiple trees to reduce the size of the task.

In this example, the root segment CNTRL of task TK1 (described in Section 3.1.1) has had two routines added to it: CNTRLX and CNTRLY. The routines are logically independent of each other, and both are approximately 4000(octal) bytes long. However, the routines have been placed in the root segment of TK1 instead of being overlaid because both routines must be accessed from modules on all paths of the tree. In a single-tree overlay structure, the root segment is the only segment common to all paths of the tree. The schematic diagram for the modified structure is shown in Figure 3-14.



ZK-407-81

Figure 3-14 Overlay Tree for Modified TK1

OVERLAY CAPABILITY

One possible overlay description for this structure is shown below:

```
.ROOT CNTRL-CNTRLX-CNTRLY-(AFCTR,BFCTR,C)
AFCTR: .FCTR A0-(A1,A2FCTR)
A2FCTR: .FCTR A2-(A21,A22)
BFCTR: .FCTR B0-(B1,B2)
.END
```

Because TK1 consists of disk-resident overlays and the new routines are concatenated within the overlay structure, the new routines add 10000(octal) bytes to both the virtual address space and physical memory requirements of the task. However, the added routines consume more virtual address space than might be expected, as shown in Figure 3-15.

The expansion of TK1's virtual address space requirements caused the task to extend 4000(octal) bytes beyond the next highest 4K-word boundary (APR 2). Because the Executive must use an additional mapping register (APR2), the apparent cost in virtual address space above APR 2 of 4000(octal) bytes is in fact 20000(octal) bytes. (Compare the diagram in Figure 3-15 with the diagram in Figure 3-3.) The shaded portion of the unused virtual address space in Figure 3-15 represents the portion of virtual address space that is allocated but is unusable as allocated.

Small tasks, such as TK1, are seldom adversely affected by the inefficient allocation of virtual address space, but larger tasks may be. For example, a large task that contains code to create dynamic regions (see Chapter 5) or that contains Executive directives to extend its task region (see the RSX-11M/M-PLUS Executive Reference Manual) requires at least 4K words of virtual address space to map each region. In such a task, using co-trees can often save virtual address space and can, therefore, be of paramount importance. TK1 can be modified to reflect this.

As noted earlier, the routines CNTRLX and CNTRLY are logically independent. Logical independence is a primary requirement for all segments that overlay each other. However, CNTRLX and CNTRLY cannot be structured into either of the main branches of TK1's tree because it is further required that the routines be accessible from modules on all paths of the tree. Therefore, the only way CNTRLX and CNTRLY can be overlaid and still meet all of these requirements is through a co-tree structure. Figure 3-16 shows the schematic representation of TK1 as a co-tree structure.

The root segment CNTRL2 of the co-tree is a null segment. It contains no code or data and has a length of 0. As noted earlier, the Task Builder requires the root segment in order to establish linkage with the overlay segments. One possible overlay description for building TK1 as a 2-tree structure is shown below.

```
.NAME CNTRL2
.ROOT CNTRL-(AFCTR,BFCTR,C),CNTRL2-(CNTRLX,CNTRLY)
AFCTR: .FCTR A0-(A1,A2FCTR)
A2FCTR: .FCTR A2-(A21,A22)
BFCTR: .FCTR B0-(B1,B2)
.END
```

You define the co-tree in the .ROOT directive by placing the comma operator outside all parentheses (immediately before CNTRL2). The .NAME directive creates the null root segment. Figure 3-16 shows the new relationship between virtual address space and physical memory.

OVERLAY CAPABILITY

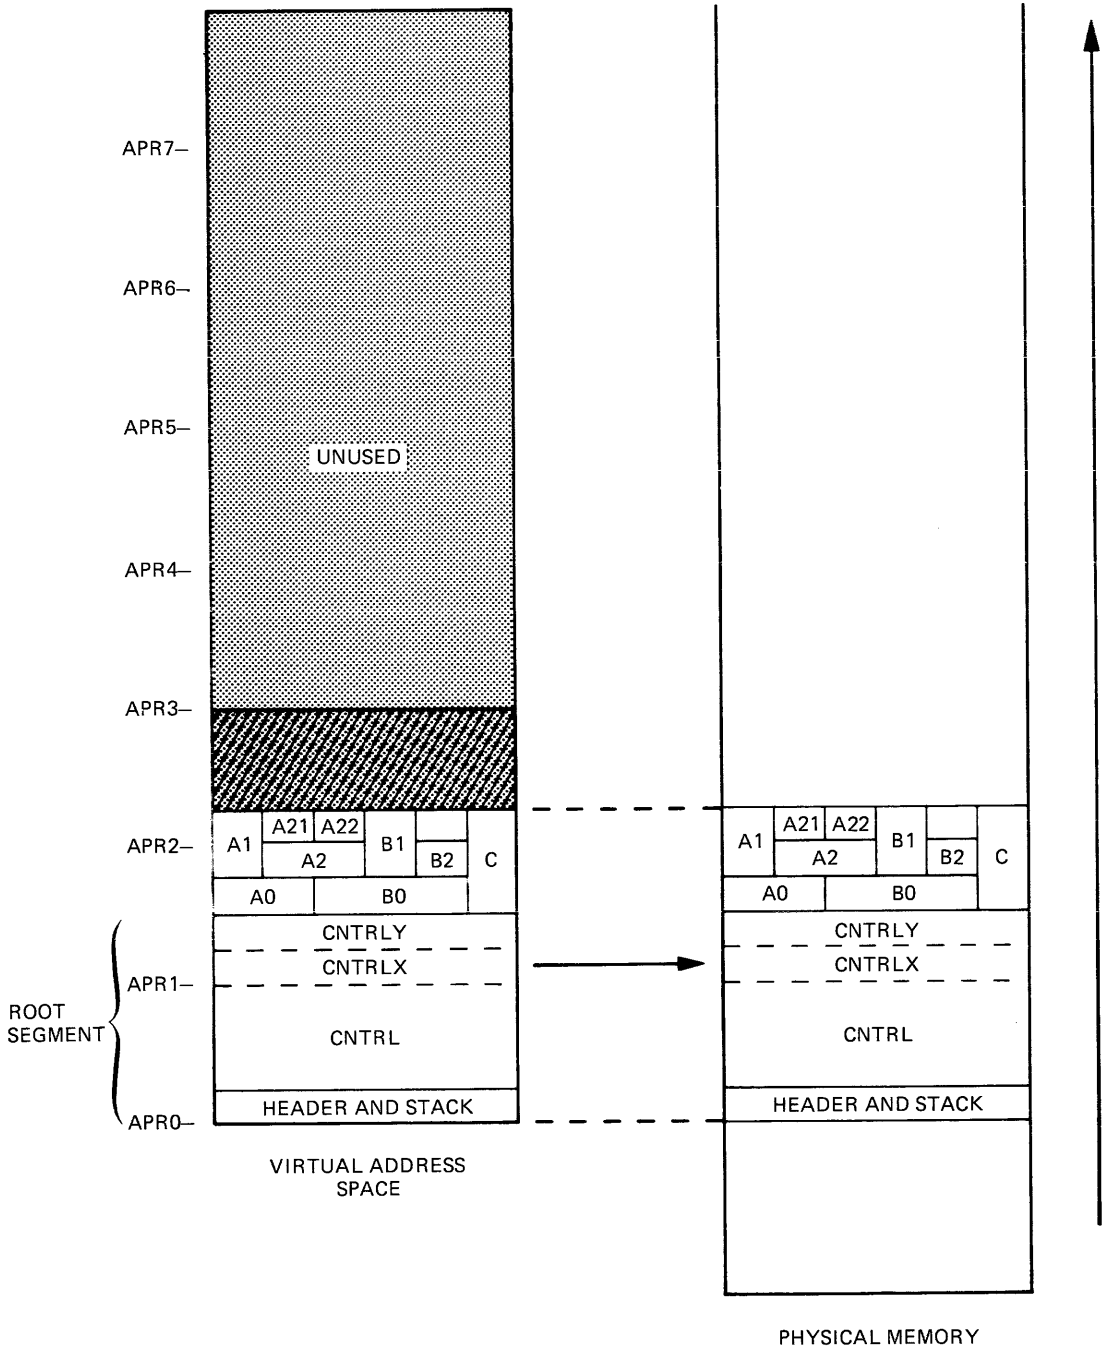
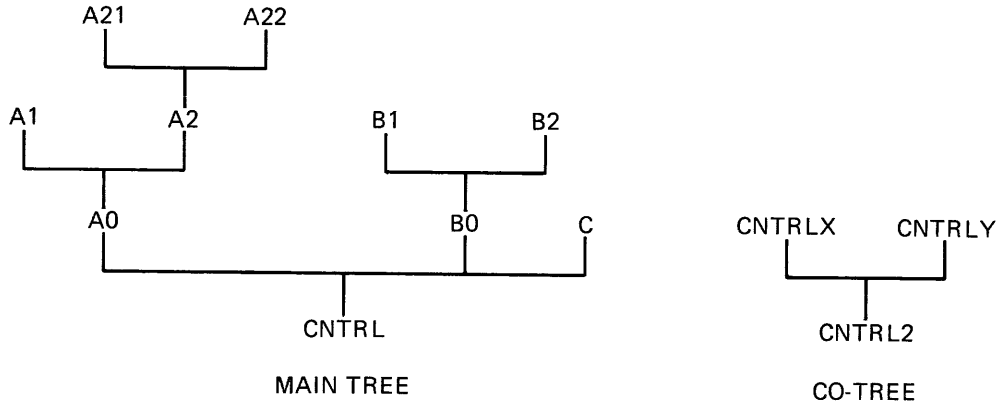


Figure 3-15 Virtual Address Space and Physical Memory for Modified TK1

ZK-408-81

OVERLAY CAPABILITY



ZK-409-81

Figure 3-16 Overlay Co-Tree for Modified TK1

The diagrams in Figure 3-17 illustrate the savings (4000(octal) bytes) in both virtual address space and physical memory that is realized by overlaying CNTRLX and CNTRLY. What may be more important in some applications, however, is that the top of TK1's task region has dropped below the 4K-word boundary of APR 2. TK1 has gained 4K words of potentially usable virtual address space.


NOTE

The numbers used in this example have been simplified for illustrative purposes. In addition, the storage required for overhead in handling the overlay structures is not reflected in this example.

Because the null root CNTRL2 is 0 bytes long, it does not require any virtual address space or physical memory and, therefore, does not appear in the diagrams in Figure 3-17.

Finally, you can define any number of co-trees. Additional co-trees can access all modules in the main tree and other co-trees.

OVERLAY CAPABILITY



ZK-410-81


Figure 3-17 Virtual Address Space and Physical Memory for TK1 As a Co-Tree

3.6 CREATING AN ODL FILE FROM A VIRTUAL ADDRESS SPACE ALLOCATION DIAGRAM

You can use a graphic method as an aid to converting a virtual address space allocation diagram into the correct Task Builder ODL file.

OVERLAY CAPABILITY

First create a virtual address space allocation diagram of your overlaid task, similar to that shown in Figure 3-18, without the dotted-line path shown in the diagram.



ZK-1052-82

Figure 3-18 Virtual Address Space Allocation Diagram

The dotted-line path will be the basis for writing the ODL statements that you need. To determine the path through your virtual address space allocation diagram, follow these steps:

1. Start in the lower left corner of the root segment.
2. Draw a dotted line upward as far as you can go without passing through the top or into "empty" virtual space, crossing into new segments as needed.
3. When you reach the top segment, proceed to the right until you reach a vertical line.
4. If the end of your dotted line is now opposite the vertical line of the lowest segment, cross the vertical line and continue again from step 2; otherwise, proceed to step 5.
5. Because the end of your dotted line is not opposite the vertical line of the lowest segment proceed downward until you reach the lowest segment.
6. If you are not in the root, cross the vertical line to the right and continue from step 2; otherwise, proceed to step 7.
7. If your dotted line is in the lower right corner of the root, you have finished the dotted-line walk.

OVERLAY CAPABILITY

Once you have drawn the dotted line, you should go back over it to verify that you followed all the steps, while doing this, draw arrowheads at each point where a line was crossed to indicate the direction of the line.

3.6.1 Creating a .ROOT Statement by Using a Virtual Address Space Allocation Diagram

Now you are ready to write the .ROOT statement. Follow these steps:

1. Write .ROOT followed by the name of the root statement (in this example, .ROOT CNTRL).
2. Follow the dotted-line path.
3. Add each successive ODL element to your root statement, using the following syntax, based on the direction of your dotted line.
 - A. At an upward crossing: -("name of new segment"
 - B. At a horizontal crossing: ,"name of new segment"
 - C. At a downward crossing:)
4. When you have returned to the root, your root statement is complete.

Using the dotted-line path in Figure 3-18 and the above associated steps for writing the .ROOT statement, you can write as shown below:

1. Step 1 : Write .ROOT CNTRL
2. Step 3A: Write .ROOT CNTRL-(A0
3. Step 3A: Write .ROOT CNTRL-(A0-(A1
4. Step 3B: Write .ROOT CNTRL-(A0-(A1,A2
5. Step 3A: Write .ROOT CNTRL-(A0-(A1,A2-(A21
6. Step 3B: Write .ROOT CNTRL-(A0-(A1,A2-(A21,A22
7. Step 3C: Write .ROOT CNTRL-(A0-(A1,A2-(A21,A22)
8. Step 3C: Write .ROOT CNTRL-(A0-(A1,A2-(A21,A22))
9. Step 3B: Write .ROOT CNTRL-(A0-(A1,A2-(A21,A22)),B0
10. Step 3A: Write .ROOT CNTRL-(A0-(A1,A2-(A21,A22))-B0-(B1
11. Step 3B: Write .ROOT CNTRL-(A0-(A1,A2-(A21,A22))-B0-(B1,B2
12. Step 3C: Write .ROOT CNTRL-(A0-(A1,A2-(A21,A22))-B0-(B1,B2)
13. Step 3B: Write .ROOT CNTRL-(A0-(A1,A2-(A21,A22))-B0-(B1,B2),C
14. Step 3C: Write .ROOT CNTRL-(A0-(A1,A2-(A21,A22))-B0-(B1,B2),C)

The steps for writing .FCTR statements and co-tree statements follow next.

OVERLAY CAPABILITY

3.6.2 Creating a .FCTR Statement by Using a Virtual Address Space Allocation Diagram

By using the steps for creating a .ROOT statement from a virtual address space allocation diagram, you created the following .ROOT statement.

```
.ROOT CNTRL-(A0-(A1,A2-(A21,A22))-B0-(B1,B2),C)
```

It may be desirable to simplify your specific .ROOT statement into one or more .FCTR statements. A technique similar to the one used to create the .ROOT statement may be used to create the .FCTR statement.

In this example, segments A0, A1, A2, A21, and A22 are selected to be in the .FCTR statement. Having selected these segments (normally related as a "stack" of segments) you are now ready to write down the .FCTR statement.

First, draw a virtual address space allocation diagram of the segments that you have selected. (You may use Figure 3-18 for this explanation.) Then follow these next steps to draw a dotted-line path through the diagram:

1. Start in the lower left corner of the lowest or "base" segment (A0) in your diagram.
2. Draw a dotted line upward as far as you can go without passing through the top or into empty virtual space, crossing into new segments as needed.
3. When you reach the top segment, proceed to the right until you reach a vertical line.
4. If the end of your dotted line is now opposite the vertical line of the lowest segment, cross the vertical line and continue again from step 2; otherwise, proceed to step 5.
5. Because the end of your dotted line is not opposite the vertical line of the lowest segment, proceed downward until you reach the lowest segment.
6. If you are not in the base segment (A0), cross the vertical line to the right and continue from step 2; otherwise, proceed to step 7.
7. If your dotted line is in the lower right corner of the base segment, you have finished the dotted-line walk.

Once you have drawn the dotted line, you should go back over it to verify that you followed all the steps. While doing this, draw arrowheads at each point where a line was crossed to indicate the direction of the line.

Now you are ready to write the .FCTR statement. Follow these next steps:

1. Write a label followed by .FCTR, which is in turn followed by the name of the first segment (A0) (in this example, AFCTR .FCTR A0)
2. Follow the dotted-line path.

OVERLAY CAPABILITY

3. Add each successive ODL element to your root statement, using the following syntax, based on the direction of your dotted line.
 - A. At an upward crossing: ("name of new segment"
 - B. At a horizontal crossing: ,"name of new segment"
 - C. At a downward crossing:)
4. When you have returned to the base segment, your .FCTR statement is complete.

Using the dotted line path and the above associated steps for writing the .FCTR statement, you can write as shown below:

1. Step 1 : Write AFCTR .FCTR A0
2. Step 3A: Write AFCTR .FCTR A0-(A1
3. Step 3B: Write AFCTR .FCTR A0-(A1,A2
4. Step 3A: Write AFCTR .FCTR A0-(A1,A2-(A21
5. Step 3B: Write AFCTR .FCTR A0-(A1,A2-(A21,A22
6. Step 3C: Write AFCTR .FCTR A0-(A1,A2-(A21,A22)
7. Step 3C: Write AFCTR .FCTR A0-(A1,A2-(A21,A22))

You have now reached the base segment and have written the two ODL statements:

```
.ROOT CNTRL-(A0-(A1,A2-(A21,A22))-B0-(B1,B2),C)
AFCTR: .FCTR A0-(A1,A2-(A21,A22))
```

The last step requires that you substitute the label, AFCTR, into the .ROOT statement, which results in:

```
.ROOT CNTRL-AFCTR-B0-(B1,B2),C)
AFCTR: .FCTR A0-(A1,A2-(A21,A22))
```

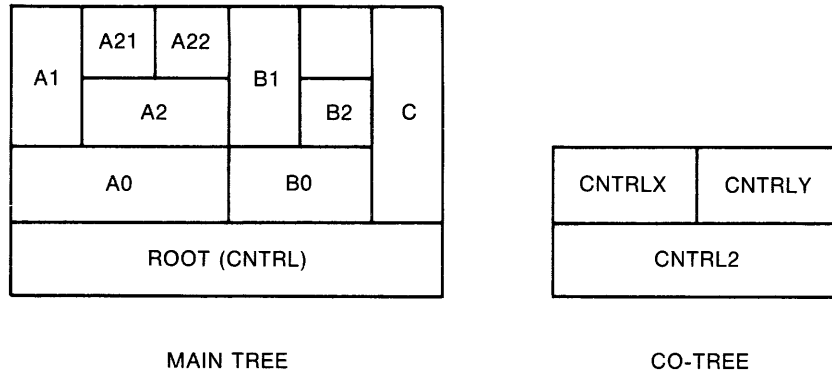
Additional .FCTR statements would be determined and written in the same way. For example, you could write a .FCTR statement labeled BFCTR for the segments B0, B1, and B2.

The following section shows how to write an ODL statement for a co-tree by using the same methods.

3.6.3 Creating an ODL Statement for a Co-Tree by Using a Virtual Address Space Diagram

Assuming that you want to write an ODL statement for a co-tree like the one in Figure 3-19, you would have two virtual address space allocation diagrams, one for the main tree and one for the co-tree. These two diagrams are shown in Figure 3-19.

OVERLAY CAPABILITY



ZK-1051-82

Figure 3-19 Virtual Address Space Allocation for a Main Tree and Its Co-Tree

From Figure 3-19 you see that the co-tree is a stack of segments also. Therefore, it is possible to write the statement for the co-tree in the same fashion and with the same rules as that described in Section 3.6. However, certain facts must be kept in mind. These are that:

- The co-tree has a null root
- A .NAME statement must be used to name the null root
- A comma must be placed outside of the parentheses and at the end of that part of the .ROOT statement that defines the main tree

Therefore, the ODL statement that we obtain before writing the co-tree part is:

```

        .NAME CNTRL2
        .ROOT CNTRL-AFCTR-B0-(B1,B2),C),
AFCTR: .FCTR A0-(A1,A2-(A21,A22))
    
```

By following the rules in Section 3.6 and by using the diagram in Figure 3-19, you can then create the ODL statement:

```

        .NAME CNTRL2
        .ROOT CNTRL-AFCTR-B0-(B1,B2),CNTRL2-(CNTRLX,CNTRLY)
AFCTR: .FCTR A0-(A1,A2-(A21,A22))
    
```

3.7 OVERLAYING PROGRAMS WRITTEN IN A HIGH-LEVEL LANGUAGE

Programs written in a high-level language usually require the use of a large number of library routines in order to execute. Unless care is taken when overlaying such programs, the following problems can occur:

- TKB throughput may be drastically reduced because of the number of library references in each overlay segment.

OVERLAY CAPABILITY

- Library references from the default object module library that are resolved across tree boundaries can result in unintentional displacement of segments from memory at run time.
- Attempts to task-build such programs can result in multiple and ambiguous symbol definitions when a co-tree structure is defined.

The following procedures are effective in solving these problems:

- You can increase TKB throughput by linking commonly used library routines into the main root segment.
- You can eliminate ambiguous definitions, multiple definitions, and cross-tree references by using the NOFU switch (the TKB default) to restrict the scope of the default library search. However, restricting the scope of the default library search may also cause problems.

If sufficient memory is available, you can effectively place the object time system in the root segment by building a memory-resident library. This also reduces total system memory requirements if other tasks are also currently using the library.

If a memory-resident library cannot be built, you can force library modules into the root by preparing a list of the appropriate global references and linking the object module into the root segment.

For other ways to reduce task size, you should consult the user's guide for the language you are using.

3.8 EXAMPLE 3-1: BUILDING AN OVERLAY

The text in this section and the figures associated with it illustrate the building of an overlay structure. For this example, the routines of the resident library LIB.TSK and the task that refers to it, MAIN.TSK (from Example 5-3, Chapter 5), are assembled as separate modules and built as an overlaid task. This task is built first with disk-resident overlays and then with memory-resident overlays. The disk-resident version of the task is named OVR.TSK and the memory-resident version is named RESOVR.TSK.

NOTE

This example is intended to provide you with a working illustration of the Overlay Description Language. It does not reflect the most efficient use of it.

Two alterations were made to each of the routines for this example:

- A .TITLE and .END assembler directive was added to each routine to establish it as a unique module.

OVERLAY CAPABILITY

- The following assembler directive was added to each arithmetic routine to increase its allocation:

```
.BLKW 1024.*3
```

This was done to make TKB allocation of address space more obvious for documentation purposes.

The operation of the overlaid task is identical to that of Example 5-3 in Chapter 5. The routines and their titles as a result of the .TITLE directives are as follows:

- The integer addition routine is named ADDOV.
- The integer subtraction routine is named SUBOV.
- The integer multiplication routine is named MULOV.
- The integer division routine is named DIVOV.
- The register save and restore routine is named SAVOV.
- The print routine is named PRNOV.
- The main calling routine is named ROOTM.

The lengths of the modules are:

Module	Length (in Octal)
ADDOV	14024 bytes
SUBOV	14024 bytes
MULOV	14024 bytes
DIVOV	14026 bytes
SAVOV	4042 bytes
PRNOV	4260 bytes
ROOTM	4104 bytes


The flow of control for OVR.TSK is as follows:

1. ROOTM calls ADDOV and ADDOV returns to ROOTM.
2. ROOTM calls PRNOV to print the result and PRNOV returns to ROOTM.
3. ROOTM calls SUBOV and SUBOV returns to ROOTM.
4. ROOTM calls PRNOV to print the result and PRNOV returns to ROOTM.
5. ROOTM calls DIVOV and DIVOV returns to ROOTM.
6. ROOTM calls PRNOV to print the result and PRNOV returns to ROOTM.

OVERLAY CAPABILITY

7. ROOTM calls MULO and MULO returns to ROOTM.
8. ROOTM calls PRNOV to print the result and PRNOV returns to ROOTM.

The print routine (contained in module PRNOV) is called between each arithmetic operation by the control routine (contained in module ROOTM). To avoid loading it into physical memory each time it is called, you can place PRNOV in the root segment of the task. In addition, each arithmetic routine calls SAVOV. Therefore, SAVOV must be on a path common to all segments in the tree. It too is placed in the root segment of the task. One possible overlay configuration for this task is shown in Figure 3-20.



ZK-490-81

Figure 3-20 Overlay Tree of Virtual Address Space for OVR.TSK

To build this overlay, first create an ODL file (OVERTREE.ODL) that contains its description:

```
.ROOT  ROOTM-PRNOV-SAVOV-*(MULO,ADDOV-(SUBOV,DIVOV))
.END
```

Then, after you have modified the modules and assembled them, you can build the task with the following command line:

```
TKB> OVR,OVR/-SP=OVRTREE/MP
```

This command instructs TKB to build a task image, OVR.TSK, and to create a map file, OVR.MAP, under the UFD that corresponds to the terminal UIC. The negated spool switch (/SP) inhibits TKB from spooling the map file to the line printer.

The overlay switch (/MP) attached to the input file tells TKB that the input file is an ODL file. Therefore, this file will be the only input file specified. Refer to Chapter 10 for a description of the switches used in this example.

A portion of the map that results from this task build is shown in Example 3-1.

OVERLAY CAPABILITY

Example 3-1 Map File for OVR.TSK

OVR.TSK Memory allocation map TKB M40.10 Page 1
 01-JAN-82 10:06

Partition name : GEN
 Identification : 01
 Task UIC : [7,62]
 Stack limits: 000260 001257 001000 00512.
 PRG xfr address: 001264
 Total address windows: 1. ②
 Task image size : 7488. words
 Task address limits: 000000 035107
 R-W disk blk limits: 000002 000073 000072 00058.

OVR.TSK Overlay description:

Base	Top	Length	
000000	005033	005034	02588. ROOTM
005034	021057	014024	06164. MULOV
005034	021057	014024	06164. ADDOV
021060	035103	014024	06164. SUBOV
021060	035107	014030	06168. DIVOV

*** Root segment: ROOTM

R/W mem limits: 000000 005033 005034 02588.
 Disk blk limits: 000002 000007 000006 00006.

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.: (RW, I, LCL, REL, CON)	001260 002514 01356.		
	001260 000102 00066. ROOTM	01	ROOTM.OBJ;1
	001362 000260 00176. PRNOV	01	PRNOV.OBJ;1
	001642 000042 00034. SAVOV	01	SAVOV.OBJ;1
ANS : (RW, D, GBL, REL, OVR)	003774 000002 00002.		
	003774 000002 00002. ROOTM	01	ROOTM.OBJ;1
	003774 000002 00002. PRNOV	01	PRNOV.OBJ;1

Global symbols:

AADD 004032-R DIVV 004052-R PRINT 001550-R SUBB 004042-R
 MULL 004022-R SAVAL 001642-R

*** Task builder statistics:

Total work file references: 6863.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 7086. words (27. pages)
 Size of work file: 3072. words (12. pages)

Elapsed time:00:00:14

OVERLAY CAPABILITY

Figure 3-21 shows the allocation of virtual address space for OVR.TSK. The circled numbers in Example 3-1 correspond to those in Figure 3-21.

Note that the root segment for OVR.TSK (ROOTM) has expanded with task building while the segments containing the arithmetic routines have not. Before task building, the sum of the modules (in octal bytes) that comprise the root segment is:

$$4104 + 4260 + 4042 = 14,426 \text{ bytes}$$

After task building, the root segment is 20,677(octal) bytes long. TKB has added a header, a stack area, and the overlay run-time routines to it. The segments containing the arithmetic routines have not changed. If there had been calls from segments nearer the root to segments farther up the tree, the Task Builder would have added data structures to the calling segments as well. (Refer to Chapter 4 for a description of the overlay loading methods.)

You can build OVR as a memory-resident overlay by simply adding the memory-resident operator (!) to the ODL file for OVR as shown below:

```
.ROOT  ROOTM-PRNOV-SAVOV-*(MULOV,ADDOV-!(SUBOV,DIVOV))
.END
```

For this example, the name of the ODL file and the task image file have been changed to RESOVR.ODL to distinguish it from the disk-resident version. You can build RESOVR with the following command line:

```
TKB> RESOVR,RESOVR/--SP=RESOVR/MP
```

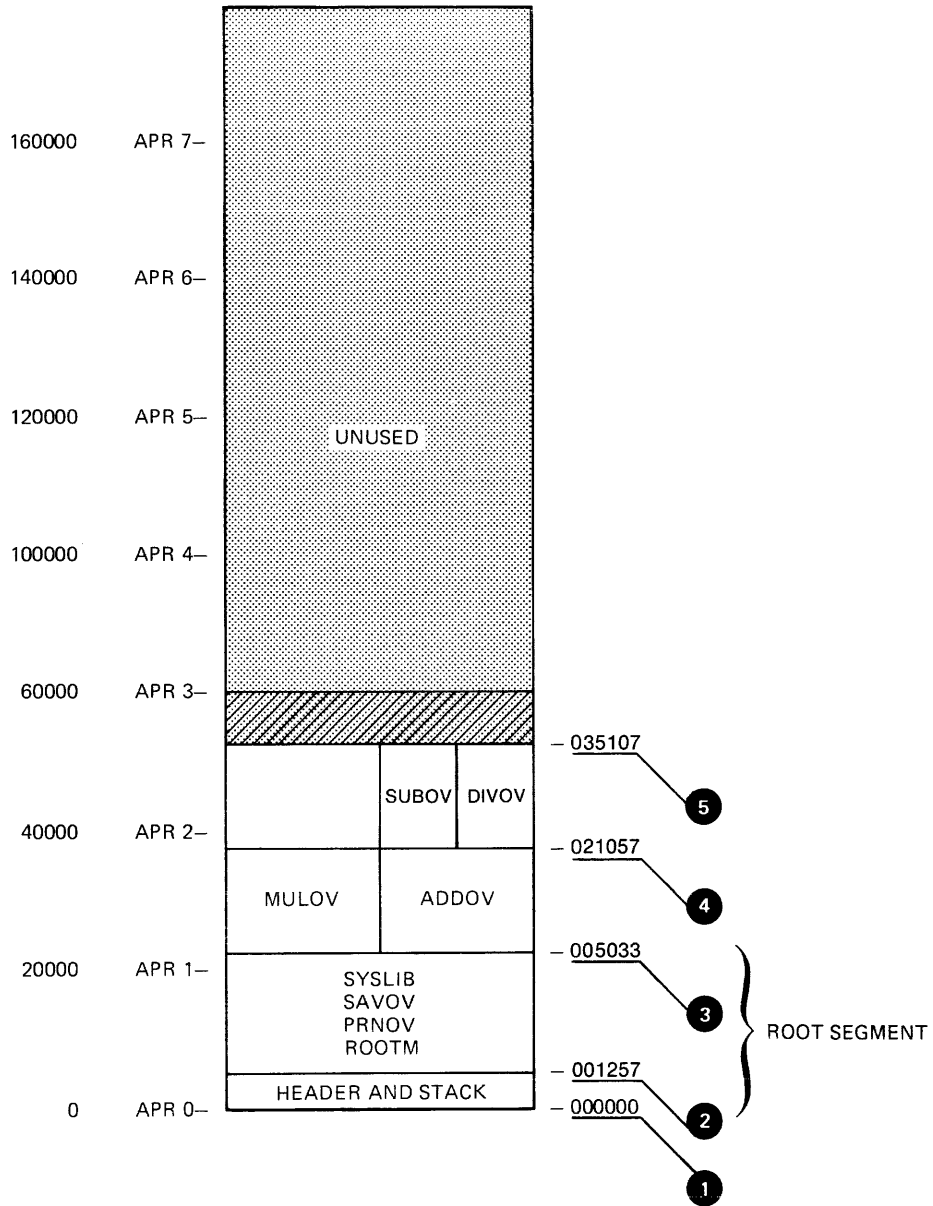
This command directs TKB to build a task named RESOVR.TSK and to create a map file named RESOVR.MAP. The negated spooling switch (/--SP) inhibits spooling of the map file.

The /MP switch on the input file tells TKB that the file is an ODL file and that it will be the only input file for this task build. Refer to Chapter 10 for a description of the switches used in this example.

A portion of the map that results from this task build is shown in Example 3-2.

Figure 3-19 shows the allocation of virtual address space for RESOVR.TSK. The circled numbers in Example 3-2 correspond to those in Figure 3-22.

OVERLAY CAPABILITY



ZK-411-81

Figure 3-21 Allocation of Virtual Address Space for OVR.TSK

OVERLAY CAPABILITY

Example 3-2 Map File for RESOVR.TSK

Partition name : GEN
 Identification : 01
 Task UIC : [7,62]
 Stack limits: 000320 001317 001000 00512.
 PRG xfr address: 001324
 Total address windows: 3. ②
 Task image size : 13920. words
 Task address limits: 000000 057777
 R-W disk blk limits: 000003 000074 000072 00058.

RESOVR.TSK Overlay description:

Base	Top	Length	
000000	005677	005700	03008. ROOTM
020000	034077	014100	06208. MULOV
020000	034077	014100	06208. ADDOV
040000	054077	014100	06208. SUBOV
040000	054077	014100	06208. DIVOV

*** Root segment: ROOTM

R/W mem limits: 000000 005677 005700 03008.
 Disk blk limits: 000003 000010 000006 00006.

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.:(RW,I,LCL,REL,CON)	001320 002514 01356.		
	001320 000102 00066. ROOTM	01	ROOTM.OBJ;1
	001422 000260 00176. PRNOV	01	PRNOV.OBJ;1
	001702 000042 00034. SAVOV	01	SAVOV.OBJ;1
ANS : (RW,D,GBL,REL,OVR)	004034 000002 00002.		
	004034 000002 00002. ROOTM	01	ROOTM.OBJ;1
	004034 000002 00002. PRNOV	01	PRNOV.OBJ;1

Global symbols:

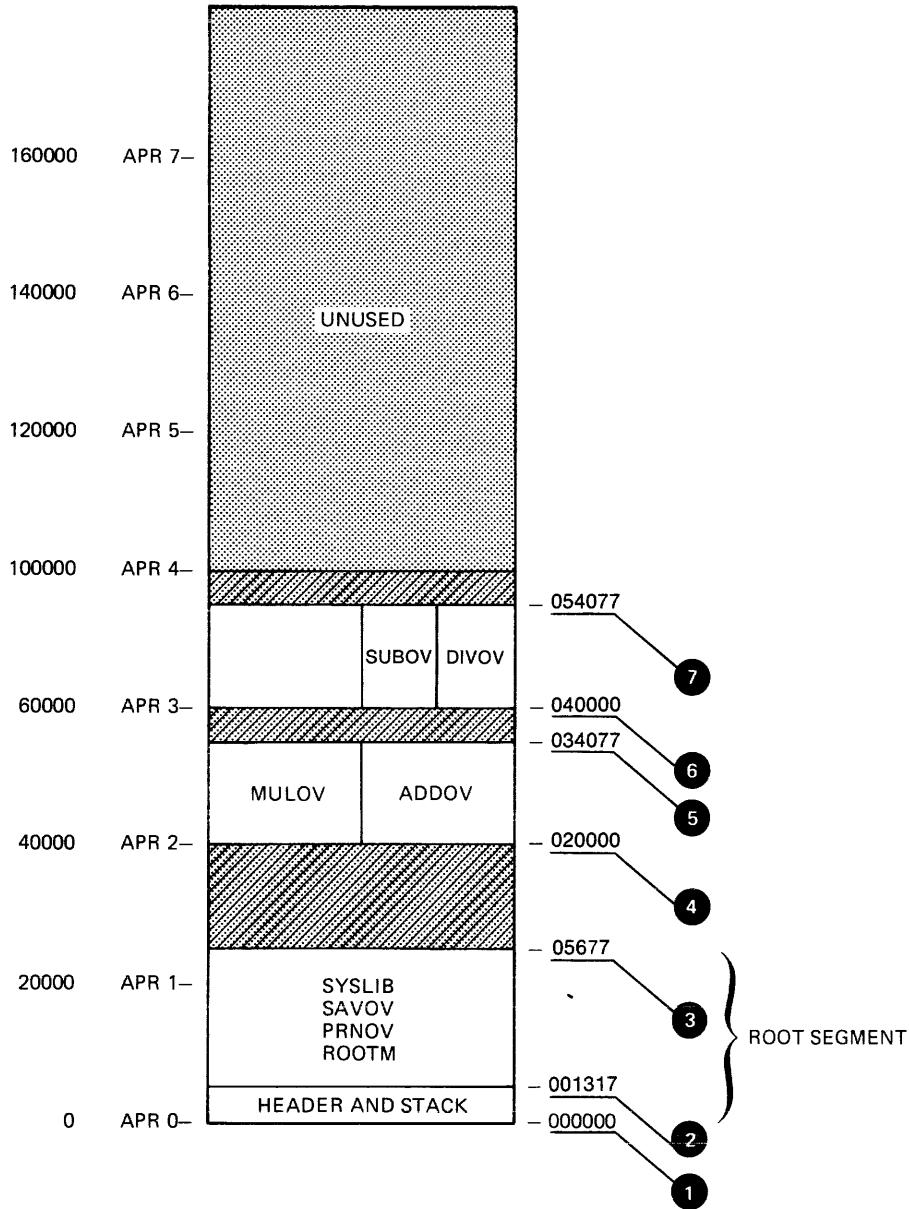
AADD 004072-R DIVV 004112-R PRINT 001610-R SUBB 004102-R
 MULL 004062-R SAVAL 001702-R

*** Task builder statistics:

Total work file references: 6938.
 Work file reads : 0.
 Work file writes : 0.
 Size of core pool: 4178. words (16. pages)
 Size of work file: 3072. words (12. pages)
 Elapsed time:00:00:21

OVERLAY CAPABILITY

Note that TKB allocates virtual address space for each level of overlay segment on a 4K-word boundary. When built as a disk-resident overlay, this structure requires 12K words of virtual address space; when built as a memory-resident overlay structure, it requires 16K words of virtual address space. As noted earlier, you must be careful when using memory-resident overlays to ensure that virtual address space is used efficiently.



ZK-412-81

Figure 3-22 Allocation of Virtual Address Space for RESOVR.TSK

OVERLAY CAPABILITY

Finally, note in Figure 3-22 that TKB has allocated three window blocks to map RESOVR.TSK. Each level of the overlay in a memory-resident overlay requires a separate window block to map it. In a disk-resident overlay, a single window block maps the entire structure regardless of how many segment levels there are within the structure. This consideration can be important when you are building an overlaid task that either creates dynamic regions or accesses a resident library or common, because of the extra window blocks required to use these features.

3.9 SUMMARY OF THE OVERLAY DESCRIPTION LANGUAGE

- An overlay structure consists of one or more trees. Each tree contains at least one segment. A segment is one or more modules containing one or more program sections that can be loaded by a single disk access.

A tree can have only one root segment, but it can have any number of overlay segments.

- An ODL file is a text file consisting of a series of overlay description directives, one directive per line. You enter this file in the TKB command line, and identify it as an ODL file by attaching the MP switch to the file name. If you enter an ODL file in the TKB command line, it must be the only input file you specify.
- The Overlay Description Language provides five directives for specifying the tree representation of the overlay structure:
 - .ROOT and .END -- There can be only one .ROOT and one .END directive; the .END directive must be the last directive because it terminates input.
 - .PSECT, .FCTR, and .NAME -- These can be used in any order in the ODL file.
- You define the tree structure using the hyphen (-), comma (,), and exclamation point (!) operators, and by using parentheses.
 - The hyphen operator (-) indicates that its arguments are to be concatenated and thus are to coexist in memory.
 - The comma operator (,) within parentheses indicates that its arguments are to overlay each other either physically, if disk resident, or virtually, if memory resident.
 - The comma operator not within parentheses delimits overlay trees.
 - The exclamation point operator (!) immediately before a left parenthesis declares the enclosed segments to be memory resident. Nested segments in parentheses are not affected by an exclamation point operator at a level closer to the root.

OVERLAY CAPABILITY

- The parentheses group segments that begin at the same point in memory. For example:

```
.ROOT A-B-(C,D-(E,F))
```

This ODL command line defines an overlay structure with a root segment consisting of the modules A and B. In this structure, there are four overlay segments: C, D, E, and F. The outer pair of parentheses indicates that the overlay segments C and D start at the same virtual address; and similarly, the inner parentheses indicate that E and F start at the same virtual address.

- The `.ROOT` directive defines the beginning overlay structure. The arguments of the `.ROOT` directive are one or more of the following:
 - File specifications as described in Chapter 1
 - Factor labels
 - Segment names
 - Program-section names
- The `.END` directive terminates input.
- The `.FCTR` directive provides a means for replacing text by a symbolic reference (the factor label). This replacement is useful for two reasons:
 - The `.FCTR` directive extends the text of the `.ROOT` directive to more than one line and thus allows complex trees to be represented.
 - The `.FCTR` directive allows you to write the overlay description in a form that makes the structure of the tree more apparent.

For example:

```
.ROOT A-(B-(C,D),E-(F,G),H)
.END
```

Using the `.FCTR` directive, you can write this overlay description as follows:

```
.ROOT A-(F1,F2,H)
F1:   .FCTR B-(C,D)
F2:   .FCTR E-(F,G)
      .END
```

The second representation makes it clear that the tree has three main branches.

- The `.PSECT` directive provides a means for directly specifying the segment in which a program section is placed. It accepts the name of the program section and its attributes. For example:

```
.PSECT ALPHA,CON,GBL,RW,I,REL
```

OVERLAY CAPABILITY

ALPHA is the program section name and the remaining arguments are the program section's attributes (program section attributes are described in Chapter 2).

The program section name (composed of the characters A-Z, 0-9, \$, or .) must appear first in the .PSECT directive, but the attributes can appear in any order or can be omitted. If an attribute is omitted, a default condition is assumed. The defaults for program section attributes are RW, I, LCL, REL, and CON.

In the example above, therefore, you need only specify the attributes that do not correspond to the defaults: .PSECT ALPHA,GBL

- The .NAME directive provides you with the means to designate a segment name for use in the overlay description, and to specify segment attributes. This directive is useful for creating a null segment, naming a segment that is to be loaded manually, or naming a nonexecutable segment that is to be autoloadable. (Refer to Chapter 4 of this manual for a description of manually loaded and automatically loaded segments.) If you do not use the .NAME directive, the Task Builder uses the name of the first file, program section, or library module in the segment to identify the segment.

The .NAME directive creates a segment name as follows:

```
.NAME segname,attr,attr
```

segname

is the designated name (composed of the characters A-Z, 0-9, and \$).

attr

is an optional attribute taken from the following: GBL, NODSK, NOGBL, DSK.

The defaults are NOGBL and DSK. The defined name must be unique with respect to the names of program sections, segments, files, and factor labels.

- You can define a co-tree by specifying an additional tree structure in the .ROOT directive. The first overlay tree description in the .ROOT directive is the main tree. Subsequent overlay descriptions are co-trees. For example:

```
.ROOT A-B-(C,D-(E,F)),X-(Y,Z),Q-(R,S,T)
```

The main tree in this example has the root segment consisting of files A.OBJ and B.OBJ. Two co-trees are defined; the first co-tree has the root segment X and the second co-tree has the root segment Q.

CHAPTER 4

OVERLAY LOADING METHODS

The RSX-11M/M-PLUS systems provide two methods for loading disk-resident and memory-resident overlays:

- Autoload -- The overlay run-time routines are automatically called to load segments you have specified.
- Manual Load -- You include in the task explicit calls to the overlay run-time routines.

When you build an overlaid task, you must decide which one of these methods to use, because both cannot be used in the same task.

The loading process depends on the kind of overlay:

- Disk resident -- A segment is loaded from disk into a shared area of physical memory, writing over whatever was present.
- Memory resident -- A segment is loaded by mapping a set of shared virtual addresses to a unique unshared area of physical memory, where the segment has been made permanently resident (after having been initially brought in from the disk).

With the autoload method, the overlay run-time routines handle loading and error recovery. Overlays are automatically loaded by being referenced through a transfer-of-control instruction (CALL, JMP, or JSR). No explicit calls to the overlay run-time routines are needed.

In the manual load method, you handle loading and error recovery explicitly. Manual loading saves space and gives you full control over the loading process, including the ability to specify whether loading is to be done synchronously or asynchronously.

In the manual load method, you must provide for loading the overlay segments of the main tree, as well as the root segments and the overlay segments of the co-trees. Once loaded, the root segment of a co-tree remains in memory.

4.1 AUTOLOAD

To specify the autoload method, you use the autoload indicator, an asterisk (*). You place this indicator in the ODL description of the task at the points where loading must occur. The execution of a transfer-of-control instruction to an autoloading segment up-tree (farther away from the root) initiates the autoload process.

OVERLAY LOADING METHODS

4.1.1 Autoload Indicator

The autoload indicator (*) marks as autoloadable the segment or other task element (as defined below). If you apply the autoload indicator to an ODL statement enclosed in parentheses, every task element within the parentheses is marked as autoloadable. Placing the autoload indicator at the outermost level of parentheses in the ODL description marks every module in the overlay segments as autoloadable.

In the TK1 example of Chapter 3, Section 3.1.1, if segment C consisted of a set of modules C1, C2, C3, C4, and C5, the tree diagram would be as shown in Figure 4-1.




Figure 4-1 Details of Segment C of TK1

Placing the autoload indicator at the outermost level of parentheses ensures that, regardless of the flow of control within the task, a module will be properly loaded when it is called. The ODL description for task TK1 would be:

```
.ROOT CNTRL-*(AFCTR,BFCTR,CFCTR)
AFCTR: .FCTR A0-(A1,A2-(A21,A22))
BFCTR: .FCTR B0-(B1,B2)
CFCTR: .FCTR C1-C2-C3-C4-C5
.END
```

When you use autoload, the root of a co-tree is loaded by path loading if one of the branches of the co-tree is called before the root. However, if the root of the co-tree is called before the branch is called, the root must have an autoload indicator.

Also, when the root segment of a co-tree is not a null segment, you must mark the co-tree's root segment (CNTRL2) as well as its outermost level of parentheses to ensure that all modules of the co-tree are properly loaded. For example, if the co-tree root (CNTRL2) of the multiple tree example, Section 3.5.2, had contained code or data, it would have been marked as follows:

```
.ROOT CNTRL-*(AFCTR,BFCTR,CFCTR),*CNTRL2-*(CNTRLX,CNTRLY)
.
.
.
```

You can apply the autoload indicator to the following elements:

- File names -- to make all the components of the file autoloadable.
- Portions of ODL tree descriptions enclosed in parentheses -- to make all the elements within the parentheses autoloadable, including elements within any nested parentheses.

OVERLAY LOADING METHODS

- Program section names -- to make the program section autoloading. The program section must have the instruction (I) attribute.
- Segment names defined by the .NAME directive -- to make all components of the segment autoloading.
- .FCTR label names -- to make the first component of the factor autoloading. All elements specified in the .FCTR statement are autoloading if they are enclosed in parentheses.

In the following example, two .PSECT directives and a .NAME directive are introduced into the ODL description for TK1. Autoload indicators are applied as follows:

```
.ROOT CNTRL-(*AFCTR,*BFCTR,*CFCTR) ①
AFCTR: .FCTR A0-*ASUB1-ASUB2-*(A1,A2-(A21,A22)) ② ③
BFCTR: .FCTR (B0-(B1,B2))
CFCTR: .FCTR CNAM-C1-C2-C3-C4-C5
       .NAME CNAM,GBL ①
       .PSECT ASUB1,I,GBL,OV ②
       .PSECT ASUB2,I,GBL,OV
       .END
```

The following notes are keyed to the example above.

- ① The autoload indicator is applied to each factor name; therefore:

- a. *AFCTR=*A0
- b. *BFCTR=*(B0-(B1,B2))
- c. *CFCTR=*CNAM

CNAM, however, is an element defined by a .NAME directive. Therefore, all components of the segment to which the name applies are made autoloading, that is, C1, C2, C3, C4, and C5.

- ② The autoload indicator is applied to the name of a program section with the instruction (I) attribute (*ASUB1), so that program section ASUB1 is made autoloading.
- ③ The autoload indicator is applied to a portion of the ODL description enclosed in parentheses:

```
*(A1,A2-(A21,A22))
```

Thus, every element within the parentheses is made autoloading (that is, files A1, A2, A21, and A22).

The net effect of this ODL description is to make every element except program section ASUB2 autoloading.


4.1.2 Path Loading

The autoload method uses path loading; that is, a call from one segment to another segment up-tree (farther away from the root) ensures that all the segments on the path from the calling segment to the called segment will reside in physical memory and be mapped. Path loading is confined to the tree in which the called segment resides.

OVERLAY LOADING METHODS

A call from a segment in one tree to a segment in another tree results in the loading of all segments on the path in the second tree from the root to the called module.

In Figure 4-2, if CNTRL calls A22, all the modules between the CNTRL and A2 are loaded. In this case, modules A0 and A2 are loaded.



ZK-414-81


Figure 4-2 Path-Loading Example

With the autoload method, the overlay run-time routines keep a record of the segments that are loaded and mapped, and issue disk-load requests only for segments that are not in memory. If CNTRL calls A2 after calling A1, A0 is not loaded again because it is already in memory and mapped.

A reference from one segment to another segment down-tree (closer to the root) is resolved directly. For example, A2 can immediately access A0 because A0 was path loaded in the call to A2.

4.1.3 Autoload Vectors

To resolve a reference up-tree to a global symbol in an autoloadable segment, TKB generates an autoload vector for the referenced global symbol. The reference in the code is changed to a definition that points to an autoload vector entry. The format for the autoload vector for conventional tasks is shown in Figure 4-3 and the format for I- and D-space tasks is shown in Figure 4-4.



ZK-415-81

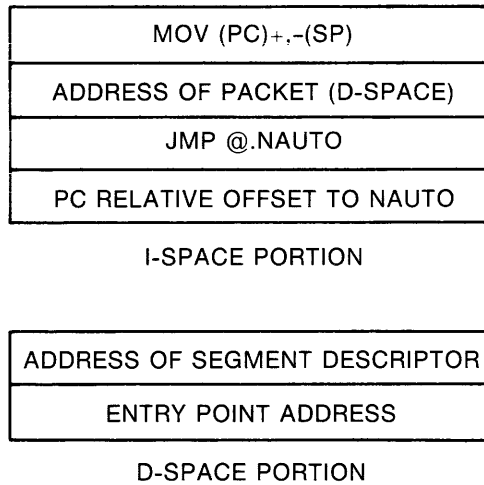
Figure 4-3 Autoload Vector Format for Conventional Tasks

For I- and D-space tasks, TKB generates the autoload vector in a format that differs from the vector in a conventional task. The I- and D-space autoload vector is six words long and consists of two parts: one part residing in I-space and the other part residing in D-space. The I-space part consists of two 2-word instructions, and the D-space part consists of two words of data. The data in the vector are the segment descriptor address and the target entry point address. The I- and D-space vector is shown in Figure 4-4.

OVERLAY LOADING METHODS

The task root and the overlay segments may contain autoload vectors; the I-space part of the root or segment contains the I-space part of the vectors, and the D-space part of the root or segment contains the D-space part of the vectors.

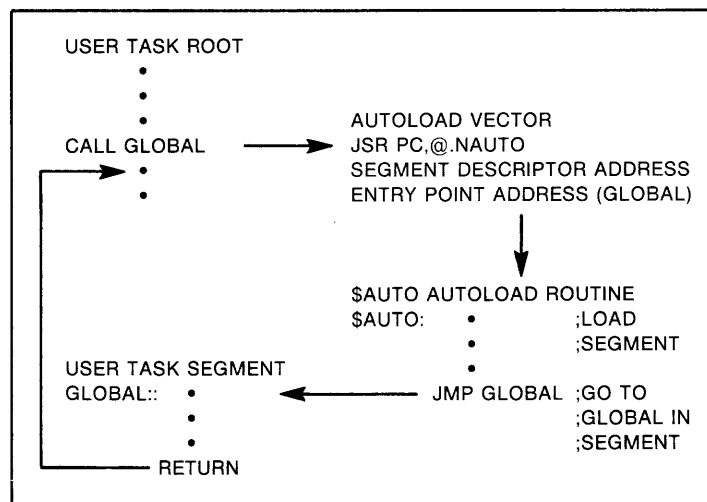
The MOV instruction in the I-space part of the vector places the address of the D-space part of the vector on the stack. The second instruction in the vector executes an indirect JMP to \$AUTO through the location .NAUTO.



ZK-1089-82

Figure 4-4 Autoload Vector Format for I- and D-space Tasks

In Figures 4-3 and 4-4, a transfer-of-control instruction to the up-tree global symbol generates an autoload vector in the shown format. An example of the code sequence used in a call to a global symbol in an autoloadable segment is shown in Figure 4-5.



ZK-416-81

Figure 4-5 Example Autoload Code Sequence for a Conventional Task

OVERLAY LOADING METHODS

An exception to the procedure for generating autoloading vectors is made in the case of a program section with the data (D) attribute. References from a segment to a global symbol up-tree in a program section with the data (D) attribute are resolved directly.

Because TKB can obtain no information about the flow of control within the task, it often generates more autoloading vectors than are necessary. However, your knowledge of the flow of control within your task, and of path loading, can help you determine where to place the autoloading indicators. By placing the autoloading indicators only at the points where loading is actually required, you can minimize the number of autoloading vectors generated for the task.

In the following example, all the calls to overlays originate in the root segment. That is, no module in an overlay segment calls outside its segment. The root segment CNTRL has the following contents:

```
PROGRAM CNTRL
CALL A1
CALL A21
CALL A2
CALL A0
CALL A22
CALL B0
CALL B1
CALL B2
CALL C1
CALL C2
CALL C3
CALL C4
CALL C5
END
```

If you place the autoloading indicator at the outermost level of parentheses, 13 autoloading vectors are generated for this task; however, because A2 and A0 are loaded by path loading to A21, the autoloading vectors for A2 and A0 are unnecessary. Moreover, because the call to C1 loads the segment that contains C2, C3, C4, and C5, autoloading vectors for C2 through C5 are unnecessary.

You can eliminate the unnecessary autoloading vectors by placing the autoloading indicator only at the points where explicit loading is required, as follows:

```
                .ROOT CNTRL-(AFCTR,*BFCTR,CFCTR)
AFCTR:          .FCTR A0-(*A1,A2-*(A21,A22))
BFCTR:          .FCTR (B0-(B1,B2))
CFCTR:          .FCTR *C1-C2-C3-C4-C5
                .END
```

With this ODL description, TKB generates seven autoloading vectors -- for A1, A21, A22, B0, B1, B2, and C1.

OVERLAY LOADING METHODS

4.1.4 Autoloadable Data Segments

You can make overlay segments that contain no executable code autoloadable, as follows. First, you must include a `.NAME` directive and specify the `GBL` attribute, as described in Section 3.4.4. For example:

```
      .ROOT A-*(B,C)
      .NAME BNAME,GBL
B:    .FCTR BNAME-BFIL
      .END
```

The global symbol `BNAME` is created and entered into the symbol table of segment `BNAME`. Because this segment is marked to be autoloaded, root segment `A` calls segment `BNAME` as follows:

```
      CALL BNAME
```

The segment is autoloaded and an immediate return to inline code occurs.

The data of `BFIL` must be placed in a program section with the data (`D`) attribute to suppress the creation of autoload vectors.

4.2 MANUAL LOAD

If you decide to use the manual-load method to load segments, you must include in your program explicit calls to the `$LOAD` routine. These load requests must supply the name of the segment to be loaded. In addition, they can include information necessary to perform asynchronous load requests, and to handle load request failures.

The `$LOAD` routine does not path load. A call to `$LOAD` loads only the segment named in the request. The segment is read in from disk and mapped. For memory-resident overlays; the segment is mapped, but only read in if it was not previously read in.

A MACRO-11 programmer calls the `$LOAD` routine directly. A FORTRAN programmer calls `$LOAD` using the FORTRAN subroutine `MNLOAD`.

4.2.1 MACRO-11 Manual Load Calling Sequence

A MACRO-11 programmer calls `$LOAD` as follows:

```
      MOV     #PBLK,R0
      CALL    $LOAD
```

`PBLK` is the address of a parameter block with the following format:

```
PBLK:  .BYTE length,event-flag
        .RAD50 /seg-name/
        .WORD [i/o-status] or 0
        .WORD [ast-trp] or 0
```

`length`

The length of the parameter block (3 to 5 words).

OVERLAY LOADING METHODS

event-flag

The event flag number, used for asynchronous loading. If the event-flag number is 0, synchronous loading is performed.

seg-name

The name of the segment to be loaded: a 1- to 6-character Radix-50 name, occupying two words.

i/o-status

The address of the I/O status doubleword. Standard QIO status codes apply.

ast-trp

The address of an asynchronous trap service routine to which control is transferred at the completion of the load request.

The condition code C-list is set or cleared on return, as follows:

- If condition code C=0, the load request was accepted.
- If condition code C=1, the load request was unsuccessful.

For a synchronous load request, the return of the condition code C=0 means that the desired segment is loaded and is ready to be executed. For an asynchronous load request, the return of the code C=0 means that the load request was successfully queued to the device driver, but the segment is not necessarily in memory. Your program must ensure that loading has been completed by waiting for the specified event flag before calling any routines or accessing any data in the segment.

4.2.2 MACRO-11 Manual Load Calling Sequence For I- and D-Space Tasks

A MACRO-11 programmer calls \$LOAD as follows:

```
MOV #PBLK,R0  
CALL $LOAD
```

PBLK is the address of a parameter block with the following format in an I- and D-space task:

```
PBLK: BYTE 3,0  
      .RAD50 /seg-name/
```

length

The length of the parameter block (3 words).

OVERLAY LOADING METHODS

event-flag

Specify this as 0. Only synchronous load requests are possible when loading I- and D-space segments.

seg-name

The name of the segment to be loaded: a 1- to 6-character Radix-50 name, occupying two words.

The condition code C-list is set or cleared on return, as follows:

- If condition code C=0, the load request was accepted.
- If condition code C=1, the load request was unsuccessful.

For a synchronous load request, which is the only one possible for I- and D-space segments, the return of the condition code C=0 means that the desired segment is loaded and is ready to be executed. Your program must ensure that loading has been successful by checking for the condition code rather than assuming that the segment has been loaded.

4.2.3 FORTRAN Manual Load Calling Sequence

To use the manual load mechanism in a FORTRAN program, your program must refer to the \$LOAD routine by means of the MNLOAD subroutine. The subroutine call has the form:

```
CALL MNLOAD(seg-name,[event-flag],[i/o-status],[ast-trp],[ld-ind])
```

seg-name

A 2-word real variable containing the segment name in Radix-50 format.

event-flag

An optional integer event flag number used for an asynchronous load request. If the event flag number is 0, the load request is synchronous.

i/o-status

An optional 2-word integer array containing the I/O status doubleword, as described for the QIO directive in the RSX-11M/M-PLUS Executive Reference Manual.

ast-trp

An optional asynchronous trap subroutine entered at the completion of a request. MNLOAD requires that all pending traps specify the same subroutine.

OVERLAY LOADING METHODS

ld-ind

An optional integer variable containing the results of the subroutine call. One of the following values is returned:

- +1 Request was successfully executed.
- 1 Request had bad parameters or was not successfully executed.

You can omit optional arguments. The following calls are legal:

Call	Effect
CALL MNLOAD (SEG1)	Loads segment named in SEG1 synchronously.
CALL MNLOAD (SEG1,0,,,LDIND)	Loads segment named in SEG1 synchronously and returns success indicator to LDIND.
CALL MNLOAD (SEG1,1,IOSTAT,ASTSUB,LDIND)	Loads segment named in SEG1 asynchronously, transferring control to ASTSUB upon completion of the load request; stores the I/O status doubleword in IOSTAT and the success indicator in LDIND.

The following example uses the program CNTRL, previously discussed in Section 4.1. In this example, there is sufficient processing between the calls to the overlay segments to make asynchronous loading effective. The autoloading indicators are removed from the ODL description and the FORTRAN programs are recompiled with explicit calls to the MNLOAD subroutine, as follows:

```
PROGRAM CNTRL
EXTERNAL ASTSUB
DATA SEG1 /6RA1 /
DATA SEGA21 /6RA21 /
.
.
.
CALL MNLOAD (SEG1,1,IOSTAT,ASTSUB,LDIND)
.
.
.
CALL A1
.
.
.
CALL MNLOAD (SEGA21,1,IOSTAT,ASTSUB,LDIND)
.
.
.
```


OVERLAY LOADING METHODS

```
CALL A21
  .
  .
  .
END
SUBROUTINE ASTSUB
DIMENSION IOSTAT(2)
  .
  .
  .
END
```

When the AST trap routine is used, the I/O status doubleword is automatically supplied to the dummy variable IOSTAT.

4.2.4 FORTRAN Manual Load Calling Sequence for I- and D-Space Tasks

To use the manual load mechanism in a FORTRAN program, your program must refer to the \$LOAD routine by means of the MNLOAD subroutine. The subroutine call has the form:

```
CALL MNLOAD(seg-name,,, [ld-ind])
```

seg-name

A 2-word real variable containing the segment name in Radix-50 format.

ld-ind

An optional integer variable containing the results of the subroutine call. One of the following values is returned:

- +1 Request was successfully executed.
- 1 Request had bad parameters or was not successfully executed.

You can omit optional arguments. The following calls are legal:

Call	Effect
CALL MNLOAD (SEGB1)	Loads segment named in SEGB1 synchronously.
CALL MNLOAD (SEGB1,,,LDIND)	Loads segment named in SEGB1 synchronously and returns success indicator to LDIND.

Only synchronous loading is possible when manually loading I- and D-space task segments.

4.3 ERROR HANDLING

If you use the autoload mechanism, a simple recovery procedure is provided that checks the Directive Status Word (DSW) for an error indication. If the DSW indicates that no system dynamic storage is available, the routine issues a Wait for Significant Event directive and tries again; if the problem is not dynamic storage, the recovery

OVERLAY LOADING METHODS

procedure generates a synchronous breakpoint trap. If the task services the trap and returns without altering the state of the program, the request will be retried.

If you select the manual-load method, you must provide error handling routines that diagnose load errors and provide appropriate recovery. A more comprehensive user-written error recovery subroutine can be substituted for the system-provided routine if the following conventions are observed:

- The error recovery routine must have the entry point name \$ALERR.
- The contents of all registers must be saved and restored.

On entry to \$ALERR, R2 contains the address of the segment descriptor that could not be loaded. Before recovery action can be taken, the routine must determine the cause of the error by examining the following words in the sequence indicated:


1. \$DSW The Directive Status Word may contain an error status code, indicating that the Executive rejected the I/O request to load the overlay segment.
2. N.OVPT The contents of this location, offset by N.IOST, point to a 2-word I/O status block containing the results of the load overlay request returned by the device driver. The status code occupies the low-order byte of word 0. For example, for a device-not-ready condition, the code will be IE.DNR. (For more information on these codes, refer to the IAS/RSX-11 I/O Operations Reference Manual.)

4.4 GLOBAL CROSS-REFERENCE OF AN OVERLAID TASK

This section illustrates a global cross-reference that has been created for an overlaid task. The task consists of a root segment containing the module ROOT.OBJ, and overlay segments composed of modules OVR1, OVR2, OVR3, and OVR4. The overlay description of the file is as follows:

```
.ROOT  ROOT-(OVR,*OVR2)
OVR: .FCTR OVR1-*(OVR3,OVR4)
```

Only segments OVR2, OVR3, and OVR4 are autoloadable. Figure 4-6 shows the resulting overlay tree.



ZK-417-81

Figure 4-6 Autoload Overlay Tree Example

OVERLAY LOADING METHODS

As shown, the global symbol OVR1 is defined in module OVR1, and a single nonautoloadable, up-tree reference is made to this symbol by the module ROOT, as indicated by the circumflex. Note that because OVR1 is not autoloadable, it depends on a call to OVR3 or OVR4 to get loaded by path loading. The asterisk indicates that the module contains an autoloadable definition. The modules shown with the asterisk define the symbol.

The asterisks preceding the modules OVR2, OVR3, and OVR4 indicate that the global symbols OVR2, OVR3, and OVR4 are autoload symbols and are referenced from the module ROOT through an autoload vector, as shown by the at-sign (@) character.

The asterisk and at-sign are shown in the cross-reference listing in Example 4-1.

Example 4-1 Cross-Reference Listing of Overlaid Task

OVRTST	CREATED BY	TKB	ON 27-JUL-82 AT 12:04	PAGE 1
GLOBAL CROSS REFERENCE				CREF V01
SYMBOL	VALUE	REFERENCES...		
N.ALER	000010	AUTO	#	OVRES
N.IOST	000004	OVCTL	#	OVRES
N.MRKS	000016	#		OVRES
N.OVLY	000000	OVCTL	#	OVRES
N.OVPT	000054	AUTO		OVCTL # VCTDF
N.RDSG	000014	#		OVRES
N.STBL	000002	#		OVRES
N.SZSG	000012	#		OVRES
OVR1	002014-R	#	OVR1	^ ROOT
OVR2	002014-R	*	OVR2	@ ROOT
OVR3	002014-R	*	OVR3	@ ROOT
OVR4	002014-R	*	OVR4	@ ROOT
ROOT	001176-R	#		ROOT
\$ALBP1	001320-R	#		AUTO
\$ALBP2	001416-R	#		AUTO
\$ALERR	001246-R	#	ALERR	OVDAT
\$AUTO	001302-R	#		AUTO
\$DSW	000046		ALERR	# VCTDF
\$MARKS	001546-R	#	OVCTL	
\$OTSV	000052	#		VCTDF
\$SAVRG	001452-R		AUTO	# SAVRG
\$VEXT	000056	#		VCTDF
.FSRPT	000050	#		VCTDF
.NALER	001442-R	#		OVDAT
.NIOST	001436-R	#		OVDAT
.NMRKS	001450-R	#		OVDAT
.NOVLY	001432-R	#		OVDAT
.NOVPT	000042	#		OVDAT
.NRDSG	001446-R	#		OVDAT
.NSTBL	001434-R	#		OVDAT
.NSZSG	001444-R	#		OVDAT

OVERLAY LOADING METHODS

Example 4-1 (Cont.) Cross-Reference Listing of Overlaid Task

```
OVRTST      CREATED BY   TKB      ON 27-JUL-82 AT 12:04      PAGE 2
SEGMENT CROSS REFERENCE                                     CREF   V01
SEGMENT NAME      RESIDENT MODULES
OVR1              OVR1
OVR2              OVR2
OVR3              OVR3
OVR4              OVR4
ROOT              ALERR   AUTO   OVCTL   CVDAT   OVRES   ROOT   SAVRG
                  VCTDF
```

Down-tree references to the global symbol ROOT are made from modules OVR1, OVR2, OVR3, and OVR4. These references are resolved directly.

The segment cross-reference shows the segment name and modules in each overlay.

4.5 USE AND SIZE OF OVERLAY RUNTIME ROUTINES

TKB, when constructing an overlaid task, incorporates certain modules from the system library to perform the actual overlay operations. An overlay run-time routine in the task loads overlays from disk or maps resident overlays by issuing QIO\$ or CRAW\$ directives.

The modules and routines described below implement the TKB autoloading mechanism as described in Section 4.1.

There are three major components to the autoloading service, as follows:

AUTO This module controls the overlay process, and the autoloading vectors indirectly call AUTO through .NAUTO. AUTO determines whether the referenced overlay segment is already in memory or mapped. It then jumps to the required entry point if the entry point is available.

The AUTO module is supplied in two variations. These variations are separately named and described as follows:

AUTO Selected by TKB by default for all overlaid tasks. It manages disk-only, PLAS, and cluster library overlay structures.

AUTOT Manually selected by you by using an explicit reference in the TKB .ODL file, as shown below. This module disables the AST traps while manipulating the overlay data structures. This is required where user task AST traps might cause modification of the overlay database. To incorporate this module in your task image, you

OVERLAY LOADING METHODS

must include the following element in the .ROOT factor of the task's ODL file:

```
-LB:[1,1]SYSLIB/LB:AUTOT-
```

In addition to including AUTOT in the .ROOT factor, the following code must be included in your task as initialization prior to the AST handling routines in your task:

```
MOV @#.NOVPT,R0  
BISB #200,N.FAST(R0)
```

MRKS This routine traverses the overlay descriptor data structure to mark any overlay segment that will be displaced by a new overlay as "out of memory" and consequently not available.

RDSG The AUTO module calls the RDSG routine repeatedly to read or map each segment along the overlay tree path into the task's virtual address space. This is referred to as "path loading." When path loading is completed, AUTO calls the required entry point.

Several versions of each component exist reflecting the various sizes as appropriate for tasks having disk-only overlays, PLAS mapped overlays, and/or cluster libraries. TKB incorporates the smallest support routines appropriate for the overlay structure of your task.

Depending on whether your task has disk-only overlays, resident overlays, or cluster libraries, TKB forces one of the following modules into your task:

OVCTL Contain the MRKS and RDSG routines optimized for disk overlays only. No support is included for memory-resident or cluster overlays. **OVCTL** is the module included for conventional tasks, and **OVIDL** is the module included for I- and D-space tasks.

OVCTR Contain MRKS and RDSG routines assembled for disk and memory resident overlays. TKB selects either of these modules if the task overlay structure includes memory-resident overlays or maps a resident library containing resident overlays. **OVCTR** is the module selected for conventional tasks and non-overlaid I- and D-space tasks. **OVIDR** is the module selected for overlaid I- and D-space tasks.

OVCTC Contain the MRKS, RDSG, and cluster library support routines. TKB includes **OVCTC** or **OVIDC** if cluster libraries are included in your task. **OVCTC** is the module selected for conventional tasks and non-overlaid I- and D-space tasks. **OVIDC** is the module selected for overlaid I- and D-space tasks.

OVERLAY LOADING METHODS

Two other modules are incorporated into your task's image. They are:

- OVDAT** A small, impure data area used by AUTO, MRKS, and RDSG routines. TKB includes OVDAT in all overlaid tasks, and its size is independent of the overlay structure of that task.
- ALERR** An error service module that AUTO invokes under one of the following circumstances:
- If an I/O error occurs while attempting to read a disk overlay into memory
 - If a directive error occurs while attempting to attach or map a region containing memory resident overlays

Table 4-1 compares the sizes of the overlay run-time support modules. You can use it to determine when it is appropriate to force certain variants into your task image.

Table 4-1
Comparison of Overlay Run-Time Module Sizes

Module	Program Section	Number of Bytes Oct/Dec	Specific Use
<p>One of the following modules is included in any overlaid task that uses autoloading and in any task that links to a PLAS overlaid resident library.</p>			
AUTO	\$\$AUTO	122/82.	All tasks that use autoloading
AUTOT	\$\$AUTO	132/90.	All tasks with ASTs
	\$\$RTQ	32/26.	disabled during autoloading
	\$\$RTR	30/24.	

One of the following modules is included in any overlaid conventional task. OVCTR or OVCTC is included in any non-overlaid task (conventional or I- and D- space) that links to a PLAS overlaid resident library.

OVCTL	\$\$MRKS	76/62.	Disk overlays only
	\$\$RDSG	160/112.	
	\$\$PDLs	2/2.	
OVCTR	\$\$MRKS	234/156.	Disk and PLAS overlays with no cluster libraries
	\$\$RDSG	332/218.	
	\$\$PDLs	12/10.	
OVCTC	\$\$MRKS	254/172.	Disk and PLAS overlays with cluster libraries
	\$\$RDSG	352/234.	
	\$\$PDLs	120/80.	

(continued on next page)

OVERLAY LOADING METHODS

Table 4-1 (Cont.)
Comparison of Overlay Run-Time Module Sizes

Module	Program Section	Number of Bytes Oct/Dec	Specific Use
<p>One of the following three modules is included in any overlaid I- and D-space task.</p>			
OVIDL	\$\$MRKS	76/62.	Disk overlays only
	\$\$RDSG	224/148.	
	\$\$PDLs	2/2.	
OVIDR	\$\$MRKS	304/196.	Disk and PLAS overlays with no cluster libraries
	\$\$RDSG	502/322.	
	\$\$PDLs	12/10.	
OVIDC	\$\$MRKS	324/212.	Disk and PLAS overlays with cluster libraries
	\$\$RDSG	522/338.	
	\$\$PDLs	120/80.	
<p>The overlay data vector OVDAT is included in any overlaid task and in any task that links to a PLAS overlaid resident library.</p>			
OVDAT	\$\$OVDT	24/20.	Included in all tasks that perform overlay operations
	\$\$SGD0	0/0.	
	\$\$SGD2	2/2.	
	\$\$RTQ	0/0.	
	\$\$RTR	0/0.	
	\$\$RTS	2/2.	
<p>The overlay error service routine ALERR is included whenever OVDAT is included.</p>			
ALERR	\$\$ALER	24/20.	Overlay error
<p>Manual overlay control (LOAD) is used in place of any AUTO routine. (See Section 4.2, Manual Load.)</p>			
LOAD	\$\$LOAD	252/170.	Manual overlay control
	\$\$AUTO	14/12.	

CHAPTER 5

SHARED REGION CONCEPTS AND EXAMPLES

The Task Builder provides you with many ways of using shared regions for tailoring your tasks to meet your specific requirements. This chapter describes some of these facilities and their applications.

This chapter contains five working examples. The discussion of the examples assumes that you are familiar with the programming concepts described in the RSX-11M/M-PLUS Guide to Program Development and with the first four chapters of this manual.

5.1 SHARED REGIONS DEFINED

A shared region is a block of data or code that resides in memory and can be used by any number of tasks. A shared region can contain data for use by several tasks or it may be an area where one task writes data for use by another task. Also, a shared region can contain routines for use by several tasks.

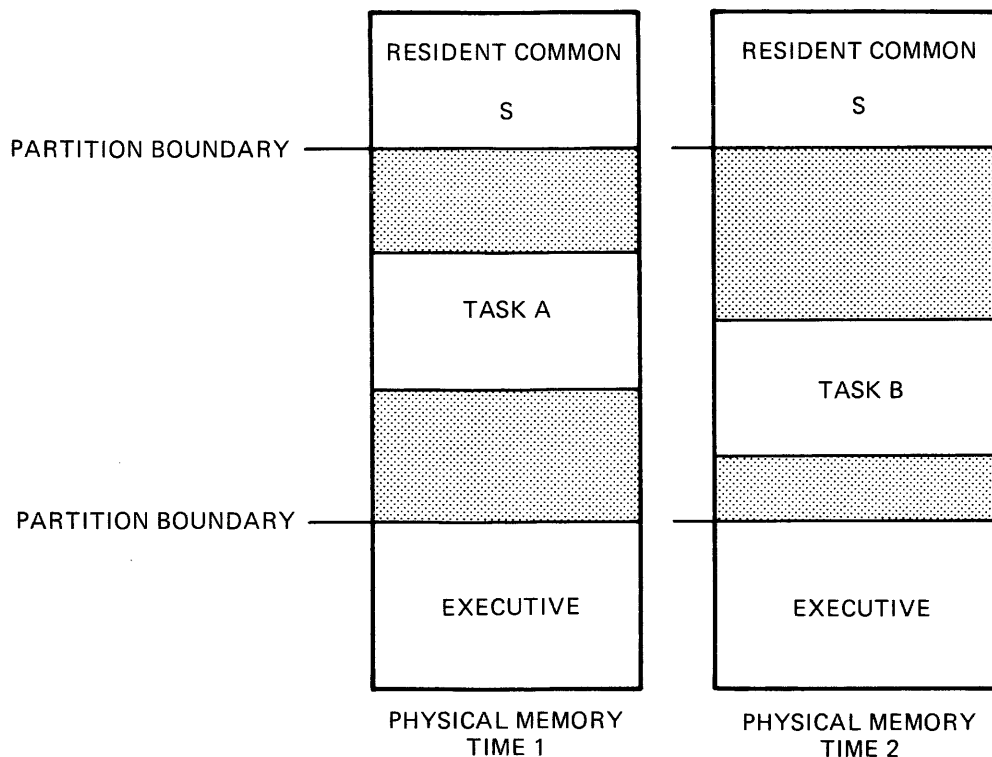
Shared regions are useful because they make more efficient use of physical memory. The two kinds of shared regions are:

- A resident common that provides a way that two or more tasks can share their data
- A resident library that provides a way that two or more tasks can share a single copy of commonly used subroutines

The term "resident" denotes a shared region that is built and installed into the system separately from the task that links to it. In other words, you use TKB to build a shared region much as you would use it to build a task. However, the region does not have a header or a stack. Also, you can use switches to designate the kind of shared region (a library or a common) to be built.

Figure 5-1 shows a typical resident common. Task A stores some results in resident common S, and Task B retrieves the data from the common at a later time.

Figure 5-2 shows a typical resident library. In this case, common reentrant subroutines are not included in each task image; instead, a single copy is shared by all tasks.



ZK-418-81

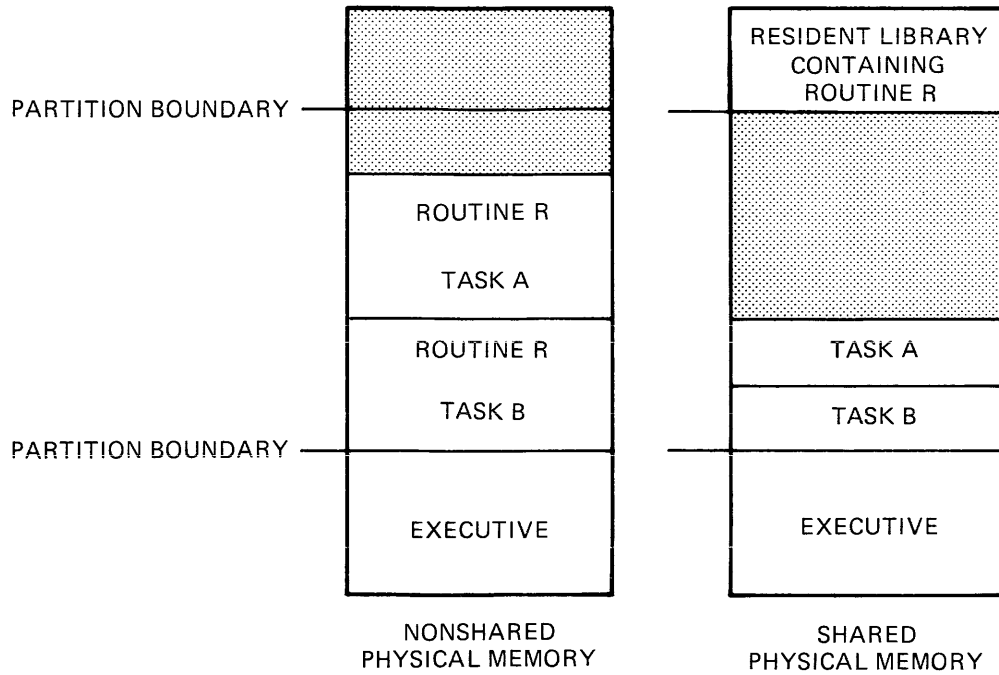
Figure 5-1 Typical Resident Common

When you build a shared region, you must specify an output image file name for the region in the TKB command sequence. But, because a shared region is not an executable unit, it is not a task, and does not require a header or a stack area. Therefore, when you build a shared region, you always attach the negated header switch (/HD) to the region's image file specification. This switch tells TKB to suppress the header within the image. To suppress the stack area in the Task Builder command sequence during option input, you specify STACK=0. (Refer to Chapters 10 and 11 for a complete description of the /HD switch, the STACK option, and other switches and options.)

In either an RSX-11M or RSX-11M-PLUS system, when you build a shared region, you use the PAR option to name the partition in which the region is to reside. You specify the partition name in the TKB command sequence during option input. (Refer to Chapter 11 for a description of the PAR option.) In an RSX-11M system, the partition named in the PAR option must previously exist when the common is installed. It need not exist when the common is being built by TKB. In an RSX-11M-PLUS system, the partition named in the PAR option need not previously exist, but the actual partition defaults to GEN. On both systems, the name used in the PAR option must be the same name as that of the region.

In an RSX-11M system, a shared region must reside in its own partition. Therefore, when you generate your system, you must consider the physical memory requirements, in terms of partitions, for any shared regions that you expect to reside within your system. If you do not consider these requirements at system generation time, later, when you build a shared region, you will be forced to go back and create a common partition for the region.

SHARED REGION CONCEPTS AND EXAMPLES



ZK-419-81

Figure 5-2 Typical Resident Library

In an RSX-11M-PLUS system, shared regions do not have to reside within partitions of their own; you can install a shared region in any partition large enough to hold it. In fact, the partition for which the shared region was built does not have to exist in the system at the time the shared region is installed. Then, it follows that a TKB command sequence or build file for a memory-resident overlaid library must contain the statement PAR=xxx, where xxx is the same name as that of the region being built. Then, when you attempt to install the shared region in a partition that does not exist, the INSTALL processor installs it in the GEN partition and displays the following message on your terminal:

```
INS--PARTITION parname NOT IN SYSTEM DEFAULTING TO GEN
```

Also, you should consider three switches when you build the region. The /PI switch determines whether the region is relocatable. You can use the /CO switch in the TKB command sequence to declare a region as a shared common. The /CO switch specifies the use of the region as a shared common rather than as a shared library. Alternatively, you can use the /LI switch when you build the region to declare the region as a shared library. Using these three switches affects the contents of the symbol definition file, which is described in Chapter 10 under the /CO, /LI, and /PI switch headings. See also Figure 5-3, Interaction of the /LI, /CO, and /PI Switches. The contents of the symbol definition file is described in the following sections.

SHARED REGION CONCEPTS AND EXAMPLES

5.1.1 The Symbol Definition File

When you build a shared region, you must specify a symbol definition (.STB) file in the TKB command sequence. This file contains linkage information about the region. (The format at a .STB file as input to TKB is the same as that of a .OBJ file. See Appendix A.) Later, when you build a task that links to the region, TKB uses this .STB file to resolve calls from within the referencing task to locations within the region.

The /PI switch declares a shared region to be relocatable. Conversely, the /-PI switch declares a shared region to be absolute. If you specify the /PI switch without the /CO or /LI switches to indicate a relocatable region, TKB defaults to building relocatable (position-independent) shared regions with all program sections declared in the .STB file. TKB also defaults to building absolute (position-dependent) shared regions with only the .ABS. program section declared in the .STB file. The contents of the .STB file when these three switches are used are described in Chapter 10 under the /CO, /LI, and /PI switch headings. See also Figure 5-3, Interaction of the /LI, /CO, and /PI Switches.

SWITCH SPECIFIED WITH /-HD	SHARED REGION		REGION PSECT NAME	.STB FILE PSECT	.STB FILE SYMBOLS
	ABSOLUTE	RELOCATABLE			
/PI/LI		YES	SAME AS LIBRARY ROOT	ONE PSECT RELOCATABLE	ALL SYMBOLS. RELATIVE TO START OF THE PSECT
/PI/CO		YES	ALL DECLARED PSECT NAMES INCLUDED	ALL DECLARED PSECTS RELOCATABLE	ALL PSECTS AND SYMBOLS
/-PI/LI*	YES		.ABS	ONE PSECT ABSOLUTE	ALL SYMBOLS ABSOLUTE
/-PI/CO*	YES		ALL DECLARED PSECT NAMES INCLUDED	ALL DECLARED PSECTS ABSOLUTE	ALL SYMBOLS ABSOLUTE
/PI		YES	SAME AS /PI/CO		
/-PI*	YES		SAME AS /-PI/LI		
NONE	YES		SAME AS /-PI/LI		

*/-PI is the default of not using /PI

ZK-420-81

Figure 5-3 Interaction of the /LI, /CO, and /PI Switches

SHARED REGION CONCEPTS AND EXAMPLES

If you do not use either /CO or /LI, the contents of an .STB file for a shared region depend on the use of the /PI switch, which determines whether the region is absolute or relocatable. The effects of declaring a shared region relocatable or absolute and the resulting contents of the .STB file are described in the following sections.

Some .STB files include an entry in the .STB file for each program section in the region. Each entry declares the program section's name, attributes, and length. In addition, TKB includes in the .STB file every symbol in the shared region and its value relative to the beginning of the section in which it resides.

5.1.2 Position-Independent Shared Regions

A position-independent shared region can be placed anywhere in a referencing task's virtual address space when the system on which the task runs has memory management hardware.

5.1.2.1 Position-Independent Shared Region Mapping - In the example of using the memory management APRs, shown in Figure 5-4, two tasks refer to the shared region S: task A and task B. The shared region S is 4K words long and therefore requires that much space in the virtual address space of both tasks. Task A is 6K words long and requires two APRs (APR0 and APR1) to map its task region. The first APR available to map the shared region is APR 2. Thus, you can specify APR 2 when task A is built.

Task B is 16.5K words long. It requires five APRs to map its task region. The first APR available to map the shared region S in task B is APR 5. Therefore, you can specify APR 5 when task B is built.

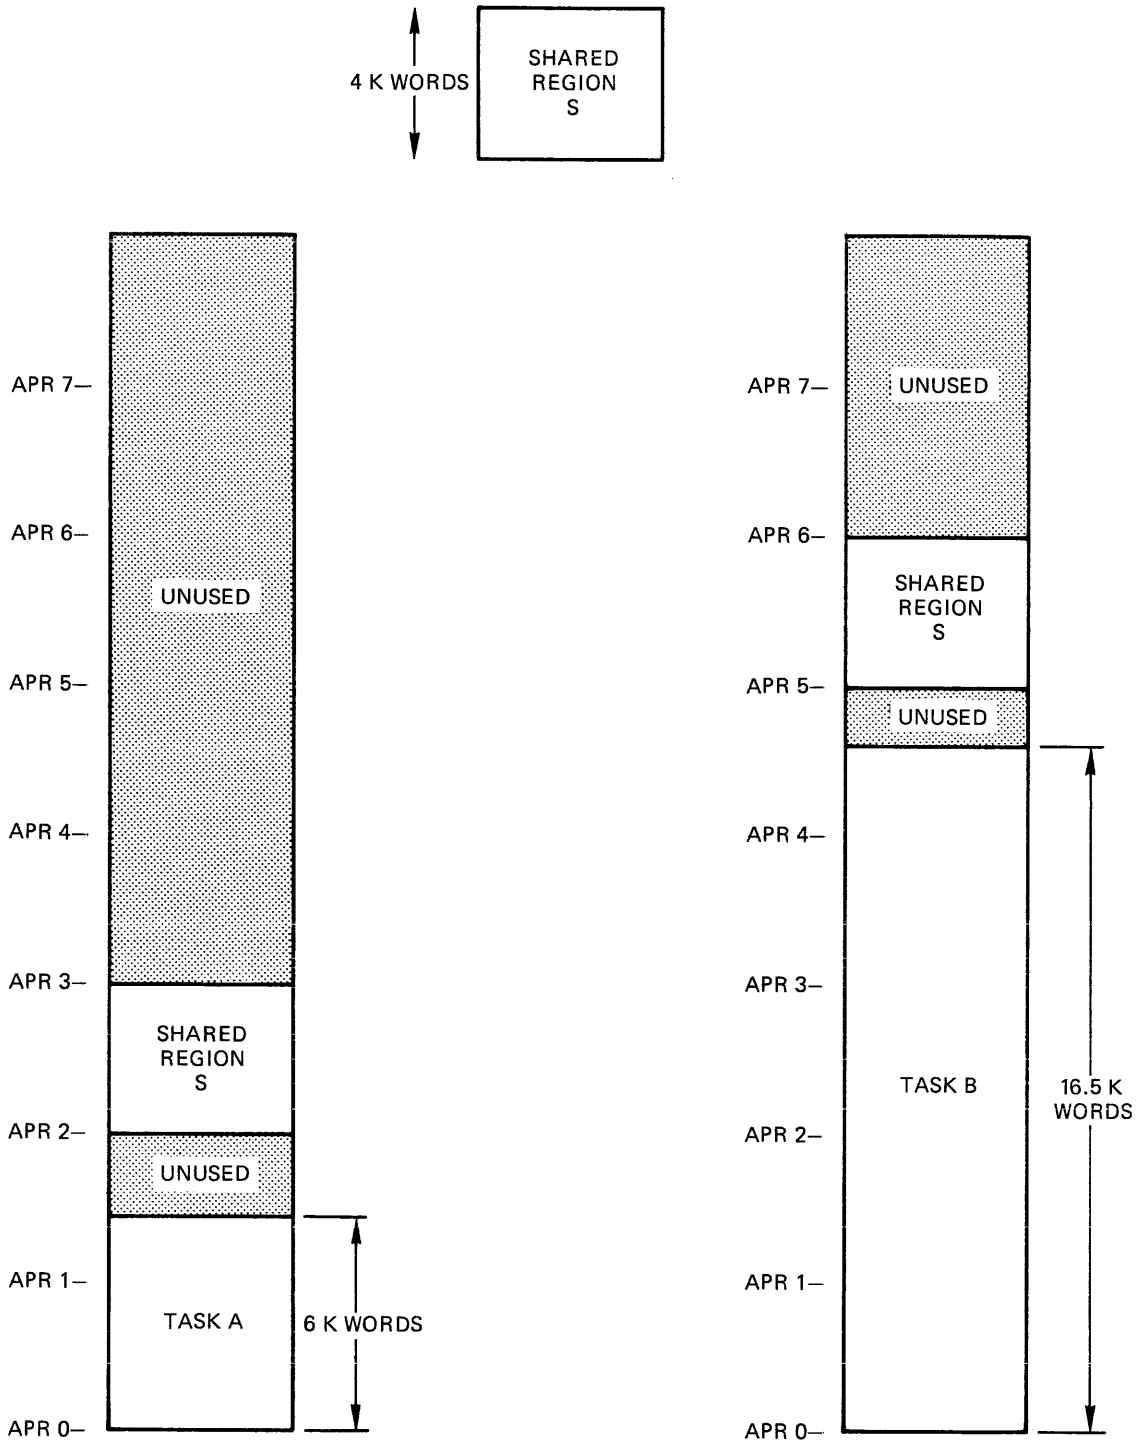
If you do not specify which APR is to be used to map a position-independent shared region, TKB selects the highest set of APRs available in the referencing task's virtual address space. In Figure 5-4, for example, if APR 2 in task A and APR 5 in task B had not been selected at task-build time, TKB would have selected APR 7 in both cases.

5.1.2.2 Specifying a Position-Independent Region - You specify that a shared region is position independent when you build it by attaching the /PI switch to the image file specification for the region. (Refer to Chapter 10 for a description of the /PI switch.)

You should declare a region position independent if:

- The region contains code that will execute correctly regardless of its location in the address space of the referencing task.
- The region contains data that is not address dependent.
- The region contains data that will be referenced by a FORTRAN program (such data must reside in a named common).

SHARED REGION CONCEPTS AND EXAMPLES



ZK-421-81

Figure 5-4 Specifying APRs for a Position-Independent Shared Region

SHARED REGION CONCEPTS AND EXAMPLES

Program section names are preserved in some shared regions. All the following switch combinations produce shared regions in which PSECT names are preserved:

`/PI/CO`, `/-PI/CO`, and `/PI`

Therefore, you should observe the following precautions when building and referring to these regions:

- No code or data in the region should be included in the blank (`. BLK.`) program section.
- No code or data in a referencing task should appear in a program section of the same name as a program section in the shared region.
- The order in which memory is allocated to program sections (alphabetic or sequential) must be the same for the shared region and its referencing tasks. (Chapter 2 describes alphabetic ordering of program sections. Refer to the description of the `/SQ` and `/SG` switches in Chapter 10 for an explanation of sequential ordering of program sections.)

5.1.3 Absolute Shared Regions

When a shared region is absolute, you select the virtual addresses for it when you build it. Thus, an absolute shared region is fixed in the virtual address space of all tasks that refer to it.

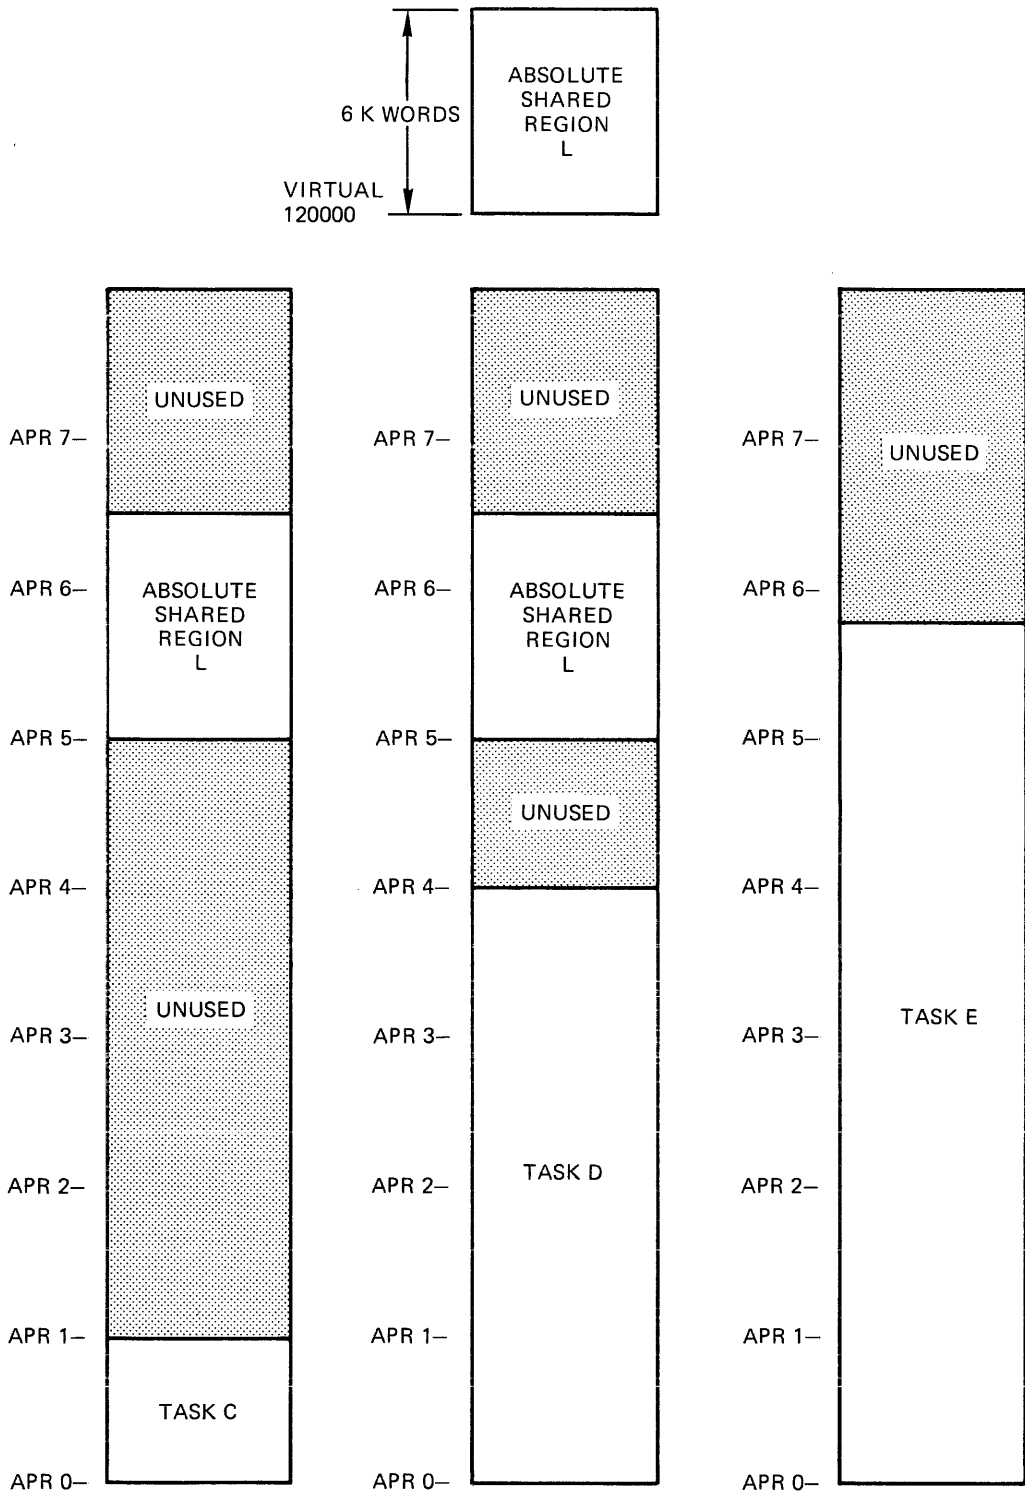
5.1.3.1 Absolute Shared Region Mapping - Figure 5-5 shows three tasks (task C, task D, and task E) and a single absolute shared region, L. The absolute shared region L is 6K words long and is built to occupy virtual addresses 120000(octal) to 150000(octal). These addresses correspond to APR 5 and APR 6, respectively. Tasks C and D can be linked to region L because at the time they are built APR 5 and APR 6 are unused in both tasks. However, task E is 23K words long and even though it has 8K words of virtual address space available to map the shared region, APR 5 (which corresponds to virtual address 120000, the base address of the shared region) has been allocated to the task region. If shared region L were position independent, task E could be linked to it.

5.1.3.2 Specifying an Absolute Shared Region - You specify that a shared region is absolute when you build it by using the `/-PI` switch or omitting the `/PI` switch from the task image file. You establish the virtual address for the region by specifying the base address of the region as a parameter of the `PAR` option.

You should build an absolute shared region if:

- The region contains code that must appear in a specific location in the address space of a referencing task.
- The region contains data that is address dependent.
- The region contains program sections of the same name as program sections in referencing tasks.

SHARED REGION CONCEPTS AND EXAMPLES



ZK-422-81

Figure 5-5 Mapping for an Absolute Shared Region

SHARED REGION CONCEPTS AND EXAMPLES

5.1.3.3 Absolute Shared Region .STB File - When a shared region is created with the `/-PI/LI` or `/-PI` switches, or just the `/-HD` switch, the only program section name that appears in the `.STB` file for the region is the absolute program section name (`.ABS.`). TKB includes in the `.STB` file for the region each symbol in the region and its value. But, because TKB does not include the program section names of an absolute shared region in its `.STB` file, all code or data in the region must be referred to by global symbol name. Also, because the program section names are not in the `.STB` file, TKB places no restrictions on the way the program sections are ordered in either the absolute shared region or the tasks that reference it. You can order program sections the way you want by using the `/SQ` and `/SG` switches described in Chapter 10.

5.1.4 Shared Regions with Memory-Resident Overlays

Shared regions with memory-resident overlays are a primary resource for conserving memory. If the shared region is larger than the available virtual address space in a task that must reference the region, you can build the region -- both position-independent and absolute -- with memory-resident overlays. All segments of the overlay structure are included in the shared region, but each task referencing the shared region can include only part of the shared region -- that is, an overlay segment or series of segments in an overlay path -- in its virtual address space. Therefore, the task need only have enough virtual address space for the largest shared region overlay segment, or series of segments in an overlay path, it is likely to access. Hence, the virtual address space of the task can be considerably smaller than the size of the shared region.

5.1.4.1 Considerations About Building an Overlaid Shared Region - In general, overlays can be disk-resident or memory-resident, but those in shared regions must, by their very nature, be memory-resident. TKB marks each overlay segment in the shared region with the `NODSK` attribute to suppress overlay load requests. When you build a shared region with memory-resident overlays, you must define the overlay structure through a conventional `ODL` file. (See Chapters 3 and 4 of this manual for information on overlays and the Overlay Description Language.) TKB does not include the overlay data base (segment descriptors, autoload vectors, and so forth) or the overlay run-time routines within the region image. Instead, this data base becomes a part of the `.STB` file that is linked to the referencing task. When this task is built, its root segment automatically includes both the data base and global references to overlay support routines residing in the system object module library.

The procedure for creating a shared region with memory-resident overlays can be summarized as follows:

- Define an overlay structure containing only memory-resident overlays.
- Include the `GLBREF` option, or provide in the root segment a module containing the appropriate global references for defining entry points within those overlay segments. TKB generates autoload vectors and global definitions for the overlay segments.

SHARED REGION CONCEPTS AND EXAMPLES

5.1.4.2 **Example of Building a Memory-Resident Overlaid Shared Region**
- The procedure for creating a shared region is illustrated in the following example. The shared region to be constructed consists of reentrant code that resides within the overlay structure defined below:

```
.ROOT A-!(B,C-D)
.NAME A
.END
```

Root segment A contains no code or data and has a length of 0. All executable code exists within memory-resident overlay segments composed of object modules B.OBJ, C.OBJ, and D.OBJ, containing global entry points B, C, and D, respectively.

You generate the .TSK, .MAP, and .STB files by using the following TKB commands:

```
TKB>A/-HD/MM,LP:,SY:A=A/MP
Enter Options:
TKB>GBLREF=B,C,D
TKB>PAR=A:160000:20000
TKB>STACK=0
TKB>/
```

NOTE

When building a shared region, you must use the same name for the partition and the .TSK and .STB files.

See the PAR, RESLIB, LIBR, RESCOM, and COMMON options in Chapter 11.

TKB inserts references to entry points B, C, and D in the root segment of the library which subsequently appear in the .STB file as definitions.

TKB resolves the definitions for symbol C directly to the actual entry point. TKB resolves the definitions for symbols B and D to autoload vectors that it includes in each referencing task.

5.1.4.3 **Options for Use in Overlaid Shared Regions** - Certain options may prove useful to you when building and linking shared regions to a task. They are described next.

GBLDEF -- You can declare the definition of a symbol by means of the GBLDEF option. The syntax of this option is:

```
GBLDEF= symbol-name:symbol-value
```

where symbol-name is a 1- to 6- character Radix-50 name of the defined symbol and symbol-value is an octal number in the range of 0 through 177777 assigned to the symbol. This option is frequently used in the TKB build file for a task or shared region to allow you to alter the value of a global symbol that resides in a module. This saves you the trouble of reassembling the source code for a module if changes are necessary.

SHARED REGION CONCEPTS AND EXAMPLES

GBLINC -- By means of this option, you force TKB to include the specified symbols in the .STB file being created by the linking process in which this option appears. The syntax of this option is:

```
GBLINC=symbol-name,symbol-name,...,symbol-name
```

where symbol-name is the symbol or symbols to be included. Use this option when you want to force particular modules to be linked to the task that references this library. The global symbol references specified by this option must be satisfied by some module or GBLDEF specification when you build the task.

GBLREF -- You can force the inclusion of a global reference in the root segment of the shared region by means of the GBLREF option. In this way, the necessary autoload vectors and definitions can be generated without explicitly including such references in an object module. The syntax of the option is:

```
GBLREF=[,name[,name...]]
```

where the name consists of from one to six Radix-50 characters. If the definition resides within an autoloadable segment, TKB constructs an autoload vector and includes it in the symbol definition file. If the definition is not autoloadable, TKB obtains the real value and defines it in the root segment. No global symbol appears in the .STB file unless the symbol is either defined in the root segment or is referenced in the root segment and defined elsewhere in the overlay structure.

GBLXCL -- You can exclude a symbol or symbols from the symbol definition file of a shared region by means of the GBLXCL option. The syntax of this option is:

```
GBLXCL=symbol-name,symbol-name,...,symbol-name
```

where symbol-name is the symbol or symbols to be excluded. You can use this option when you do not want the task to be aware of specific symbols within the library. This option is particularly useful when you cluster overlaid libraries together (see the CLSTR option in Chapter 11 and the Cluster Libraries section in this chapter).

5.1.4.4 Autoload Vectors and .STB Files for Overlaid Shared Regions -

When TKB builds a task image file containing memory-resident overlays, TKB allocates autoload vectors in the task image. If the task links to a shared region, autoload vectors for the shared region are also allocated in the task image. TKB allocates the autoload vectors in the task's root segment, but not in the shared region. Therefore, the shared region cannot reference unloaded (unmapped) segments of its overlay structure.

When the task executes, the shared region is effectively part of the task. In fact, when the task loads overlay segments, it makes no distinction between overlay segments of the task and those of the shared region. They are loaded as needed in a procedure that is transparent insofar as the execution of the task is concerned.

For the Fast Task Builder (FTB) and older versions of TKB that do not support overlaid I- and D-space tasks, each autoload vector in the shared region's .STB file is allocated in the root of the task being linked to the region, whether or not the entry point is referenced by the task.

SHARED REGION CONCEPTS AND EXAMPLES

However, if you use a version of TKB that supports overlaid I- and D-space tasks and the library was built with one of these versions, TKB allocates autoload vectors in the root of the task only for those autoloadable entry points in the library that the task references. The .STB file contains ISD records that allow TKB to create dynamically autoload vectors when linking the task to the library. TKB ignores the autoload vectors in the .STB file if the ISD records are present. Therefore, tasks that link to overlaid shared regions and that are built with newer versions of TKB tend to be smaller and use less virtual address space than those that are built by FTB or older versions of TKB.

NOTE

Libraries created with older versions of TKB do not have the ISD records in the .STB file that newer versions of TKB use to include autoload vectors in the task from the .STB file. Therefore, TKB must create autoload vectors for every entry point in the library.

If you are using one of these older libraries and you are linking an I- and D-space task to it, TKB will give you the fatal error message:

```
"Module      module-name      contains
incompatible autoload vectors."
```

This message occurs because the .STB file contains conventional autoload vectors that are not usable by an I- and D-space task.

Only those global symbols defined or referenced in the root segment of the shared region appear in the .STB file. The .STB file also contains the data base required by the overlay run-time system in relocatable object module format. This data base includes:

- All autoload vectors
- Segment tables (linked as described in Appendix B)
- Window descriptors
- A single region descriptor

The overlay structure, as reflected in the segment table linkage, is preserved and conveyed to the referencing task by the .STB file. Thus, path loading for the shared region can occur exactly as it does within a task. Aside from address space restrictions, there are no limitations on the overlay structures that can be defined for a shared region.

SHARED REGION CONCEPTS AND EXAMPLES

5.1.5 Run-Time Support for Overlaid Shared Regions

Memory-resident overlays within a shared region require little additional support from the overlay run-time system. The shared region overlay data base that is linked within the image of the referencing task has a structure that is identical to the equivalent data created for an overlaid task. Therefore, memory-resident overlays within the shared region are indistinguishable from memory-resident overlays that form a part of the task image. The only additional processing is that required to attach the shared region and obtain its identification for use by the mapping directives.

Once this initialization is complete, all further processing is identical to memory-resident overlay processing performed on task overlays.

Several restrictions apply to shared regions existing as memory-resident overlays:

- A shared region cannot use the autoloading facility to reference memory-resident overlays within itself or any other region. If each segment is uniquely named, overlays can be mapped through the manual load facility.
- Named program sections in a shared region overlay segment cannot be referenced by the task. If reference to the storage is required, such sections must be included in the root segment of the region (with resultant loss of virtual address space).
- For FTB, and libraries built with versions of TKB that do not support I- and D-space overlaid tasks, the number of autoloading vectors is independent of the entry points actually referenced. The maximum number of vectors will be allocated within each referencing task. In some cases the size of the allocation will be large.
- There is an overhead of six instructions per autoloading call, even when the segment is mapped. The overhead is seven instructions for an overlaid I- and D- space task.

As implied by the previous items, great care must be exercised if an efficient memory-resident overlay structure for library routines such as the FORTRAN IV OTS is to be implemented.

5.1.6 Linking to a Shared Region

When you build a task that links to a shared region, you must indicate to TKB the name of the shared region and the type of access the task requires to it (read/write or read-only). In addition, if the shared region is position independent, you can specify which APR TKB is to allocate for mapping the region into the task's virtual address space. Four options are available for this action:

- RESLIB (resident library)
- RESCOM (resident common)
- LIBR (system-owned resident library)
- COMMON (system-owned resident common)

SHARED REGION CONCEPTS AND EXAMPLES

RESLIB and RESCOM accept a complete file specification as one of their arguments. Thus, you can specify a device and UFD indicating to TKB the location of the region's image file and, by implication, its symbol definition file. (Refer to Chapter 1 for more information on file specifications and defaults.)

LIBR and COMMON accept a 1- to 6-character name. When you specify either of these options, the shared region's image file and symbol definition file must reside under UFD [1,1] on device LB0:.

The RESLIB and RESCOM options require that all users of the shared region know the UFD under which the shared region's image file and .STB file reside. The LIBR and COMMON options require only that the users of the shared region know the name of the shared region. When you specify either LIBR or COMMON, by default, TKB expects to find the shared region's image and .STB files on device LB: under UFD [1,1].

All four options accept two additional arguments:

- The type of access that the task requires (RO or RW).
- The first APR that TKB is to allocate for mapping the region into the task's virtual address space. As stated earlier, this argument is valid only when the shared region is position independent.

When you specify any of these options, TKB expects to find a symbol definition file of the same name as that of the shared region, but with an extension of .STB, on the same device and under the same UFD as those of the shared region's image file.

The syntax of these options is given in Chapter 11.

When TKB builds a task, it processes first any options that appear in the TKB command sequence. When TKB processes one of the four options above, it locates the disk image of the shared region named in the option. The disk image of a shared region does not have a header, but it does have a label block that contains the allocation information about the shared region (for example, its base address, load size, and the name of the partition for which it was built). TKB extracts this data from the shared region's label block and places it in the LIBRARY REQUEST section of the label block for the referencing task.

The .STB file associated with the shared region is an object module file. TKB processes it as an input file. If the shared region is position independent, its .STB file contains program section names, attributes, and lengths. However, the program section names are flagged within the file as "library" program sections and TKB does not add their allocations to the task image it is building.

If the task links to only one shared region, and if neither the shared region nor the task that links to it contain memory-resident overlays, the Task Builder allocates two window blocks in the header of the task. (Overlays are described in Chapter 3.) When the task is installed, the INSTALL processor will initialize these window blocks as follows:

- Window block 0 will describe the range of virtual addresses (the window) for the task region.
- Window block 1 will describe the window for the shared region.

Figure 5-6 shows the window-to-region relationship of such a task.

SHARED REGION CONCEPTS AND EXAMPLES




Figure 5-6 Windows for Shared Region and Referencing Task

SHARED REGION CONCEPTS AND EXAMPLES

A shared region need not be installed before a task that links to it is built. The .STB file that you specify when you build the shared region contains all the information required by TKB to resolve references from within a task to locations within the shared region. The only requirement is that you install a shared region before you install a task that links to it.

Unless you use the /LI switch, there is a restriction on the way TKB processes tasks that link to relocatable shared regions. TKB places all program section names into its internal control section table. The program section names include those from the .STB file of the shared region as well as those from the other input modules. A conflict can arise when building a task that contains program sections of the same name as those in the shared region to which the task links. The conflict arises because TKB tries to add the program section allocation in the task to the already existing allocation for the program section of the same name in the region. This is not possible because the region's image has already been built, is outside the address space of the task currently being built, and cannot be modified. Therefore, to avoid this conflict, the program section names within a task that links to a relocatable shared region must normally be unique with respect to program section names within the shared region.

TKB displays an error message under the following conditions:

- A program section in the task and a program section in the shared region have the same name.
- The program section in the task contains data.
- TKB tries to initialize the program section in the task.

The error message occurs when TKB tries to store data in an image outside the address limits of the task it is building. If this conflict occurs, TKB prints the following message:

```
TKB--*DIAG*-load addr out of range in module module-name
```

One exception to the above restriction develops when all of the following conditions exist:

- Both program sections (in the shared region and in the referencing task) have the (D) data and the OVR (overlay) attributes.
- The program section in the task is equal to or shorter than the program section in the shared region.
- The program section in the task does not contain data.

When all of these conditions exist, there is nothing to be initialized within the shared region. TKB binds the base address of the program section in the task to the base address of the program section in the shared region. If the program section in the task contains global symbols, TKB assigns addresses to them that reflect their location relative to the beginning of the program section. You can use this technique to establish symbolic offsets into resident commons. Examples 5-1 and 5-2 in the following sections illustrate how to establish these offsets.

SHARED REGION CONCEPTS AND EXAMPLES

5.1.7 Number and Size of Shared Regions

The number of shared regions to which a task can link is a function of the number of window blocks required to map the task and the regions. In an RSX-11M operating system, if a task is 4K words or less, and each shared region to which the task links is 4K words or less, then a nonprivileged task can access as many as seven shared regions.

In an RSX-11M-PLUS operating system, if a task is 4K words or less, and each shared region to which the task links is 4K words or less, a nonprivileged task can refer to as many as 15 shared regions: 7 in user mode and 8 in supervisor mode. (Supervisor-mode libraries are described in Chapter 8.)

5.1.8 Example 5-1: Building and Linking to a Common in MACRO-11

The text in this section and the figures associated with it illustrate the development of a MACRO-11 position-independent resident common and the development of two MACRO-11 tasks that share the common. The steps in building a position-independent common can be summarized as follows:

1. You create a source file that allocates the amount of space required for the common. In MACRO-11, either of the assembler directives, `.BLKB` or `.BLKW`, provide the means of allocating this space.
2. You assemble the source file.
3. You build the assembled module, specifying both a task image file and a symbol definition file.

You specify the `/-HD` (no header) switch and declare the common with `/CO`. You specify the common to be position independent with the `/PI` switch.

Under options you specify:

```
STACK=0
PAR=parname
```

The `parname` in this `PAR` option is the name of the partition in which the common is to reside. (The `/HD` and `/PI` switches and the `STACK` and `PAR` options are described in Chapter 10.)

If your system is an RSX-11M system, the common must reside within a common partition of the same name as that of the common.

If your system is an RSX-11M-PLUS system, the common can reside within any partition large enough to hold it.

4. You install the common.

Example 5-1 below shows a MACRO-11 source file that, when assembled and built, creates a position-independent resident common area named `MACCOM`. The common area consists of two program sections named `COM1` and `COM2`, respectively. Each program section is 512(decimal) words long.

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-1, Part 1 Common Area Source File in MACRO-11

```
                .TITLE MACCOM
;
;                COM1 - 512 WORDS
;                COM2 - 512 WORDS
;
                .PSECT COM1,RW,D,GBL,REL,OVR
                .BLKW 512.
                .PSECT COM2,RW,D,GBL,REL,OVR
                .BLKW 512.

                .END
```

Once this common has been assembled, the Task Builder command sequence shown below can be used to build it.

```
>TKB
TKB>MACCOM/PI/-HD/CO,MACCOM/-SP,MACCOM=MACCOM
TKB>/
Enter Options:
TKB>STACK=0
TKB>PAR=MACCOM:0:4000
TKB>//
```

This command sequence directs TKB to build a position-independent (/PI), headerless (/HD) common image file named MACCOM.TSK. It also specifies that the Task Builder is to create a map file, MACCOM.MAP, and a symbol definition file, MACCOM.STB. TKB creates all three files -- MACCOM.TSK, MACCOM.MAP, and MACCOM.STB -- on device SY: under the UFD that corresponds to the terminal UIC. Because /-SP is attached to the map file, TKB will not spool a map listing to the line printer.

Under options, STACK=0 suppresses the stack area in the common's image. The PAR option specifies that the common area will reside within a common partition of the same name as that of the common, MACCOM. In addition, the parameters in the PAR option specify a base of 0 and a length of 4000 octal bytes for the common. (Refer to Chapters 10 and 11 for descriptions of the switches and options used in this example.)

Example 5-1, Part 2 shows the map resulting from this command sequence.

The task attributes section of this map reflects the switches and options of the command string. It indicates that the common resides in a partition named MACCOM, that it was built under terminal UIC [7,62], that it is headerless and position independent, and that it requires one window block to map. The total length of the common is 1024(decimal) words and its address limits range from 0 to 3777(octal). The common image (that portion of the disk image file that eventually will be read into memory) begins at file-relative disk block 2 ①. The last block in the file is file-relative disk block 5 ② and the common image is four blocks long ③.

The memory allocation synopsis details the Task Builder's allocation for and the attributes of the program sections within the common. For example, reading from left to right, the map indicates that the

SHARED REGION CONCEPTS AND EXAMPLES

program section COM1 permits read/write access, that it contains data, and that its scope is global. It also indicates that COM1 is relocatable and that all contributions to COM1 are to be overlaid. Because COM1 has the overlay attribute, the total allocation for it will be equal to the largest allocation request from the modules that contribute to it. (For more information on program section attributes, see Chapter 2.)

Continuing to the right, the first 6-digit number is COM1's base address, which is 0 ④. The next two digits are its length (bytes) in octal and decimal, respectively.

The next line down lists the first object module that contributes to COM1. In this case there is only one: the module MACCOM from the file MACCOM.OBJ;1. The numbers on this line indicate the relative base address of the contribution and the length of the contribution in octal and decimal ⑤. If there had been more than one module input to TKB that contained a program section named COM1, TKB would have listed each module and its contribution in this section.

Notice that there is a program section named .BLK. shown on the map just above the field for COM1. This is the "blank" program section that is created automatically by the language translators. The attributes shown are the default attributes. The allocation for .BLK. is 0 because the program sections in MACCOM were explicitly declared. If the program sections had not been explicitly declared, all of the allocation for the common would have been within this program section.

Example 5-1, Part 2 Task Builder Map for MACCOM.TSK

MACCOM.TSK;1 Memory allocation map TKB M40.10 Page 1
17-NOV-82 16:05

Partition name : MACCOM
 Identification :
 Task UIC : [7,62]
 Task attributes: -HD,PI
 Total address windows: 1.
 Task image size : 1024. WORDS
 Task address limits: 000000 003777
 R-W disk blk limits: 000002 000005 000004 000004.

*** Root segment: MACCOM ① ② ③

R/W mem limits: 000000 003777 004000 02048.
 Disk blk limits: 000002 000005 000004 000004.

Memory allocation synopsis:

Section	Title	Ident	File
.BLK.:(RW,I,LCL,REL,CON)	000000	000000	00000.
COM1 : (RW,D,GBL,REL,OVR)	000000	002000	01024.
	000000	002000	01024.
COM2 : (RW,D,GBL,REL,OVR)	002000	002000	01024.
	002000	002000	01024.

⑥ ④ ⑤

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

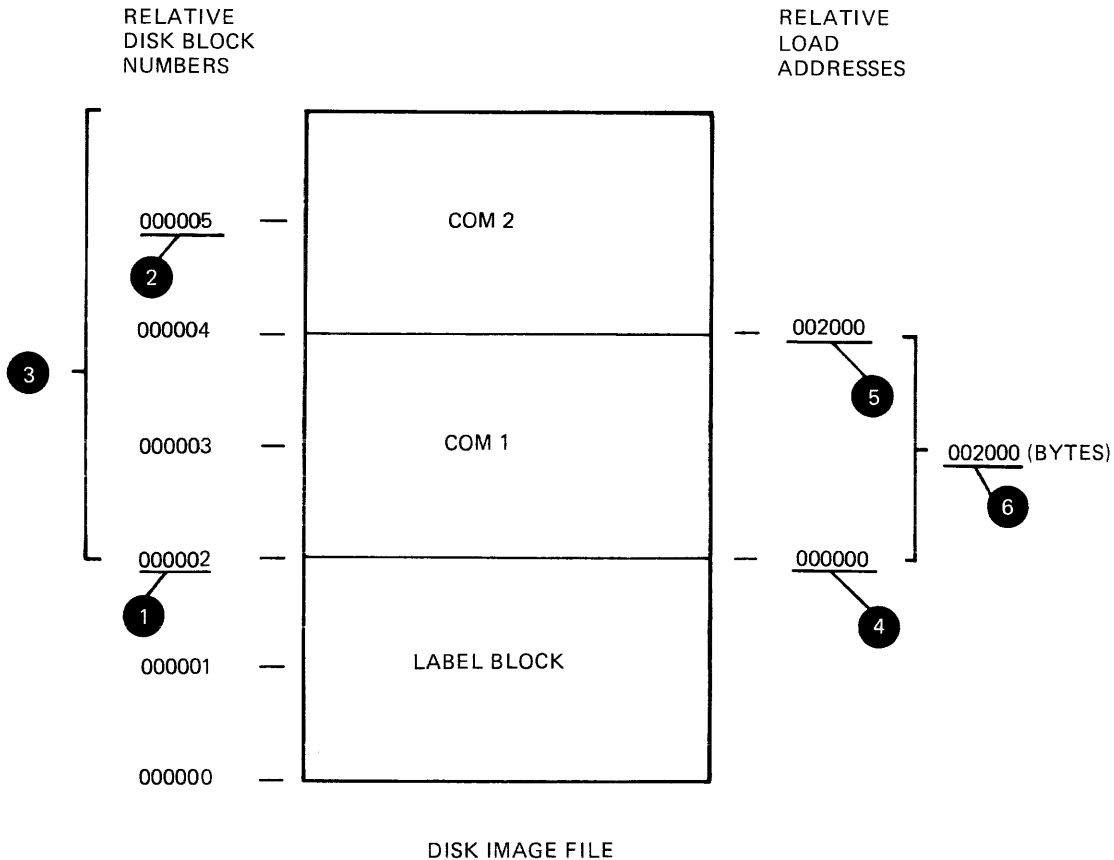
Example 5-1, Part 2 (Cont.) Task Builder Map for MACCOM.TSK

*** Task builder statistics:

Total work file references: 183.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 7086. WORDS (27. PAGES)
 Size of work file: 768. WORDS (3. PAGES)

Elapsed time:00:00:05

Figure 5-7 is a diagram that represents the disk image file for MACCOM. The circled numbers in Figure 5-7 correspond to the circled numbers in Example 5-1, Part 2.



ZK-424-81

Figure 5-7 Allocation Diagram for MACCOM.TSK

Once you have built MACCOM, you can install it. If your system is an RSX-11M system, the common is loaded into memory when you install it.

If your system is an RSX-11M-PLUS system, it remains there until you explicitly remove it with the MCR command REMOVE. The common will not be loaded until either one of the following occurs:

- A task that is linked to it is run.
- You explicitly fix the common in memory with the MCR command FIX.

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-1, Parts 3 and 4 show two programs: MCOM1 and MCOM2, respectively. Both of these programs reference the common area MACCOM created above. MCOM1 in Example 5-1, Part 3 accesses the COM1 portion of MACCOM. It inserts into the first 10 words of COM1 the numbers 1 through 10 in ascending order. It then issues an Executive directive request for the task MCOM2 and suspends itself.

When MCOM2 runs, it adds together the integers left in COM1 by MCOM1 and leaves the sum in the first word of COM2. It then issues a resume directive for MCOM1 and exits.

When MCOM1 resumes, it retrieves the answer left in COM2 and calls the system library routine \$EDMSG (edit message) to format the answer for output to device TI:.

All of the Executive directives for both programs (RQST\$C, SPND\$\$, QIOW\$\$, RSUM\$C, and EXIT\$\$) are documented in the RSX-11M-PLUS Executive Reference Manual. The system library routine \$EDMSG is documented in the IAS/RSX-11 System Library Routines Reference Manual.

Example 5-1, Part 3 MACRO-11 Source Listing for MCOM1

```
.TITLE MCOM1
.IDENT /01/

.MCALL EXIT$$,SPND$$,RQST$C,QIOW$$

OUT: .BLKW 100. ; SCRATCH AREA
FORMAT: .ASCIZ /THE RESULT IS %D./
MES: .ASCII /ERROR FROM REQUEST/
LEN = . - MES
.EVEN

; PSECT - COM1 IS USED TO ACCESS THE FIRST 512. WORDS OF THE
; COMMON.

.PSECT COM1,GBL,OVR,D
INT: .BLKW 10.

; PSECT - COM2 IS USED TO ACCESS THE SECOND 512. WORDS OF THE
; COMMON. IT WILL CONTAIN THE RESULT

.PSECT COM2,GBL,OVR,D
ANS: .BLKW 1

.PSECT
START:
MOV #10.,R0 ; NUMBER OF INTEGERS TO SUM
MOV #1,R1 ; START WITH A 1
MOV #INT,R3 ; PLACE VALUES IN 1ST 10 WORDS
; OF COMMON
```

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-1, Part 3 (Cont.) MACRO-11 Source Listing for MCOM1

```

10$:  MOV      R1,(R3)+          ; INITIALIZE COMMON
      INC      R1              ; NEXT INTEGER
      DEC      R0              ; ONE LESS TIME
      BNE     10$              ; TO INITIALIZE
      RQST$C  MCOM2           ; REQUEST THE SECOND TASK
      BCS     ERR1            ; REQUEST FAILED
      SPND$$  ;              ; WAIT FOR MCOM2 TO SUM THE INTEGERS
      MOV     #OUT,R0          ; ADDRESS OF SCRATCH AREA
      MOV     #FORMAT,R1       ; FORMAT SPECIFICATION
      MOV     #ANS,R2          ; ARGUMENT TO CONVERT
      CALL    $EDMSG           ; DO CONVERSION
      QIOW$$  #IO.WVB,#5,#1,,, <#OUT,R1,#40>
      EXIT$$

ERR1:  QIOW$$  #IO.WVB,#5,#1,,, <#MES,#LEN,#40>
      EXIT$$
      .END    START

```

Example 5-1, Part 4 MACRO-11 Source Listing for MCOM2

```

      .TITLE  MCOM2
      .IDENT  /01/

      .MCALL  EXIT$$,QIOW$$,RSUM$C

MES:   .ASCII  /ERROR FROM RESUME/
      LEN = . - MES
      .EVEN

;      PSECT - COM1 IS USED TO ACCESS THE FIRST 10. WORDS OF THE
;      COMMON.

      .PSECT  COM1,GBL,OVR,D
INT:   .BLKW  10.

;      PSECT - COM2 IS USED TO ACCESS THE SECOND 10. WORDS OF THE
;      COMMON. IT WILL CONTAIN THE RESULT

      .PSECT  COM2,GBL,OVR,D
ANS:   .BLKW  1

      .PSECT

START:
      MOV     #10.,R0          ; NUMBER OF INTEGERS TO SUM
      MOV     #INT,R3         ; PLACE VALUES IN 1ST 10 WORDS
                                ; OF COMMON
10$:   CLR     ANS              ; INITIALIZE ANSWER
      ADD     (R3)+,ANS        ; ADD IN VALUES
      DEC     R0              ; ONE LESS VALUE
      BNE     10$              ; TO SUM

      RSUM$C  MCOM1           ; RESUME MCOM1
      BCS     ERR              ; RESUME FAILED
      EXIT$$

ERR:   QIOW$$  #IO.WVB,#5,#1,,, <#MES,#LEN,#40>
      EXIT$$
      .END    START

```

SHARED REGION CONCEPTS AND EXAMPLES

Note that both tasks MCOM1 and MCOM2 contain .PSECT declarations establishing program section names that are the same as program section names within the position-independent common to which the task is linked (MACCOM). As stated earlier, in most circumstances this would be illegal. In this application, however, the .PSECT directives have been placed into the tasks to establish symbolic offsets in the resident common. When either task is built, TKB assigns to the symbol INT: the base address of program section COM1, and to the symbol ANS: the base address of program section COM2. Figure 5-8 illustrates this assignment.

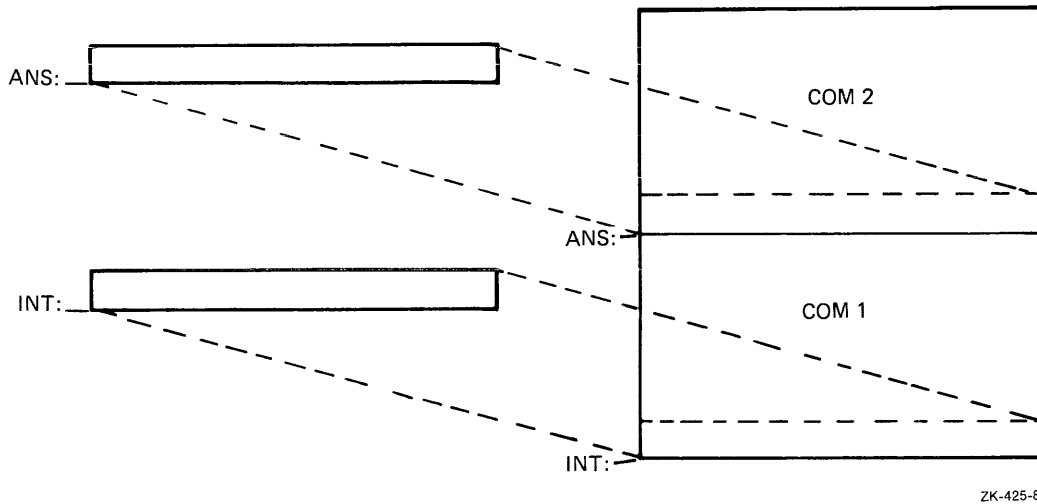


Figure 5-8 Assigning Symbolic References within a Common

Once you have assembled MCOM1 and MCOM2, you can build them with the following TKB command sequences:

```
>TKB
TKB>MCOM1,MCOM1/-SP=MCOM1
TKB>/
Enter Options:
TKB>RESCOM=MACCOM/RW
TKB>//
```

```
>TKB
TKB>MCOM2,MCOM2/-SP=MCOM2
TKB>/
Enter Options:
TKB>RESCOM=MACCOM/RW
TKB>//
```

Under options in both of these command sequences, the RESCOM option tells TKB that these programs intend to reference a common data area named MACCOM and that the tasks require read/write access to it. Because the RESCOM option is used, TKB expects to find the image file and the symbol definition file for the common on device SY: under the UFD that corresponds to the terminal UIC. In addition, because the optional APR specification was omitted from the RESCOM option, TKB allocates virtual address space for the common starting with APR7 in both tasks (the highest APR available in both tasks).

SHARED REGION CONCEPTS AND EXAMPLES

The TKB map for MCOM1 is shown in Example 5-1, Part 5. The map for MCOM2 is not essentially different from that of MCOM1 and is therefore not included here.

Example 5-1, Part 5 Task Builder Map for MCOM1.TSK

MCOM1.TSK;1 Memory allocation map TKB M40.10 Page 1
11-DEC-82 16:12

Partition name : GEN
 Identification : 01
 Task UIC : [7,62]
 Stack limits: 000274 001273 001000 00512.
 PRG xfr address: 001650
 Total address windows: 2.
 Task image size : 1184. WORDS
 Task address limits: 000000 004407
 R-W disk blk limits: 000002 000006 000005 00005.

*** Root segment: MCOM1

R/W mem limits: 000000 004407 004410 02312.
 Disk blk limits: 000002 000006 000005 00005.

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.:(RW,I,LCL,REL,CON) 001274 002664 01460.			
001274 000574 00380. MCOM	01	MCOM1.OBJ;1	
COM1 : (RW,D,GBL,REL,OVR) 160000 002000 01024.			
160000 000024 00020. MCOM	01	MCOM1.OBJ;1	
COM2 : (RW,D,GBL,REL,OVR) 162000 002000 01024.			
162000 000002 00002. MCOM	01	MCOM1.OBJ;1	
\$DPB\$\$:(RW,I,LCL,REL,CON) 004160 000016 00014.			
004160 000016 00014. MCOM	01	MCOM1.OBJ;1	
\$\$RESL:(RO,I,LCL,REL,CON) 004176 000212 00138.			

*** Task builder statistics:

Total work file references: 1924.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 7086. WORDS (27. PAGES)
 Size of work file: 1024. WORDS (4. PAGES)

Elapsed time:00:00:04

SHARED REGION CONCEPTS AND EXAMPLES

Note that TKB has placed two window blocks in MCOM1's header. When MCOM1 is installed, the INSTALL processor will initialize these window blocks as follows:

- Window block 0 will describe the range of virtual addresses (the window) for MCOM1's task region.
- Window block 1 will describe the window for the shared region MACCOM.

5.1.9 Linking Shared Regions Together

Shared regions can link to other shared regions. You may find it convenient to have code in a shared library and have access to routines in another shared library to which it links.

The following text describes, as an example for a mapped system, the TKB command sequence for building a resident library named FILEB. That text is followed by a TKB command sequence that shows an example of building another resident library named FORCOM that links to FILEB. Following after that, a TKB command sequence shows the building of a task that links to FORCOM. In the TKB command sequences to follow, it is assumed that you know the contents of the libraries and the task. The examples show the linkage only.

The first shared region to be built is called FILEB. The library FILEB is a position-dependent library. You use the /-PI switch to signify that the library is absolute. You build the library with the /-HD switch to indicate that the library has no header. The /LI switch indicates that FILEB is to be a shared library. The program section name of the library is .ABS, which is the only one in the library. FILEB is to be loaded into a user-controlled partition on a mapped system. The name of the partition in which FILEB resides has the same name, FILEB, that you specify in the PAR option. The PAR option also specifies the base address and the length of the partition. Because FILEB is absolute, a base address must be specified; here, the base address is 160000. The length in this example is 4K bytes. If neither the base nor the length is specified, TKB tries to determine the length. The TKB command sequence follows:

```
>TKB
TKB>FILEB/-PI/-HD/LI,FILEB/-SP,FILEB=FILEB.OBJ
TKB>/
Enter Options:
TKB>STACK=0
TKB>PAR=FILEB:160000:40000
TKB>//
```

The next TKB command sequence specifies a shared library called FORCOM. FORCOM links to the read-only library called FILEB. You build FORCOM with the /LI switch to specify a library to the Task Builder. FORCOM is relocatable. You specify in the RESLIB option that the resident library to which FORCOM links is called FILEB. The

SHARED REGION CONCEPTS AND EXAMPLES

access required is read-only, which /RO specifies in the RESLIB option line. The TKB command sequence follows:

```
>TKB
TKB>FORCOM/-HD/LI/PI,FORCOM/-SP,FORCOM=FORCOM.OBJ
TKB>/
Enter Options:
TKB>STACK=0
TKB>PAR=FORCOM:0:4000
TKB>RESLIB=FILEB/RO
TKB>//
```

The next TKB command sequence builds the task and specifies that the task links to the library called FORCOM. The RESLIB option line specifies the link to the resident library called FORCOM.

```
>TKB
TKB>FOTASK,FOTASK/-SP,FOTASK=FOTASK.OBJ
TKB>/
Enter Options:
TKB>RESLIB=FORCOM/RW
TKB>//
```

Build the libraries before you build the task, and install the libraries before you run or install the task. See Chapters 10 and 11 for a description of the /PI, /HD, /CO, and /LI switches and the PAR, RESCOM, and RESLIB options.

5.1.10 Example 5-2: Building and Linking to a Device Common in MACRO-11

A device common is a special type of common that occupies physical addresses on the I/O page. When mapped into the virtual address space of a task, a device common permits the task to manipulate peripheral device registers directly.

NOTE

Because any access to the I/O page is potentially hazardous to the running system, you must exercise extreme caution when working with device commons.

The remaining text in this section and the figures associated with it illustrate the development and use of a device common. Example 5-2, Part 1 shows an assembly listing for a position-independent device common named TTCOM. When installed, TTCOM will map the control and data registers of the console terminal. Its physical base address will be 777500.

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-2, Part 1 Assembly Listing for TTCOM

```
.TITLE TTCOM
.PSECT TTCOM,GBL,D,RW,OVR
.=.+60
$RCSR:: .BLKW 1
$RBUF:: .BLKW 1
$XCSR:: .BLKW 1
$XBUF:: .BLKW 1
.END
```

The PDP-11 Peripherals Handbook defines the control and data register addresses for the console terminal. In Example 5-2, Part 1, the register addresses and the symbol names that correspond to them are as follows:

Register	Address	Symbol
Keyboard Status	777560	\$RCSR
Keyboard Data	777562	\$RBUF
Printer Status	777564	\$XCSR
Printer Data	777566	\$XBUF

The double colon (::) following each symbol in Example 5-2, Part 1 establishes the symbol as global. The first symbol, RCSR, is offset from the beginning of TTCOM by 60(octal) bytes. Each symbol thereafter is one word removed from the symbol that precedes it. Thus, when TTCOM is installed at 777500, each symbol will be located at its proper address.

Once you have assembled TTCOM, you can build it using the following TKB command sequence:

```
>TKB
TKB>LB:[1,1]TTCOM/-HD/PI, LB:[1,1]TTCOM/-WI/SP, LB:[1,1]TTCOM=TTCOM
TKB>/
Enter Options:
TKB>STACK=0
TKB>PAR=TTCOM:0:100
TKB>//
```

This command sequence directs TKB to create a common image named TTCOM.TSK and a symbol definition file named TTCOM.STB. TKB places both files on device LB: under UFD [1,1]. The command sequence also specifies that TKB is to spool a map listing to the line printer. The /SP switch need not be present because it is the default. The /-WI switch specifies an 80-column line printer listing format.

NOTE

For the command sequence above to work in a multiuser protection system, it must be input from a privileged terminal.

The STACK=0 option suppresses the stack area in the common's image file. The PAR option specifies that the device common will reside within a partition of the same name as that of the common. As with the data common in Example 5-1 (Section 5.1.7), this is a requirement of the RSX-11M system; in an RSX-11M-PLUS system it is not. The PAR option also specifies that the base of the common is 0 and that it is 100(octal) bytes long.

SHARED REGION CONCEPTS AND EXAMPLES

The TKB map for TTCOM that results from the command sequence above is shown in Example 5-2, Part 2. The task attributes section of this map indicates that the common is position independent and that no header is associated with it. The common's image and symbol definition file reside on device LB: under UFD [1,1].

The map in Example 5-2, Part 2 shows the global symbols defined in the common with their relative offsets into the common region. You establish the virtual base address for the common and the virtual addresses for the symbols within it when you build the tasks that link to the common.

You establish the physical addresses for the common with the MCR command SET. The keyword that you use with the SET command depends on which system you are running. If your system is an RSX-11M system, use the command

```
>SET /MAIN=TTCOM:7775:1:DEV
```

If your system is one that uses 22-bit physical addresses, use the command

```
>SET /MAIN=TTCOM:177775:1:DEV
```

If your system is an RSX-11M-PLUS system, use the command

```
>SET /PAR=TTCOM:177775:1:DEV
```

These previous SET command sequences create a main partition named TTCOM that begins at physical address 777500 in 18-bit systems and physical address 1777750 in 22-bit systems. The partition is one 64-byte block long, (100(octal) bytes). The argument DEV identifies the partition type. With the common built and the partition for it created, you must install TTCOM in an RSX-11M system before using it. For example, use

```
>INS LB:[1,1]TTCOM
```

You can establish the partition for a device common at any time in both the RSX-11M and the RSX-11M-PLUS systems. Partitions created to accommodate a device common are not a system generation consideration because they represent areas of physical address space above memory and therefore cannot conflict with memory partitions.

Example 5-2, Part 2 Task Builder Map for TTCOM

```
TTCOM.TSK;1  Memory allocation map  TKB M40.10      Page 1  
              1-DEC-82   17:02
```

```
Partition name : TTCOM  
Identification :  
Task UIC       : [7,62]  
Task attributes: -HD,PI  
Total address windows: 1.  
Task image size : 32. WORDS  
TASK  
ATTRIBUTES  
SECTION
```

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-2, Part 2 (Cont.) Task Builder Map for TTCOM

Task address limits: 000000 000067
R-W disk blk limits: 000002 000002 000001 00001.

*** Root segment: TTCOM

R/W mem limits: 000000 000067 000070 00056.
Disk blk limits: 000002 000002 000001 00001.

Memory allocation synopsis:

Table with 4 columns: Section, Title, Ident, File. It lists memory segments for BLK., TTCOM, and .MAIN.

Global symbols:

\$RBUF 000062-R \$RCSR 000060-R \$XBUF 000066-R \$XCSR 000064-R

*** Task builder statistics:

Total work file references: 214.
Work file reads: 0.
Work file writes: 0.
Size of core pool: 6666. WORDS (26. PAGES)
Size of work file: 768. WORDS (3. PAGES)

Elapsed time:00:00:02

Example 5-2, Part 3 shows an assembly listing for a demonstration program named TEST. When built and installed, TEST will print the letters A through Z on the console terminal by directly accessing the console terminal status and data registers. It will access the status and data registers through the device common TTCOM.

Example 5-2, Part 3 Assembly Listing for TEST

.TITLE TEST
.IDENT /01/
.MCALL EXIT\$\$
START: MOV #15,R0 ; START WITH A CARRIAGE RETURN
CALL OUTBYT ; PRINT IT
MOV #12,R0 ; THEN A LINE FEED
CALL OUTBYT ; PRINT IT
MOV #101,R0 ; FIRST LETTER IS AN "A"
MOV #26.,R1 ; NUMBER OF LETTERS TO PRINT

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-2, Part 3 (Cont.) Assembly Listing for TEST

```
OUTPUT: CALL    OUTBYT    ; PRINT CURRENT LETTER
         DEC     R1       ; ONE LESS TIME
         BNE    OUTPUT    ; AGAIN
         MOV     #15,R0   ; ANOTHER CARRIAGE RETURN
         CALL   OUTBYT
         MOV     #12,R0   ; ANOTHER LINE FEED
         CALL   OUTBYT
EXIT$$
OUTBYT: TSTB    $XCSR    ; OUTPUT BUFFER READY?
         BPL    OUTBYT   ; IF NOT WAIT
         MOV    R0,$XBUF  ; MOVE CHARACTER TO OUTPUT BUFFER
         INC    R0       ; INITIALIZE NEXT LETTER
         RETURN
         .END    START
```

Once you have assembled TEST, you can build it with the following TKB command sequence:

```
>TKB
TKB>TEST,TEST/-WI/MA=TEST
TKB>/
Enter Options:
TKB>COMMON=TTCOM:RW:1
TKB>//
```

The COMMON option in this command sequence tells TKB that TEST intends to access the device common TTCOM and that TEST will have read/write access to it. It also directs TKB to reserve APR 1 for mapping the common into TEST's virtual address space.

The TKB map that results from the command sequence above is shown in Example 5-2, Part 4.

This map contains a global symbols section. TKB included it because the /MA switch was applied to the memory allocation file at task-build time. Note that the global symbols in this section, which were defined in TTCOM, now have virtual addresses assigned to them. The addresses assigned by TKB are the result of the APR 1 specification in the COMMON= keyword during the task build.

It is important to remember that programs like TEST, which access the I/O page, take complete control of the registers they reference. Therefore, coding errors in such programs can disable the devices they reference and can even make it impossible for the device drivers to regain control of the device. If this happens, you must reboot the system.

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-2, Part 4 Memory Allocation Map for TEST

TEST.TSK;1 Memory allocation map TKB M40.10 Page 1
1-DEC-82 17:03

Partition name : GEN
Identification : 01
Task UIC : [7,62]
Stack limits: 000274 001273 001000 00512.
PRG xfr address: 001274
Total address windows: 2.
Task image size : 384. WORDS
Task address limits: 000000 001377
R-W disk blk limits: 000002 000003 000002 00002.

*** Root segment: TEST

R/W mem limits: 000000 001377 001400 00768.
Disk blk limits: 000002 000003 000002 00002.

Memory allocation synopsis:

Section	Title	Ident	File
-----	-----	-----	-----
. BLK.: (RW,I,LCL,REL,CON)	001274 000100 00068.		
	001274 000100 00068.	.MAIN.	TEST.OBJ;1
TTCOM : (RW,D,GBL,REL,OVR)	200000 000070 00056.		
	200000 000070 00056.	TTCOM	TTCOM.STB;1

Global symbols:

\$RBUF 020062-R \$RCSR 020060-R \$XBUF 020066-R \$XCSR 020064-R

*** Task builder statistics:

Total work file references: 243.
Work file reads: 0.
Work file writes: 0.
Size of core pool: 6666. WORDS (26. pages)
Size of work file: 768. WORDS (3. pages)

Elapsed time:00:00:03

5.1.11 Example 5-3: Building and Linking to a Resident Library in MACRO-11

Resident libraries consist of subroutines that are shared by two or more tasks. When such tasks reside in physical memory simultaneously, resident libraries provide a considerable memory savings because the subroutines within the library appear in memory only once.

SHARED REGION CONCEPTS AND EXAMPLES

The text in this section and the figures associated with it illustrate the development and use of a resident library, called LIB.

Example 5-3, Part 1 shows five FORTRAN-callable subroutines:

- An integer addition routine, AADD
- An integer subtraction routine, SUBB
- An integer multiplication routine, MULL
- An integer division routine, DIVV
- A register save and restore coroutine, SAVAL

These subroutines are contained in a single source file, LIB.MAC. When assembled and built, they constitute an example of a resident library. FORTRAN-callable routines were used in this example so that the routines can be accessed by either FORTRAN or MACRO-11 programs.

Example 5-3, Part 1 Source Listing for Resident Library LIB.MAC

```
.TITLE LIB
.IDENT /01/

.PSECT AADD,RO,I,GBL,REL,CON

;** FORTRAN CALLABLE SUBROUTINE TO ADD TWO INTEGERS

AADD:: CALL    $SAVAL          ; SAVE R0-R5
        MOV     @2(R5),R0      ; FIRST OPERAND
        MOV     @4(R5),R1      ; SECOND OPERAND
        ADD     R0,R1          ; SUM THEM
        MOV     R1,@6(R5)      ; STORE RESULT
        RETURN                 ; RESTORE REGISTERS AND RETURN

.PSECT SUBB,RO,I,GBL,REL,CON

;** FORTRAN CALLABLE SUBROUTINE TO SUBTRACT TWO INTEGERS

SUBB:: CALL    $SAVAL          ; SAVE R0-R5
        MOV     @2(R5),R0      ; FIRST OPERAND
        MOV     @4(R5),R1      ; SECOND OPERAND
        SUB     R1,R0          ; SUBTRACT SECOND FROM FIRST
        MOV     R0,@6(R5)      ; STORE RESULT
        RETURN                 ; RESTORE REGISTERS AND RETURN

.PSECT MULL,RO,I,GBL,REL,CON

;** FORTRAN CALLABLE SUBROUTINE TO MULTIPLY TWO INTEGERS

(continued on next page)
```

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-3, Part 1 (Cont.) Source Listing for Resident Library LIB.MAC

```
MULL:: CALL    $SAVAL          ; SAVE R0-R5
        MOV     @2(R5),R0      ; FIRST OPERAND
        MOV     @4(R5),R1      ; SECOND OPERAND
        MUL     R0,R1          ; MULTIPLY
        MOV     R1,@6(R5)      ; STORE RESULT
        RETURN                 ; RESTORE REGISTERS AND RETURN
```

```
.PSECT DIVV,RO,I,GBL,REL,CON
```

```
;** FORTRAN CALLABLE SUBROUTINE TO DIVIDE TWO INTEGERS
```

```
DIVV:: CALL    $SAVAL          ; SAVE REGS R0-R5
        MOV     @2(R5),R3      ; FIRST OPERAND
        MOV     @4(R5),R1      ; SECOND OPERAND
        CLR     R2             ; LOW ORDER 16 BITS
        DIV     R1,R2          ; DIVIDE
        MOV     R2,@6(R5)      ; STORE RESULT
        RETURN                 ; RESTORE REGISTERS AND RETURN
```

```
.PSECT SAVAL,RO,I,GBL,REL,CON
```

```
**ROUTINE TO SAVE REGISTERS
```

```
$SAVAL::
        MOV     R4,-(SP)       ;SAVE R4
        MOV     R3,-(SP)       ;SAVE R3
        MOV     R2,-(SP)       ;SAVE R2
        MOV     R1,-(SP)       ;SAVE R1
        MOV     R0,-(SP)       ;SAVE R0
        MOV     12(SP),-(SP)   ;COPY RETURN
        MOV     R5,14(SP)      ;SAVE R5
        CALL    @(SP)+         ;CALL THE CALLER
        MOV     (SP)+,R0        ;RESTORE R0
        MOV     (SP)+,R1        ;RESTORE R1
        MOV     (SP)+,R2        ;RESTORE R2
        MOV     (SP)+,R3        ;RESTORE R3
        MOV     (SP)+,R4        ;RESTORE R4
        MOV     (SP)+,R5        ;RESTORE R5
        RETURN
        .END
```

Once you have assembled LIB, you can build it with the following TKB command sequence:

```
TKB>LIB/PI/-HD/LI,LIB/-WI,LIB=LIB
TKB>/
Enter Options:
TKB>STACK=0
TKB>PAR=LIB:0:200
TKB>//
```


SHARED REGION CONCEPTS AND EXAMPLES

This command sequence instructs TKB to build a position-independent (/PI), headerless (/HD) library image named LIB.TSK. It instructs TKB to create a map file LIB.MAP and to output an 80-column listing (/WI) to the line printer. It also specifies that TKB is to create a symbol definition file, LIB.STB. TKB creates all three files -- LIB.TSK, LIB.MAP, and LIB.STB -- on device SY: under the UFD that corresponds to the terminal UIC. The /LI and /PI switches used together cause TKB to name the program section LIB, which is the root segment of the library. LIB becomes the only named program section in the library.

If you used the command sequence above without the /LI switch, TKB would create a common by default.

The STACK=0 option suppresses the stack area within the resident library's image. The PAR option tells TKB that the resident library will reside within a partition of the same name as that of the library. AS with all shared regions, this is a requirement in an RSX-11M system; in an RSX-11M-PLUS system it is not. In addition, the PAR option specifies that the base of the library is 0 and that it is 200(octal) bytes long. (For more information on the switches and options used in this example, refer to Chapters 10 and 11.)

Example 5-3, Part 2 shows the TKB map that results from the command sequence above.

Note in the global symbols section of the map in Example 5-3, Part 2 that TKB has assigned offsets to the symbols for each library function. When the task that links to this library is built, TKB will assign virtual addresses to these symbols.

The program MAIN in Example 5-3, Part 3 exercises the routines in the resident library LIB.TSK. When you assemble and build it, MAIN will call upon the library routines to add, subtract, multiply, and divide the integers contained in the labels OP1 and OP2 within the program. MAIN will print the results of each operation to device TI:.

Example 5-3, Part 2 Task Builder Map for LIB.TSK

```
LIB.TSK;1      Memory allocation map  TKB M40.10      Page 1
              11-DEC-82   13:50
```

```
Partition name : LIB
Identification : 01
Task UIC       : [7,62]
Task Attributes: -HD,PI
Total address windows: 1.
Task image size : 64. WORDS
Task address limits: 000000 000163
R-W disk blk limits: 000002 000002 000001 000001.
```

```
*** Root segment: LIB
```

```
R/W mem limits: 000000 000163 000164 00116.
Disk blk limits: 000002 000002 000001 000001.
```

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-3, Part 2 (Cont.) Task Builder Map for LIB.TSK

Memory allocation synopsis:

Section	Title	Ident	File
-----	-----	-----	-----
. BLK.: (RW,I,LCL,REL,CON) 000000 000000 00000.			
AADD : (RO,I,GBL,REL,CON) 000000 000024 00020.			
	000000 000024 00020.	LIB	01 LIB.OBJ;2
DIVV : (RO,I,GBL,REL,CON) 000024 000026 00022.			
	000024 000026 00022.	LIB	01 LIB.OBJ;2
MULL : (RO,I,GBL,REL,CON) 000052 000024 00020.			
	000052 000024 00020.	LIB	01 LIB.OBJ;2
SAVAL : (RO,I,GBL,REL,CON) 000076 000042 00034.			
	000076 000042 00034.	LIB	01 LIB.OBJ;2
SUBB : (RO,I,GBL,REL,CON) 000140 000024 00020.			
	000140 000024 00020.	LIB	01 LIB.OBJ;2

Global symbols:

```
AADD 000000-R MULL 000052-R SUBB 000140-R
DIVV 000024-R
```

*** Task builder statistics:

```
Total work file references: 368.
Work file reads: 0.
Work file writes: 0.
Size of core pool: 7086. WORDS (27. PAGES)
Size of work file: 768. WORDS (3. PAGES)
```

Elapsed time:00:00:03

Example 5-3, Part 3 Source Listing for MAIN.MAC

```
.TITLE MAIN
.IDENT /01/

;+
; **MAIN - CALLING ROUTINE TO EXERCISE THE ARITHMETIC ROUTINES
; FOUND IN THE RESIDENT LIBRARY, LIB.TSK.
;-

.MCALL QIOW$$,EXIT$$

OP1: .WORD 1 ; OPERAND 1
OP2: .WORD 1 ; OPERAND 2
ANS: .BLKW 1 ; RESULT
OUT: .BLKW 100. ; FORMAT MESSAGE
```

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-3, Part 3 (Cont.) Source Listing for MAIN.MAC

```

FORMAT: .ASCIZ  /THE ANSWER = %D./
        .EVEN
        .ENABL  LSB
START:
        MOV     #ANS,-(SP)           ; TO CONTAIN RESULT
        MOV     #OP2,-(SP)         ; OPERAND 2
        MOV     #OP1,-(SP)         ; OPERAND 1
        MOV     #3,-(SP)           ; PASSING 3 ARGUMENTS
        MOV     SP,R5              ; ADDRESS OF ARGUMENT BLOCK
        CALL    AADD               ; ADD TWO OPERANDS
        CALL    PRINT              ; PRINT RESULTS
        MOV     SP,R5              ; ADDRESS OF ARGUMENT BLOCK
        CALL    SUBB               ; SUBTRACT SUBROUTINE
        CALL    PRINT              ; PRINT RESULTS
        MOV     SP,R5              ; ADDRESS OF ARGUMENT BLOCK
        CALL    MULL               ; MULTIPLY SUBROUTINE
        CALL    PRINT              ; PRINT RESULTS
        MOV     SP,R5              ; ADDRESS OF ARGUMENT BLOCK
        CALL    DIVV               ; DIVIDE SUBROUTINE
        CALL    PRINT              ; PRINT RESULTS
        EXIT$$

;+
; ** PRINT - PRINT RESULT OF OPERATION.
;-

PRINT:  MOV     #OUT,R0            ; ADDRESS OF SCRATCH AREA
        MOV     #FORMAT,R1        ; FORMAT SPECIFICATION
        MOV     #ANS,R2           ; ARGUMENT TO CONVERT
        CALL    $EDMSG            ; FORMAT MESSAGE
        QIOW$$ #IO.WVB,#5,#1,,,,<#OUT,R1,#40>
        RETURN                    ; RETURN FROM SUBROUTINE
        .END     START

```

Once you have assembled MAIN, you can use the following TKB command sequence to build it:

```

TKB>MAIN,MAIN/MA/--WI/--SP=MAIN
TKB>/
Enter Options:
TKB>RESLIB=LIB/RO:3
TKB>//

```

This command sequence instructs TKB to build a task file named MAIN.TSK on device SY: under the UFD that corresponds to the terminal UIC. It also specifies that TKB is to create a map file MAIN.MAP. The /MA switch requests an extended map format. In this example, /MA was applied to the device specification so that TKB would include in the map for the task the symbols within the library LIB. The negated form of the wide listing switch (/WI) was appended to the map specification to obtain an 80-column map format. In this example, TKB will not output a map listing to the line printer

The RESLIB option specifies that the task MAIN is to access the library LIB and that it requires read-only access to LIB. TKB uses APR3 to map the library.

The TKB map that results from this command sequence is shown in Example 5-3, Part 4.

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-3, Part 4 Task Builder Map for MAIN.TSK

MAIN.TSK;1 Memory allocation map TKB M40.10 Page 1
 11-DEC-82 13:51

Partition name : GEN
 Identification : 01
 Task UIC : [7,62]
 Stack limits: 000274 001273 001000 00512.
 PRG xfr address: 001634
 Total address windows: 2.
 Task image size : 1152. WORDS
 Task address limits: 000000 004327
 R-W disk blk limits: 000002 000006 000005 00005.

*** Root segment: MAIN

R/W mem limits: 000000 004327 004330 02264.
 Disk blk limits: 000002 000006 000005 00005.

Memory allocation synopsis:

Section	Title	Ident	File
-----	-----	-----	-----
. BLK.: (RW, I, LCL, REL, CON)	001274 002620 01424.		
	001274 000530 00344. MAIN	01	MAIN.OBJ;1
	002024 001050 00552. EDTMG	15	SYSLIB.OLB;1034
	003074 000216 00142. CBTA	04.3	SYSLIB.OLB;1034
	003312 000074 00060. CATB	03	SYSLIB.OLB;1034
	003406 000250 00168. EDDAT	03	SYSLIB.OLB;1034
	003656 000126 00086. CDDMG	00	SYSLIB.OLB;1034
	004004 000110 00072. C5TA	02	SYSLIB.OLB;1034
AADD : (RO, I, GBL, REL, CON)	060000 000024 00020.		
	060000 000024 00020. LIB	01	LIB.STB;17
DIVV : (RO, I, GBL, REL, CON)	060024 000026 00022.		
	060024 000026 00022. LIB	01	LIB.STB;17
MULL : (RO, I, GBL, REL, CON)	060052 000024 00020.		
	060052 000024 00020. LIB	01	LIB.STB;17
SAVAL : (RO, I, GBL, REL, CON)	060076 000042 00034.		
	060076 000042 00034. LIB	01	LIB.STB;17
SUBB : (RO, I, GBL, REL, CON)	060140 000024 00020.		
	060140 000024 00020. LIB	01	LIB.STB;17
\$\$RESL: (RO, I, LCL, REL, CON)	004114 000212 00138.		
	004114 000024 00020. SAVRG	04	SYSLIB.OLB;1034
	004140 000066 00054. ARITH	03.04	SYSLIB.OLB;1034
	004226 000100 00064. DARITH	0007	SYSLIB.OLB;1034

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-3, Part 4 (Cont.) Task Builder Map for MAIN.TSK

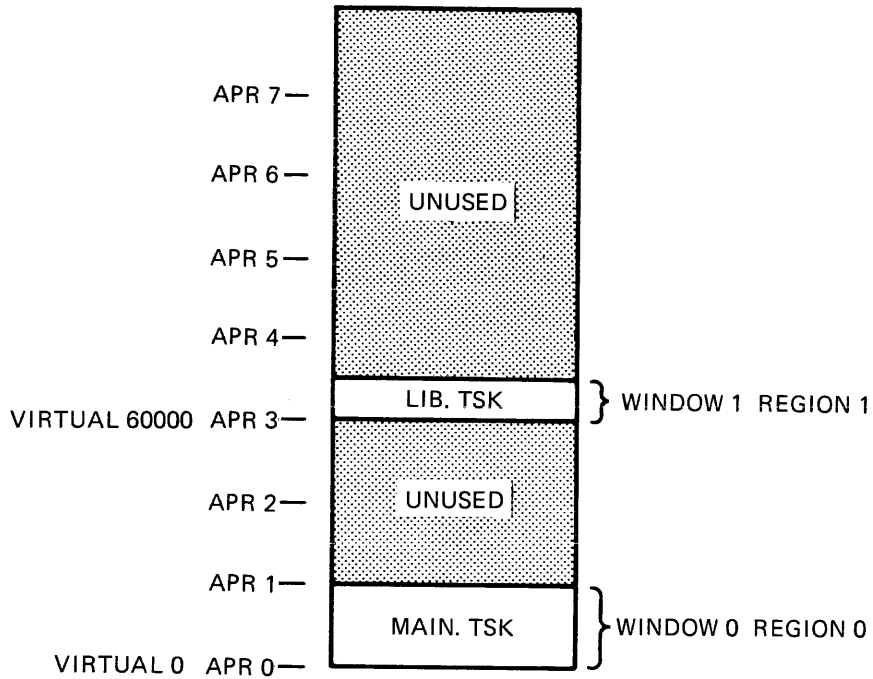
Global symbols:

AADD	060000-R	\$CBDSG	003110-R	\$CDTB	003312-R	\$EDMSG	002122-R
DIVV	060024-R	\$CBOMG	003116-R	\$COTB	003320-R	\$MUL	004140-R
IO.WVB	011000	\$CBOSG	003124-R	\$C5TA	004004-R	\$SAVRG	004114-R
MULL	060052-R	\$CBTA	003154-R	\$DAT	003452-R	\$TIM	003532-R
SUBB	060140-R	\$CBTMG	003132-R	\$DDIV	004264-R		
\$CBDAT	003074-R	\$CBVER	003116-R	\$DIV	004170-R		
\$CBDMG	003102-R	\$CDDMG	003656-R	\$DMUL	004226-R		

*** Task builder statistics:

Total work file references: 2218.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 2066. words (8. pages)
 Size of work file: 1024. words (4. pages)
 Elapsed time:00:00:19

This map contains a global symbols section. Note that the symbols within the library now have virtual addresses assigned to them and that these addresses begin at 60000(octal), the virtual base address of APR 3. The Task Builder's allocation of virtual address space for MAIN.TSK is represented diagrammatically in Figure 5-9.



ZK-426-81

Figure 5-9 Allocation of Virtual Address Space for MAIN.TSK

SHARED REGION CONCEPTS AND EXAMPLES

The library LIB is position independent and can therefore be mapped anywhere in the referencing task's virtual address space. APR 3 was used in this example to contrast this mapping arrangement with the mapping of MACCOM in the virtual address space of task MCOM1 in Example 5-1 (Section 5.1.7). If the optional APR parameter in the RESLIB option above had been left blank, TKB would have allocated the highest available APR to map the library.

5.1.11.1 Resolving Program Section Names in a Shared Region - As described in earlier sections of this chapter, program section names within position-independent shared regions must normally be unique with respect to program section names within tasks that reference them. When a shared region is a position-independent resident common and you explicitly declare the program section names within it, avoiding program section name conflicts is an easy matter. However, when a shared region is a position-independent resident library that contains calls to routines within an object module library (SYSLIB, for example), conflicts may develop that are not apparent to you. The problem arises when the position-independent resident library and one or more tasks that link to it contain calls to separate routines residing within the same program section of an object module library.

When TKB resolves a call from within a module that it is processing to a routine within an object module library, it places the routine from the library into the image it is building. It also enters into its internal table the name of the program section in the object module library within which the routine resides. If a position-independent resident library contains a call to a routine within a given program section of SYSLIB, for example, and then subsequently a task that links to the resident library contains a call to a different routine within the same program section of SYSLIB, both the resident library and the referencing task will contain the program section name. When you build the referencing task, the library's .STB file will contain the program section name and a program section conflict will develop. (Refer to Section 5.1.6 for additional information on the sequence in which TKB processes tasks and the potential program section name conflicts that can result.)

This situation and one possible solution to it can be illustrated with Example 5-3. When this example was first created, only the arithmetic routines were included in the source file of the resident library (LIB.MAC in Example 5-3, Part 1). The system library coroutine (\$SAVAL) was resolved from SYSLIB. Because the first instruction of each arithmetic routine called \$SAVAL, TKB included a copy of it in the resident library's image at task-build time. This turned out to be unsatisfactory because of a call to the SYSLIB routine \$EDMSG (edit message) within the program MAIN that links to the resident library. Both routines (\$SAVAL and \$EDMSG) reside within the unnamed or blank program section (. BLK.) within SYSLIB. Therefore, a program section name conflict developed when MAIN was built.

To circumvent this problem, the source code for \$SAVAL was included in the source file for the resident library under the explicitly declared program section name SAVAL.

Another solution would have been to build the resident library absolute. In this case, TKB would not have included program section names from the resident library into the .STB file for the resident library.

SHARED REGION CONCEPTS AND EXAMPLES

It is important to note that the above program section name conflict develops only when two different routines residing within the same program section of an object module library are involved. It presents no problem when a resident library and a task that links to it contain a call to the same routine in an object module library. In that case, TKB copies the routine and the program section name in which it resides into the resident library when the library is built. Then, when the task that calls the same routine is built, TKB resolves the reference to the routine in the resident library instead of in the object module library.

5.1.12 Example 5-4: Building a Task That Creates a Dynamic Region

In all the examples of tasks shown thus far in this chapter, TKB has automatically constructed and placed in the header of the task all of the window blocks necessary to map all of the regions of the task's image. The INSTALL processor has been responsible for initializing the window blocks when the task was installed. In all the examples, this has been possible because both TKB and the INSTALL processor have had all the information concerning the regions available to them.

When a task creates regions while it is running (dynamic regions), the information concerning the regions is not available to either the Task Builder or INSTALL. Therefore, when TKB builds such a task, it does not automatically create window blocks for the dynamic regions. It creates only the window blocks necessary to map the task region (the region containing the header and stack) and any shared regions that the task references.

Dynamic regions are created and mapped with Executive directives that are imbedded in the task's code. When you build a task that creates dynamic regions, you must explicitly specify to TKB how many window blocks (in excess of those created by TKB for the task region and any shared regions) it is to place in the task's header. The Executive will initialize these window blocks when it processes the region and mapping directives. In all (including window blocks for the task region and shared regions), you can include as many as 8 window blocks to a task in an RSX-11M system and as many as 16 in an RSX-11M-PLUS system.

The text in the remainder of this section and the figures associated with it illustrate the development of a task that creates dynamic regions. Example 5-4 shows a task (DYNAMIC.MAC) that creates a 128-word dynamic region. This task simply creates an unnamed region, maps to it, and fills it with an ascending sequence of numbers beginning at the region's base and moving upwards. When the region is full, DYNAMIC detaches from it and prints the following message on your terminal:

```
DYNAMIC IS NOW EXITING
```

The region is automatically deleted on detach.

All of the Executive directives used by DYNAMIC (RDBBK\$, WDBBK\$, DTRG\$\$, EXIT\$\$, CRG\$\$, CRAW\$\$, QIOW\$\$, and QIOW\$C) to create and manipulate the region are described in the RSX-11M/M-PLUS Executive Reference Manual. These directives are SYSGEN options on RSX-11M systems.

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-4, Part 1 Source Listing for DYNAMIC.MAC

```
.TITLE DYNAMIC
.IDENT /V01/
.MCALL RDBBK$,WDBBK$,DTRG$$,EXIT$$,CRRG$$,CRAW$$
.MCALL QIOW$C,QIOW$$

.NLIST BEX

; REGION DESCRIPTOR BLOCK
; WORD 0 SIZE OF REGION IN 32 DECIMAL WORD BLOCKS
; WORD 1 REGION NAME
; WORD 2 ""
; WORD 3 NAME OF SYSTEM CONTROLLED PARTITION IN
; WORD 4 WHICH REGION WILL BE CREATED
; WORD 5 STATUS WORD
; WORD 6 PROTECTION WORD

RDB: RDBBK$ 128.,,GEN,<RS.MDL!RS.ATT!RS.DEL!RS.RED!RS.WRT>,170017

; WINDOW DESCRIPTOR BLOCK
; WORD 0 APR TO BE USED TO MAP REGION
; WORD 1 SIZE OF WINDOW IN 32-WORD BLOCKS
; WORD 2 REGION ID
; WORD 3 OFFSET INTO REGION TO START MAPPING
; WORD 4 LENGTH IN 32-WORD BLOCKS TO MAP
; WORD 5 STATUS WORD

WDB: WDBBK$ 7,128.,0,0,,WS.MAP!WS.WRT>
MES1: .ASCIZ /DYNAMIC IS NOW EXITING/
S1 = . - MES1
ERR1: .ASCII /CREATE REGION FAILED/
SIZ1 = . - ERR1
ERR2: .ASCII /CREATE ADDRESS WINDOW FAILED/
SIZ2 = . - ERR2
ERR3: .ASCII /DETACH REGION FAILED/
SIZ3 = . - ERR3
.EVEN
.PAGE
.ENABL LSB

START:
CRRG$$ #RDB ; CREATE A 128 WORD UNNAMED REGION
BCS 1$ ; FAILED TO CREATE REGION
MOV RDB+R.GID,WDB+W.NRID ; COPY REGION ID INTO WINDOW BLOCK
CRAW$$ #WDB ; CREATE ADDR WINDOW AND MAP
BCS 2$ ; FAILED TO CREATE ADDR WINDOW
MOV WDB+W.NBAS,R0 ; BASE ADDR OF CREATED REGION
MOV WDB+W.NSIZ,R2 ; NUMBER OF 32. WORDS IN REGION
.REPT 5 ; MULTIPLY
ASL R2 ; BY
.ENDR ; 32.
MOV #1,R1 ; INITIAL VALUE TO PLACE IN REGION
20$: MOV R1,(R0)+ ; MOVE VALUE INTO REGION
INC R1 ; NEXT VALUE TO PLACE IN REGION
DEC R2 ; ONE LESS WORD LEFT
BGT 20$ ; TO FILL IN
DTRG$$ #RDB ; DETACH AND DELETE REGION
BCS 3$ ; DETACH FAILED
QIOW$C IO.WVB,5,1,,,,<MES1,S1,40>
EXIT$$ ;
```

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-4, Part 1 (Cont.) Source Listing for DYNAMIC.MAC

```

;
;      ERROR ROUTINES
;
1$:   MOV     #ERR1,R0      ; CREATE FAILED
      MOV     #SIZ1,R1      ; SIZE OF MESSAGE
      BR      6$           ; WRITE MESSAGE
2$:   MOV     #ERR2,R0      ; CREATE ADDRESS WINDOW FAILED
      MOV     #SIZ2,R1      ; SIZE OF MESSAGE
      BR      6$
3$:   MOV     #ERR3,R0      ; DETACH FAILED
      MOV     #SIZ1,R1      ; SIZE OF MESSAGE
6$:   QIOW$$  #IO.WVB,#5,#1,,,,<R0,R1,#40>
      EXIT$$
      .END      START

```

Once you have assembled DYNAMIC, you can build it with the following TKB command sequence:

```

TKB>DYNAMIC,DYNAMIC/-WI/--SP=DYNAMIC
TKB>/
Enter Options:
TKB>WWDWS=1
TKB>//

```

This command sequence directs TKB to create a task image named DYNAMIC.TSK and an 80-column (/WI) map file named DYNAMIC.MAP on device SY: under the terminal UIC. Because /-SP is attached to the map file, TKB does not output the file to the line printer.

Under options, the WWDWS option directs TKB to create one window block over and above that required to map the task region. Note that one window block must be created for each region the task expects to be mapped to simultaneously.

The map that results from this command sequence is shown in Example 5-4, Part 2.

Note that creating dynamic regions always involves the assumption that there will be enough room in the partition named in the task's region descriptor block to create the region when the task is run. In this example, if DYNAMIC were to be run in a system whose partition GEN was not large enough to accommodate the region it creates, the CREATE REGION directive would fail.

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-4, Part 2 Task Builder Map for DYNAMIC.TSK

DYNAMIC.TSK;1 Memory allocation map TKB M40.10 Page 1
11-DEC-82 16:05

Partition name : GEN
Identification : V01
Task UIC : [7,62]
Stack limits: 000274 001273 001000 00512.
PRG xfr address: 001470
Total address windows: 2.
Task image size : 512. WORDS
Task address limits: 000000 001753
R-W disk blk limits: 000002 000003 000002 00002.

*** Root segment: DYNAMI

R/W mem limits: 000000 001753 001754 01004.
Disk blk limits: 000002 000003 000002 00002.

Memory allocation synopsis:

Section	Title	Ident	File
-----	-----	-----	-----
. BLK.:(RW,I,LCL,REL,CON) 001274 000430 00280.			
001274 000430 00280.	DYNAMI	V01	DYNAMIC.OBJ;1
§DPB§§:(RW,I,LCL,REL,CON) 001724 000030 00024.			
001724 000030 00024.	DYNAMI	V01	DYNAMIC.OBJ;1

*** Task builder statistics:

Total work file references: 549.
Work file reads: 0.
Work file writes: 0.
Size of core pool: 7086. words (27. pages)
Size of work file: 768. words (3. pages)

Elapsed time:00:00:06

5.2 CLUSTER LIBRARIES


The term "cluster libraries" refers to both a function and a structure created by the Task Builder (TKB) that allow a task to dynamically map memory-resident shared regions at run time. Cluster libraries permit a task to use, for example, a F77CLS library, an FMS-11 library, and an FCS-11 library, all mapped through the same task address window. The run-time routines put into the task by the Task Builder remap the library regions so that, instead of occupying 48K bytes of virtual address space, they share 16K bytes of virtual address space.

One task address window (window 1) maps the libraries into the same span of virtual address space (48Kb to 64Kb). TKB maps your task from virtual 0 upward.

SHARED REGION CONCEPTS AND EXAMPLES

TKB implements the cluster library function in two parts. The first part, revectoring of interlibrary calls, is independent of the actual remap mechanism but is required for remapping to work. The second part executes the required MAP\$ directives to map the appropriate library.

The following examples use the library and task structure shown in Figure 5-10. Note that in the following examples, the FMS-11/RSX V1.0 and FORTRAN-77 software products are sold under separate license and are not included with the RSX-11M or RSX-11M-PLUS system. Cluster library support may be used with RMS-11 V2.0 or later versions, and operates in a fashion similar to the FCS-11 example. Also, the particular FCSRES used below is generated by SYSGEN. It consists of two PLAS overlays and a null root.



ZK-492-81

Figure 5-10 Example Library and Task Structure

5.2.1 Building the Libraries

You must follow several rules when designing and building shareable clustered libraries. The rules are summarized next and discussed in detail following the summary.

5.2.1.1 Summary of Rules for Building the Libraries -

- All libraries but the first require resident overlays.
- User task vectors indirectly resolve all interlibrary references.
- Revectoring entry point symbols must not appear in the "upstream" .STB file.
- A called library procedure must not require parameters on the stack.
- All the libraries must be PIC or built for the same address.
- Trap or asynchronous entry into a library is not permitted.


The rules are discussed in detail as follows.

SHARED REGION CONCEPTS AND EXAMPLES

5.2.1.2 Rule 1: All Libraries but the First Require Resident Overlays - The first library is the first named library in the CLSTR option line. To obtain the required run-time overlay data structures in your task, you must define all the libraries except possibly the first by using memory resident overlays. Although it can be an overlaid library, the first library need not be and can be a single-segment structure. All the libraries, except the first, must have a null root if overlaid. You can achieve this in cases where a library is not normally overlaid by creating an unbalanced overlay structure with a null module. For example, the following ODL specification for FMSCLS and a null module would suffice:

```
.NAME FMSCLS
.ROOT FMSCLS-*(NULL,FMSLIB)
NULL: .FCTR LB:[1,1]SYSLIB/LB:NULL ;NULL MODULE
FMSLIB: .FCTR SY:FMSLIB-LB:[1,1]FDVLIB/LB ;FMS-11 ROUTINES
.END
```

The above ODL specification creates an unbalanced tree in the form shown in Figure 5-11:

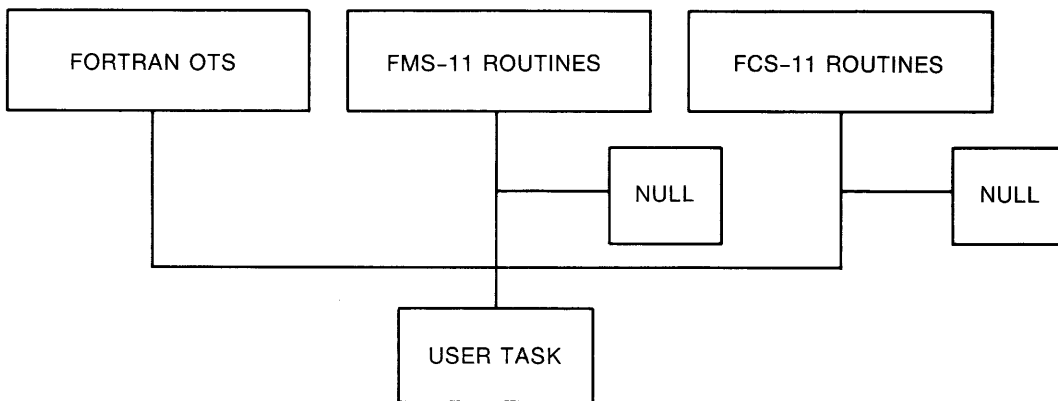


ZK-427-81

Figure 5-11 Example of an Unbalanced Tree with Null Segment

The effect, after you build your task, is an overlay structure that is represented in the Figure 5-12.

TKB provides the cross-library linkage that it creates from the overlay segment data contained in the individual .STB files of each library.




ZK-428-81

Figure 5-12 Example of an Overlay Cluster Library Structure

SHARED REGION CONCEPTS AND EXAMPLES

5.2.1.3 Rule 2: User Task Vectors Indirectly Resolve all Interlibrary References - Figure 5-13 below illustrates rule 2 and is a part of the example in Figure 5-12. In Figure 5-13, if the FORTRAN OTS library references an FCS-11 entry point .OPEN, the transfer of control from the FORTRAN OTS library to the FCS-11 library must be resolved by a jump vector in your task. Or, to state it in another way, the CALL instruction in the FORTRAN OTS library must not reference directly the target address (the address of .OPEN) in the FCS-11 library. The system library contains the modules that perform the indirect transfer for FCS-11 based libraries and user tasks. If you want to duplicate the indirect referencing mechanism for your own purposes, Figure 5-13 and the following text describe the control flow for FCS-11.



Sample code from FCSVEC module:

```

.OPEN::  MOV    #30,-(SP)           ; STACK OFFSET INTO USER TASK
        BR     DISPATCH           ; JUMP TABLE
        .
        .
DISPATCH: MOV    R0,-(SP)         ; SAVE REGISTER
        MOV    @#.FSRPT,R0        ; GET FCS-11 POINTER
        ADD    A.JUMP(R0),2(SP)    ; ADD VECTOR BASE TO OFFSET
        MOV    (SP)+,R0           ; RESTORE REGISTER
        MOV    @(SP)+,-(SP)       ; PICK UP ADDRESS OF TARGET
        RETURN                    ; AND TRANSFER TO TARGET
  
```

ZK-429-81

Figure 5-13 Example of a Vectored Call Between Libraries

In this example, the module FCSVEC defines the .OPEN entry point. The code at that location stacks an offset or "entry number" and joins common dispatch code. The dispatch code, using the low core FCS-11 impure pointer called .FSRPT, obtains the address of the FCS-11 impure data area. At offset A.JUMP in that area is the address of a vector of FCS-11 entry points. A return is executed, which transfers control to the routine whose address is now on top of the stack. If the target routine is an overlaid library, the run-time support (\$AUTO) loads the appropriate overlay and relays the transfer of control.

You may use this vectoring mechanism to isolate the linkages between two libraries whether or not you use them in the cluster library scheme. You can replace either the FORTRAN OTS or the FCS-11 library in your system without relinking the other library. However, you must relink your task when you replace either of these libraries.

SHARED REGION CONCEPTS AND EXAMPLES

5.2.1.4 Rule 3: Revectored Entry Point Symbols Must Not Appear in the "Upstream" .STB File - This rule means that the GBLXCL=symbol option must appear for each revectored symbol, as in FORTRAN OTS in this example. In the brief example above, the following line must appear in the build file for the FORTRAN OTS library:

```
GBLXCL=.OPEN
```

5.2.1.5 Rule 4: A Called Library Procedure Must Not Require Parameters on the Stack - This rule applies to routines contained in libraries other than the "default" library, as represented by the FMSCLS and FCSRES libraries of the above example. In addition, the called procedures must use the JSR PC and RTS PC call and return convention. The flow of control for a call into a cluster library member other than the default proceeds as follows.

Only your task can call and reference the FCSRES library routine .OPEN. All references from other libraries are revectored as described above. TKB resolves all such references to an appropriate task resident autoload vector. As in the example, when the FORTRAN OTS library calls .OPEN, the code revector the call through your task and hence to the autoload vector. At this point, the TKB run-time routine \$AUTO gets control and searches the overlay segment descriptor tree, noting which segments are resident and which must be loaded or mapped to access the target routine.

Next, \$AUTO notes that a member of a library cluster must be unmapped to comply with the map adjustments required to access the target routine. The reference to the unmapped library and the segment within the library is placed on the stack, the target library is mapped, and the target routine is accessed through a JSR PC instruction. That target routine must not attempt to access parameters by offsets from the stack pointer (SP) because of the presence of \$AUTO saved information. Upon return from the target by an RTS PC instruction, the target library is unmapped, and the previous library remapped using the saved segment and library data on the stack. Finally, \$AUTO executes an RTS PC instruction to return to the caller.

Note that if your task contains a mix of cluster libraries and noncluster libraries, the call format rule applies only to control transfers to cluster library routines. Other noncluster libraries that you create may use any appropriate call and parameter passing convention.

5.2.1.6 Rule 5: All the Libraries Must be PIC or Built for the Same Address - TKB must be able to place each library of the cluster at the same virtual address. To do this, the libraries must be built as position independent or be built to the exact address specified in the CLSTR command described below.

5.2.1.7 Rule 6: Trap or Asynchronous Entry Into a Library is not Permitted - A routine built as part of a library that is to be used in a cluster may not be specified as the service routine for a synchronous trap, or for asynchronous entry as a result of I/O completion or Executive service. This restriction is required because at the moment of the trap or fault, the appropriate library may not be the one that is mapped. For example, if the default library contains the service routine to display an error message upon odd address trap (the odd address fault occurs within one of the other libraries of the

SHARED REGION CONCEPTS AND EXAMPLES

cluster), the routine will not be available to service the trap. It will have been unmapped by the run-time routines to map the called library.

I/O completion and fault service vectors and routines must be placed in libraries or task segments that are resident at all times that the fault, trap, or I/O completion may occur.

5.2.2 Building Your Task

After building the individual libraries and placing the .TSK and .STB files for all the libraries into the LB:[1,1] directory, you may build your task. The TKB option line that you must use for your task has the following syntax:

```
CLSTR=library_1,library_2,...library_n:switch:apr
```

library_n

The first specification denotes the first or the default library, which is the library to which the task maps when the task starts up and remaps after any call to another library.

In an RSX-11M or RSX-11M-PLUS system, the total number of libraries to which a task may map is seven. The number of the component libraries in clusters is limited to a maximum of six. A cluster must contain a minimum of two libraries. It is possible to have two clusters of three libraries each or three clusters of two libraries each; any combination of clusters and libraries must equal at least two or a maximum of six. If six libraries are used in clusters, the task may map to only one other, separate library.

:switch:apr

The switch :RW or :RO indicates whether the cluster is read-only or read-write for this particular task. The APR specification is optional and indicates which APR is to be used as the starting APR when mapping to cluster libraries. If not specified, TKB assigns the highest available APRs and as many as required to map the library.

5.2.3 Examples

The sample build files for F77CLS, FDVRES, and FCSRES, and for the FMS-11 demonstration task FMSDEM are appended as an example of the cluster library-build process.

5.2.3.1 F77CLS -- Build the Default Library for the FORTRAN-77 OTS -

```
>TKB
TKB>F77CLS/-HD,F77CLS/CR/-SP/MA,F77CLS=F77RES
TKB>LB:[1,1]F77OTS/LB
TKB>LB:[1,1]SYSLIB/LB:FCSVEC           ; INCLUDE THE FCS JUMP VECTOR
TKB>/
Enter Options:
STACK=0
PAR=F77CLS:140000:40000
```

SHARED REGION CONCEPTS AND EXAMPLES

```
;
; FORCE THE JUMP TABLE TO BE LOADED FROM THE SYSTEM
; LIBRARY WHEN THE USER TASK IS BUILT
;
GBLINC=.FCSJT                                ; REFERENCE SYMBOL DEFINED IN
                                              ; THE MODULE SYSLIB/LB:FCSJMP
;
; PREVENT DEFINITIONS FOR FCS-11 ENTRY POINTS FROM APPEARING
; IN THE .STB FILE FOR THIS LIBRARY OR OTHER SYSTEM LIBRARY
;
GBLXCL=.CLOSE
GBLXCL=.CSI1
GBLXCL=.CSI2
GBLXCL=.DLFNB
GBLXCL=.FINIT
GBLXCL=.GET
GBLXCL=.GETSQ
GBLXCL=.GTDID
GBLXCL=.MRKDL
GBLXCL=.OPFNB
GBLXCL=.PARSE
GBLXCL=.POINT
GBLXCL=.POSRC
GBLXCL=.PRINT
GBLXCL=.PUT
GBLXCL=.PUTSQ
GBLXCL=.SAVR1
GBLXCL=.READ
GBLXCL=.WAIT
//
```

The GBLINC option as shown above forces TKB to add a global reference entry in the library .STB file. This ensures that TKB links certain modules required by the library, such as impure data areas or root-only routines, without further user action. These modules should be in the system library (LB:[1,1]SYSLIB.OLB) or in a library always referenced by your task, so that this forced loading mechanism is entirely invisible to you.

5.2.3.2 FDVRES -- Build an FMS-11/RSX V1.0 Shareable Library -

```
; TITLE OF THE EXAMPLE COMMAND FILE THAT BUILDS THE FORMS
; MANAGEMENT PLAS-RESIDENT LIBRARY FOR USE WITH THE
; TASK BUILDER CLSTR OPTION.
;
; FDVRES.CMD
;
; THE FOLLOWING CODE IS THE EXAMPLE COMMAND FILE:
;
LB:[1,1]FDVRES/-HD/MM/SG,MP:[1,34]FVRES/MA/-SP,LB:[1,1]FDVRES=
SY:[1,24]FDVRESBLD/MP
STACK=0
PAR=FDVRES:140000:40000
TASK=FDVRES
;
; THE FOLLOWING LINE FORCES THE FCS JUMP TABLE TO BE INCLUDED IN THE
; SYMBOL TABLE FILE FOR THE FORMS MANAGEMENT LIBRARY.
;
GBLINC=.FCSJT
```


SHARED REGION CONCEPTS AND EXAMPLES

```

;
; THE FOLLOWING LINES FORCE LIBRARY ENTRY POINTS AND DEFINITIONS INTO
; THE TASK ROOT:
;
GBLREF=CB$CUR, CB$REV, CB$TST, CB$132, DV$BLD, DV$BLK, DV$DHW, DV$DWD
GBLREF=DV$GRA, DV$REV, DV$UND, D$ATT1, D$ATT2, D$CLRC, D$FID, D$FXLN
GBLREF=D$LNCL, D$PICT, D$PLEN, D$RLEN, D$VATT, D$2ATT, D1$ALN, D1$ALP
GBLREF=D1$ARY, D1$COM, D1$MIX, D1$NUM, D1$SCR, D1$SNM, D2$DEC, D2$DIS
GBLREF=D2$FUL, D2$NEC, D2$REQ, D2$RTJ, D2$SPO, D2$TAB, D2$VRT, D2$ZFL
GBLREF=FC$ALL, FC$ANY, FC$CLS, FC$CSH, FC$DAT, FC$GET, FC$GSC, FC$LST
GBLREF=FC$OPN, FC$PAL, FC$PSC, FC$PUT, FC$RAL, RC$RTN, FC$SHO, FC$SLN
GBLREF=FC$SPF, FC$SPN, FC$TRM, FE$ARG, FE$DLN, FE$DNM, FE$DSP, FE$FCD
GBLREF=FE$FCH, FE$FLB, FE$FLD, FE$FNM, FE$FRM, FE$FSP, FE$ICH, FE$IFN
GBLREF=FE$IMP, FE$INI, FE$IOL, FE$IOR, FE$LIN, FE$NOF, FE$NSC, FE$STR
GBLREF=FE$UTR, FE$INC, FS$SUC, FT$ATB, FT$KPD, FT$NTR, FT$NXT, FT$PRV
GBLREF=FT$SBK, FT$SFW, FT$SNX, FT$SPR, FT$XBK, FT$XFW, F$ASIZ, F$CHN
GBLREF=F$FNC, F$IMP, F$LEN, F$NAM, F$NUM, F$REQ, F$RSIZ, F$STS
GBLREF=F$TRM, F$VAL, IS$ALT, IS$CLR, IS$DEC, IS$DSP, IS$ERR, IS$HFM
GBLREF=IS$HLP, IS$INS, IS$LST, IS$MED, IS$NMS, IS$SCR, IS$SGN, I$ADVO
GBLREF=I$ALLC, I$BADR, I$BEND, I$BPTR, I$BSIZ, I$CFRM, I$CURC, I$CURP
GBLREF=I$DISP, I$DLN1, I$DLN2, I$FADR, I$FBLK, I$FCHN, I$FDES, I$FDST
GBLREF=I$FDS1, I$FDS2, I$FIXD, I$FMST, I$FOFF, I$FORM, I$FSIZ, I$FXD1
GBLREF=I$FXD2, I$HLEN, I$HLPF, I$ILEN, I$IMPA, I$LCOL, I$LINE, I$LLIN
GBLREF=I$LNCL, I$LPTR, I$LVID, I$NBYT, I$NDAT, I$NFLD, I$PATN, I$PBLN
GBLREF=I$RESP, I$ROFF, I$STAT, I$STKP, I$SVST, I$VATT, L$CLSZ, L$FDES
GBLREF=L$LNCL, L$RESP, $F$DVT
GBLREF=$FDV
;
; THE FOLLOWING LINES PREVENT THE DEFINITIONS FOR FCS-11 ENTRY POINTS
; FROM APPEARING IN THE FORMS MANAGEMENT LIBRARY .STB FILE:
;
GBLXCL=.ASCPP
GBLXCL=.ASLUN
GBLXCL=.CLOSE
GBLXCL=.CTRL
GBLXCL=.DELET
GBLXCL=.DLFNB
GBLXCL=.ENTER
GBLXCL=.EXTND
GBLXCL=.FATAL
GBLXCL=.FCTYP
GBLXCL=.FIND
GBLXCL=.FINIT
GBLXCL=.FLUSH
GBLXCL=.GET
GBLXCL=.GETSQ
GBLXCL=.GTDID
GBLXCL=.GTDIR
GBLXCL=.MARK
GBLXCL=.MBFCT
GBLXCL=.MRKDL
GBLXCL=.OPEN
GBLXCL=.OPFID
GBLXCL=.OPFNB
GBLXCL=.PARSE
GBLXCL=.POINT
GBLXCL=.POSIT
GBLXCL=.POSRC
GBLXCL=.PPASC
GBLXCL=.PPR50
GBLXCL=.PRINT

```

SHARED REGION CONCEPTS AND EXAMPLES

```
GBLXCL=.PRSDI
GBLXCL=.PRSDV
GBLXCL=.PRSFN
GBLXCL=.PUT
GBLXCL=.PUTSQ
GBLXCL=.RDFDR
GBLXCL=.RDFFP
GBLXCL=.RDFUI
GBLXCL=.SAVR1
//
```

5.2.3.3 FDVRESBLD.ODL -- Overlay Description for FMS-11/RSX V1.0 Cluster Library -

```
;
; THE FOLLOWING LINE IS THE FILENAME OF THE .ODL FILE FOR THE
; PLAS-RESIDENT FORMS MANAGEMENT LIBRARY:
;
; FDVRESBLD.ODL
;
; THE FOLLOWING LINES OF CODE ARE CONTAINED IN THE .ODL FILE FOR THE
; PLAS-RESIDENT FORMS MANAGEMENT LIBRARY:
;
      .NAME      FDVROT
      .ROOT      FDVROT-*(MAIN,NULO)
NULO:  .FCTR     LB:[1,1]SYSLIB/LB:NULL
FCSV:  .FCTR     LB:[1,1]SYSLIB/LB:FCSVEC
MAIN:  .FCTR     LB:[1,1]FDVLIB/LB:FDV-LB:[1,1]FDVLIB/LB-FCSV
      .END
```

5.2.3.4 **FCSRES Library Build** - FCSRS1BLD.BLD is distributed with the RSX-11M and RSX-11M-PLUS distribution kits. Refer to the build command and overlay description contained in the files FCSRS1BLD.CMD and FCSRS1BLD.ODL, which can be generated by SYSGEN if you want.

5.2.3.5 F77TST.CMD -- File to Build the FMS-11/RSX V1.0 FORDEM Test Task -

```
FORDEM/FP, FORDEM/MA/-SP=FORDEM, HLLFOR
LB:[1,1]FDVLIB/LB
LB:[1,1]F77OTS/LB
/
EXTSCT=$$FSR1:2000
CLSTR=F77CLS,FDVRES,FCSRES:RO
STACK=200
//
```

5.2.4 Overlay Run-Time Support Requirements

The Task Builder uses the .STB files of the cluster libraries to obtain the information needed to create the overlay data base. For each PLAS overlaid cluster library TKB places autoload vectors, segment descriptors, window descriptors, and a region descriptor in

SHARED REGION CONCEPTS AND EXAMPLES

the root of the task. This information comprises the overlay run-time support for the cluster libraries. In Appendix B, Figure B-9 and the accompanying text describe this information. Table 5-1 describes the space needed for the overlay run-time system support that includes cluster libraries. For a complete description of overlay run-time routine sizes, see Section 4.5.

Using cluster libraries conserves virtual space and may require only one window.

Table 5-1
Comparison of Overlay Run-Time Module Sizes

Module	Program Section	Number of Bytes Oct/Dec	Specific Use
One of the following modules is included in any overlaid task that uses autoloader and in any task that links to a PLAS overlaid resident library.			
AUTO	\$\$AUTO	122/82.	All tasks that use autoloader
AUTOT	\$\$AUTO	132/90.	All tasks with AST's disabled during autoloader
	\$\$RTQ	32/26.	
	\$\$RTR	30/24.	

One of the following modules is included in any overlaid conventional task. OVCTR or OVCTC is included in any non-overlaid task (conventional or I- and D- space) that links to a PLAS overlaid resident library.

OVCTL	\$\$MRKS	76/62.	Disk overlays only
	\$\$RDSG	160/112.	
	\$\$PDLs	2/2.	
OVCTR	\$\$MRKS	234/156.	Disk and PLAS overlays with no cluster libraries
	\$\$RDSG	332/218.	
	\$\$PDLs	12/10.	
OVCTC	\$\$MRKS	254/172.	Disk and PLAS overlays with cluster libraries
	\$\$RDSG	352/234.	
	\$\$PDLs	120/80.	

One of the following three modules is included in any overlaid I- and D-space task.

OVIDL	\$\$MRKS	76/62.	Disk overlays only
	\$\$RDSG	224/148.	
	\$\$PDLs	2/2.	
OVIDR	\$\$MRKS	304/196.	Disk and PLAS overlays with no cluster libraries
	\$\$RDSG	502/322.	
	\$\$PDLs	12/10.	
OVIDC	\$\$MRKS	324/212.	Disk and PLAS overlays with cluster libraries
	\$\$RDSG	522/338.	
	\$\$PDLs	120/80.	

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

Table 5-1 (Cont.)
Comparison of Overlay Run-Time Module Sizes

Module	Program Section	Number of Bytes Oct/Dec	Specific Use
<p>The overlay data vector OVDAT is included in any overlaid task and in any task that links to a PLAS overlaid resident library.</p>			
OVDAT	\$\$OVDT	24/20.	Included in all tasks that perform overlay operations
	\$\$SGD0	0/0.	
	\$\$SGD2	2/2.	
	\$\$RTQ	0/0.	
	\$\$RTR	0/0.	
	\$\$RTS	2/2.	
<p>The overlay error service routine ALERR is included whenever OVDAT is included.</p>			
ALERR	\$\$ALER	24/20.	Overlay error
<p>Manual overlay control (LOAD) is used in place of any AUTO routine. (See Section 4.2, Manual Load.)</p>			
LOAD	\$\$LOAD	252/170.	Manual overlay control
	\$\$AUTO	14/12.	

5.3 VIRTUAL PROGRAM SECTIONS

A virtual program section is a special TKB storage allocation facility that permits you to create and refer to large data structures by means of the mapping directives. Virtual program sections are supported in TKB through the VSECT option and in FORTRAN through a set of FORTRAN-callable subroutines that issue the necessary mapping directives at run time. With the TKB VSECT option, you can specify the following parameters for a relocatable program section or FORTRAN common block that you have defined in your object module:

- Base virtual address
- Virtual length (window size)
- Physical length

By specifying the base address, you can align the program section on a 4K address boundary as required by the mapping directives. Thereafter, references within the program need only point to the base of the program section or to the first element in the common block to ensure proper boundary alignment.

By specifying the window size, you can fix the amount of virtual address space that TKB allocates to the program section. If the allocation made by a module causes the total size to exceed this limit, the allocation wraps around to the beginning of the window.

SHARED REGION CONCEPTS AND EXAMPLES

By specifying the physical size, you can allocate, before run time, the physical memory that the program section will be mapped into at run time. TKB allocates this physical memory within an area that precedes the task image. This area is called the mapped array area.

The physical length parameter is optional. If you intend to allocate physical memory at run time through the Create Region directive, you can specify a value of 0.

Note that when you specify a nonzero value for the physical memory parameter, the resulting allocation affects only the task's memory image, not its disk image.

Note also that TKB attaches the virtual attribute to a relocatable program section you have specified in the VSECT option only if the section is defined in the root segment of your task through either a FORTRAN COMMON or a MACRO-11 .PSECT statement. The relocatable program section with the virtual attribute in the root does not use address space in your task; using this procedure merely assigns an address, window size, and physical length to a region yet to be mapped at run time by your task. For example:

```
TKB>VSECT=MARRAY:160000:20000:2000
```

In this example, virtual program section MARRAY is allocated with a window size of 4K words (20000 (octal) bytes) and a base virtual address of 160000. In physical memory, 32K words are reserved for mapping the section at run time.

Assume the program is written in FORTRAN, and includes the following statement:

```
COMMON /MARRAY/ARRAY(4)...
```

This statement generates a program section to which TKB attaches the virtual attribute. However, this program section is not a FORTRAN virtual array. A reference to the first element of the section, ARRAY(1), is translated by TKB to the virtual address 160000.

Figure 5-14 shows the effect of this use of the VSECT option.

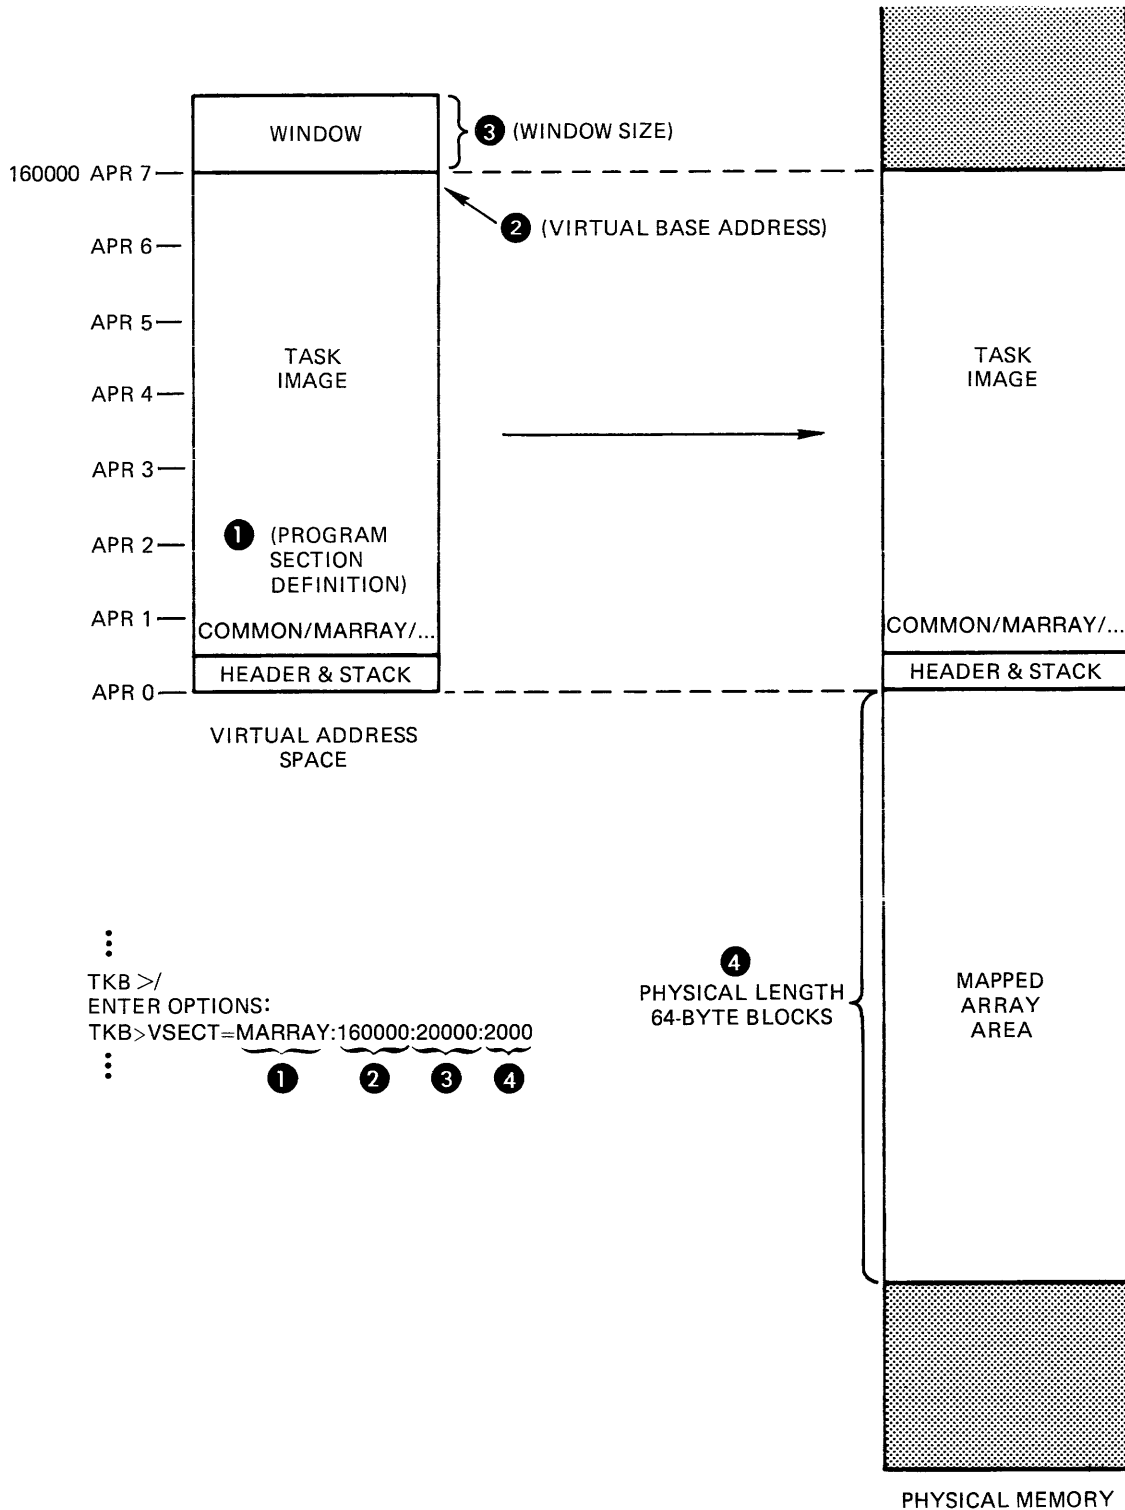
As mentioned previously, TKB restricts the amount of virtual address space allocated to the section to a value that is less than or equal to the window size, wrapping around to the base if the window size is exceeded.

This process is illustrated in the following example, in which three modules, A, B, and C, each contain a program section named VIRT that is 3000 words long. A window size of 4K words has been set through the VSECT option. If the program section has the concatenate attribute, the Task Builder allocates memory to each module as follows:

Module	Low Limit	Length	High Limit
A	160000	14000	174000
B	174000	14000	170000
C	170000	14000	164000

The address limits for modules B and C illustrate the effect of address wrap-around when a component of the total allocation exceeds the window boundary. Note that the addresses generated will be properly aligned with the contents of physical memory if the virtual section is remapped in increments of the window size.

SHARED REGION CONCEPTS AND EXAMPLES



ZK-430-81

Figure 5-14 VSECT Option Usage

SHARED REGION CONCEPTS AND EXAMPLES

5.3.1 FORTRAN Run-Time Support for Virtual Program Sections

FORTRAN supports subroutines to make use of the mapping directives. FORTRAN also supports calls to the following subroutines, which are related to virtual program sections:

Subroutine	Function
ALSCT	Allocates a portion of physical memory for use as a virtual section
RLSCT	Releases all physical memory allocated to a virtual section

As mentioned earlier, the effect of one or more VSECT= declarations at task-build time is to create a pool of physical memory below the task image (the mapped array area). Before a virtual section is referred to, the task must allocate a portion of this memory through a call to ALSCT. When space is no longer required, it is released through a call to RLSCT.

Note that these subroutines issue no mapping directives. They allocate and release space using region and window descriptor arrays that you supply. The resulting physical offsets are used in the task's subsequent calls that perform the actual mapping.

The subroutine ALSCT is called to allocate physical memory to a virtual program section as follows:

```
CALL ALSCT (ireg,iwnd[,ists])
```

ireg

A one-dimensional integer array that is nine words long. Elements 1 through 8 of the array contain a region descriptor for the physical memory to be mapped. The descriptor has the following format:

- ireg(1) Region ID.
- ireg(2) Size of region in units of 64-byte blocks.
- ireg(3) Name of region in Radix-50 format (first three characters).
- ireg(4) (Second three characters).
- ireg(5) Name of main partition containing region.
- ireg(6) The name is in Radix-50 format.
- ireg(7) Region status word.
- ireg(8) Region protection code.
- ireg(9) Thread word: This element links window descriptors that are used to map portions of the region. It is maintained by the subroutine.

The elements of the array that you set up consist of ireg(1) and ireg(3) through ireg(8). The thread word, ireg(9), must be 0 on the initial call; thereafter, the subroutine maintains it.

SHARED REGION CONCEPTS AND EXAMPLES

When your task makes an allocation, ireg(1) and ireg(2) must be 0 on the initial call. In this case, ALSCT obtains and stores the region size in ireg(2). When the allocation is being made from a separate region, the caller must supply both region ID and size. The subroutine does not refer to elements 3 through 8 but rather the caller must set them up as required by the applicable system directives. For a detailed description of these parameters, refer to the RSX-11M/M-PLUS Executive Reference Manual.

iwnd

A one-dimensional array that is 11 words long. The first eight words contain a window descriptor in the following format:

- iwnd(1) Base APR in bits 8 through 15; the Executive sets bits 0 through 7 when the appropriate mapping directives are issued.
- iwnd(2) Virtual base address.
- iwnd(3) Window size in units of 64-byte blocks.
- iwnd(4) Region ID.
- iwnd(5) Offset into the region, in units of 64-byte blocks.
- iwnd(6) Length to map, in units of 64-byte blocks.
- iwnd(7) Status word.
- iwnd(8) Address of send/receive buffer.
- iwnd(9) Base offset of physical block allocated to section in units of 64-byte blocks.
- iwnd(10) Length of block in units of 64-byte blocks (supplied by caller); set to maximum block offset by subroutine.
- iwnd(11) Thread word: This element links window descriptors that are used to map other portions of the region. It is maintained by the subroutine.

You must set up IWND(10) before calling ALSCT.

The following array elements are supplied as output from the subroutine:

iwnd(4), iwnd(5), iwnd(9), iwnd(10), and iwnd(11)

The remaining elements must be set up as required by the Executive directives. Consult the RSX-11M/M-PLUS Executive Reference Manual for a detailed description of these parameters.

ists

An area that receives the result of the call. One of the following values is returned:

- +1 Block successfully allocated. In this case, the region and window descriptor arrays are set up as described above.
- 200. Insufficient physical memory was available for allocating the block

The subroutine RLSCT is called to deallocate the physical memory assigned to a virtual section as follows:

```
CALL RLSCT (ireg,iwnd)
```

ireg

A one-dimensional integer array that is nine words long. The contents of the array are the same as those described for subroutine ALSCT.

iwnd

A one-dimensional integer array that is 11 words long. The contents of the array are the same as those described for subroutine ALSCT.

Upon return, element iwnd(10) is the length of the deallocated region in units of 64-byte blocks.

The procedure for using these subroutines can be summarized as follows:

1. You allocate storage in the program for one window descriptor per VSECT, and for a single region descriptor.
2. Your task calls the subroutine ALSCT to reserve the physical memory to which the program section will be mapped.
3. Your task issues the mapping directives to map the virtual address space into a portion of the physical memory. It is the task's responsibility to ensure that the physical memory to be mapped is always within the limits defined by iwnd(9) and iwnd(10).
4. When the space is no longer required, the task unmaps it and releases it with a call to RLSCT.

5.3.2 Example 5-5: Building a Program that Uses a Virtual Program Section

Example 5-5, Part 1 shows the FORTRAN source file for a task named VSECT.FTN. It illustrates the use of the ALSCT FORTRAN subroutine. When you build, install, and run VSECT, it will allocate the mapped array area below its header, create a 4K-word window, and map to the area through the window. ALSCT will then initialize the area and prompt for an array subscript at your terminal by printing:

```
SUBSCRIPT?
```

SHARED REGION CONCEPTS AND EXAMPLES

When you input a subscript, it responds with ELEMENT= and the contents of the array element for the subscript you typed. VSECT continues to prompt until you type CTRL/Z. Upon receiving a CTRL/Z, VSECT exits.

Once you have compiled VSECT, you can build it with the following Task Builder command sequence:

```
TKB>VSECT,VSECT/-SP=VSECT,LB:[1,1]FOROTS/LB
TKB>/
Enter Options:
TKB>WNDWS=1
TKB>VSECT=MARRAY:160000:20000:200
TKB>//
```

This command sequence directs TKB to create a task image file named VSECT.TSK and a short (by default) map file VSECT.MAP. Because /-SP is appended to the map file, TKB does not output the map to the line printer.

The library switch (/LB) specifies that TKB is to search the FORTRAN run-time library FOROTS.OLB to resolve any undefined references in the input module VSECT.OBJ. Because the library switch was applied to the FORTRAN library file without arguments, TKB extracts from the library and includes in the task image any modules in which references are defined.

The WNDWS option directs TKB to add a window block to the header in the task image. The Executive initializes this window block when it processes the mapping directives within the task.

The VSECT option directs TKB to establish for the program section named MARRAY a base address of 160000 (APR 7) and a length of 20000 (octal) bytes (4K words). The program section VIRT is defined within the task through the FORTRAN COMMON statement. The VSECT option also specifies that TKB is to allocate 200 64-byte blocks of physical memory in the task's mapped array area below the task's header. (For more information on the switches and options used in this example, refer to Chapters 10 and 11.)

The map that results from this command sequence is shown in Example 5-5, Part 2.

Example 5-5, Part 1 Source Listing for VSECT.FTN

```
C
C
C
VSECT.FTN
INTEGER *2 SUB,IRDB(9),IWDB(11),DSW
INTEGER *2 IARRAY(4096)
COMMON /MARRAY/IARRAY
IWDB (1) = "3400           !USE APR 7 FOR WINDOW
IWDB (3) = 128           !WINDOW SIZE = 128*32 WORDS = 4K
IWDB (5) = 0             !OFFSET
IWDB (7) = "402         !STATUS = WS.64B!WS.WRT
IWDB (10) = 128         !SIZE TO ALLOCATE
```

(continued on next page)

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-5, Part 1 (Cont.) Source Listing for VSECT.FTN

```

C
C   ALLOCATE 4K MAPPED ARRAY TO IWDB,IRDB
C
C   CALL ALSCT (IRDB,IWDB,DSW)
C   IF (DSW .NE. 1) GOTO 100
C
C   CREATE A 4K ADDRESS WINDOW
C
C   CALL CRAW (IWDB,DSW)
C   IF (DSW .NE. 1) GOTO 200
C
C   MAP 4K MAPPED ARRAY
C
C   CALL MAP (IWDB,DSW)
C   IF (DSW .NE. 1) GOTO 300
C   DO 1 I=1,4096
C     IARRAY (I) = I
C
C   MAPPED ARRAY IS INITIALIZED, PROMPT FOR A SUBSCRIPT
C
C   WRITE (5,5)
C   5   FORMAT ('$SUBSCRIPT?')
C   READ (5,4,END=1000)SUB
C   4   FORMAT (I7)
C   WRITE (5,6)IARRAY (SUB)
C   6   FORMAT ( ' ELEMENT = ',I7)
C   GOTO 3
C
C   ERROR ROUTINES
C
C   100  WRITE (5,101)DSW
C   101  FORMAT ( ' ERROR FROM ALSCT. ERROR = ',I7)
C       GOTO 1000
C   200  WRITE (5,201)DSW
C   201  FORMAT ( ' ERROR FROM CREATING ADDRESS WINDOW. ERROR = ',I7)
C       GOTO 1000
C   300  WRITE (5,301)DSW
C   301  FORMAT ( ' ERROR FROM MAPPING. ERROR = ',I7)
C   1000 CALL EXIT
C       END

```

SHARED REGION CONCEPTS AND EXAMPLES

Example 5-5, Part 2 Task Builder Map for VSECT.TSK

VSECT.TSK;1 Memory allocation map TKB M40.10 Page 1
 11-DEC-82 16:12

Partition Name : GEN
 Identification : FORV02
 Task UIC : [303,1]
 Stack limits: 000300 001277 001000 00512.
 PRG xfr address: 016270
 Total address windows: 2.
 Mapped array area: 4096. words
 Task image size : 8736. words
 Task address limits: 000000 042043
 R-W disk blk limits: 000002 000044 000043 00035.

*** Root segment: VSECT

R/W mem limits: 000000 042043 042044 17444.
 Disk blk limits: 000002 000044 000043 00035.

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.: (RW,I,LCL,REL,CON) 001300 001160 00624.			
MARRAY: (RW,D,GBL,REL,OVR) 160000 020000 08192.			
	160000 020000 08192.	.MAIN.	FORV02 VSECT.OBJ;3
OTSS\$F : (RW,I,GBL,REL,CON) 002460 002332 01242.			
	002460 000406 00262.	\$CONVI	F40003 FOROTS.OLB;2
	003066 001724 00980.	\$FIO	F40006 FOROTS.OLB;2
OTSS\$I : (RW,I,LCL,REL,CON) 005012 011220 04752.			
	.		
	.		
	.		

Global symbols:

ADI\$IA 005032-R CAL\$ 005140-R ICI\$ 022466-R MOI\$PS 006050-R
 .
 .

*** Task builder statistics:

Total work file references: 27855.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 7086. words (27. PAGES)
 Size of work file: 4325. words (17. PAGES)

Elapsed time:00:00:29

CHAPTER 6

PRIVILEGED TASKS

6.1 INTRODUCTION

This chapter discusses privileged tasks: what they are, their possible hazards, how they are mapped, and an example of their usage.

6.2 PRIVILEGED AND NONPRIVILEGED TASK DISTINCTION

RSX-11M/M-PLUS systems have two classes of tasks: privileged and nonprivileged. The distinction between privileged and nonprivileged tasks is primarily based upon system-access capabilities. Because all tasks in an unmapped system have access to all of memory, this distinction is not hardware enforceable. Therefore, if your system is unmapped, your task must be responsible for observing the access rules of your system.

In a mapped system, privileged tasks have special device and memory access rights that nonprivileged tasks do not have. A privileged task can, with certain exceptions, access the Executive routines and data structures; a nonprivileged task cannot. Some privileged tasks have automatic I/O page mapping available to them; nonprivileged tasks do not. Finally, a privileged task can bypass system security features, whereas a nonprivileged task cannot.

6.3 PRIVILEGED TASK HAZARDS

Because of their special access rights, privileged tasks are potentially hazardous to a running system. A privileged task with coding errors can corrupt the Executive or system data structures. Moreover, problems caused by such a privileged task can be obscure and difficult to isolate. For these reasons, you must exercise caution when developing and running a privileged task.

Make certain that your privileged task has completed its operation when you log off the system (type BYE). BYE does not abort privileged tasks as it does nonprivileged tasks because the privileged task may be in the process of changing the system data base. Therefore, it must be allowed to complete its processing. Also, if the privileged task is in system state, neither BYE nor any other task can execute until the privileged task completes its processing while in system state. However, when the privileged task leaves system state, BYE runs and logs you off the system, leaving the privileged task still in operation.

PRIVILEGED TASKS

If a processor trap occurs in a privileged task while the task is in user state, the Executive aborts the task. However, if the processor trap occurs in the privileged task while the task is in system state, the system crashes. However, even while in user state the privileged task that is mapped to the Executive can cause a system crash by incorrectly changing system data. Please note that a privileged task in user state should not be modifying system data.

All tasks in an unmapped system can access all of memory. The privileged or nonprivileged designation has no particular meaning in an unmapped system. Therefore, be just as careful about modifying Executive, device, or user data in an unmapped system.

6.4 SPECIFYING A TASK AS PRIVILEGED

You designate a task as privileged with the /PR (privileged) TKB switch (this switch is described in Chapter 10). TKB allocates address space for a privileged task based on the memory management APR that you specify as an argument to this switch. The argument is optional; the default is 5 but you can change it by modifying the TKBBLD.CMD file and rebuilding TKB. TKB accepts three arguments: 0, 4, and 5. Choosing which of these arguments to specify is based on the considerations described below.

6.5 PRIVILEGED TASK MAPPING

When you specify an argument of 0, your task is marked as privileged but not mapped to the Executive or I/O page. Virtual address space begins at virtual address 0 and extends upward as far as 32K words. Your task cannot access the Executive routines or data structures, and TKB does not reserve an APR to map the I/O page.

When you specify /PR:4 or /PR:5, TKB reserves APR 7 for mapping the I/O page. Moreover, TKB makes the Executive available to your task by reserving the APRs necessary to map the Executive into your task's virtual address space. Therefore, if your task requires access to the Executive, you must specify an argument of either 4 or 5. However, 5 is the default.

The choice between APR 4 and 5 is dictated by the size of the Executive area. If the Executive is 16K words or less, you may specify an argument of 4 or 5. The value specified depends on the task size. A /PR:4 task can be 12K in size and map the I/O page. TKB applies a bias of 100000 (16K) to all addresses within your task.

If the Executive is 20K words, you must specify an argument of 5. TKB applies a bias of 120000 (20K) to all addresses within your task.


The mapping for privileged tasks is shown in Figure 6-1.

The mapping for APR 4 and 5 is shown in Figure 6-2.


When you specify an argument of 4, there will be 12K words of address space between the beginning of the task and the start of the mapping for the I/O page. If your task expects to access the I/O page, it must not exceed this 12K-word limit. If it does, TKB uses APR 7 to map the task instead of the I/O page.

When you specify an argument of 5, there will be 8K words of address space between the beginning of the task and the start of the mapping for the I/O page. In this case, the task must not be greater than 8K words if it expects to access the I/O page.

PRIVILEGED TASKS



MAPPING FOR 8K PRIVILEGED TASK IN USER STATE AND 20K EXECUTIVE




MAPPING FOR 8K PRIVILEGED TASK IN SYSTEM STATE AND 20K EXECUTIVE

ZK-431-81

Figure 6-1 Privileged Task Mapping

PRIVILEGED TASKS



ZK-432-81

Figure 6-2 Mapping for /PR:4 and /PR:5

When a task overlaps the I/O page, TKB does not generate an error message. Before TKB generates an error message, a task designated to be mapped with APR 4 must be greater than 16K words; a task designated to be mapped with APR 5 must be greater than 12K words. Only when you install a task that overlaps the I/O page does INSTALL generate the following message:

INS--WARNING--PRIVILEGED TASK NOT MAPPED TO THE I/O PAGE

While this is not a fatal error message, you should consider the condition to be fatal if you expect your task to access the I/O page.

You can use the /-IP switch to inform TKB that the task is purposely over 12K and does not need to be mapped to the I/O page.

A /PR:4 or /PR:5 task can access all of the Executive, system control blocks, and I/O page. It can use Executive routines and do logical block I/O to a volume that is physically mounted on a device. Also, the task can issue a \$SWSTK macro to change from user to system state. This allows the task to access the Executive or system data structures without interruptions or fear of the data being changed while it is being accessed.

6.6 /PR:0 PRIVILEGED TASK

Using the /PR:0 switch causes TKB to build the task in the same way as any other task. Virtual address space begins at virtual address 0 and extends upwards as far as 32K minus 32 words. This task cannot access the Executive routines and system data structures or directly access the I/O page because the Task Builder has not reserved APRs for these purposes.

PRIVILEGED TASKS

There are advantages to using a /PR:0 task and having it mapped into APR 0. A /PR:0 task can:

- Bypass file protection.
- Use the alter priority (ALTP\$) directive.
- Issue any directive that has a target task.
- Specify a device name in spawn directives.
- Write logical block I/O to a physically mounted volume, regardless of who issued the Mount or Allocate command. For example, the VMR task is a /PR:0 task and writes to mounted volumes during the SYSGEN process. However, this advantage can be hazardous for obvious reasons.

A /PR:0 task runs in user state and cannot switch to system state. Also, a /PR:0 task is not mapped to the Executive. If you want to write a privileged task that does I/O processing, it is advantageous to use the /PR:0 switch for your task because there is less chance of corrupting the Executive or system code and data.

6.7 /PR:4 PRIVILEGED TASK

If you want your privileged task to map to the Executive and I/O page, and your Executive is 16K or less, use the /PR:4 switch in TKB command line. If you specify /PR:4 for your task, TKB reserves APR 7 to map the I/O page and reserves APRs 0 through 3 to map the Executive as part of your task's virtual address space. The /PR:4 switch can be used only if your Executive size is 16K or less, because the 16K Executive uses APRs 0 through 3 and your task is assigned mapping that starts with APR 4. Therefore, TKB applies a bias of 100000 (16K decimal) to all virtual addresses within the task. This specific mapping of APRs 0 through 4 and 7 occurs whether the task is in user or system state.

Up to 12K words of virtual address space are possible in a /PR:4 task. The beginning of the task marks the end of the Executive code. If the task is 12K words in size, the end of the task marks the start of the I/O page. If the task is going to access the I/O page through APR 7, the task cannot exceed the 12K limit. If the task does exceed the limit, TKB is forced to assign APR 7 to the task code. When building the task, TKB does not give you an error message if your task exceeds the 12K limit. However, when you install the task, INSTALL sends you the following message:

```
"INS -- WARNING -- PRIVILEGED TASK NOT MAPPED TO THE I/O PAGE"
```

6.8 /PR:5 PRIVILEGED TASK

If you want your privileged task to map to the Executive and I/O page, and your Executive is between 16K and 20K, use the /PR:5 switch in TKB command line. If you specify /PR:5 for your task, TKB reserves APR 7 to map the I/O page and reserves APRs 0 through 4 to map the Executive as part of your task's virtual address space. The /PR:5 switch can be used only if your Executive size is between 16K and 20K, because the 20K Executive uses APRs 0 through 4 and your task is assigned APR 5. (APR 5 may be used if the Executive is less than 16K, but this wastes virtual address space.) Therefore, TKB applies a bias of 120000 (20K)

PRIVILEGED TASKS

to all virtual addresses within the task. This specific mapping of APRs 0 through 5 and 7 occurs whether the task is in user or system state.

Up to 8K words of virtual address space (12K if the I/O page is overmapped) are possible in a /PR:5 task. The beginning of the task marks the end of the Executive code. If the task is 8K words in size, the end of the task marks the start of the I/O page. If the task is going to access the I/O page through APR 7, the task cannot exceed the 8K limit. If the task does exceed the limit, TKB is forced to assign APR 7 to the task code. When building the task, TKB does not give you an error message if your task exceeds the 8K limit. However, when you install the task, INSTALL sends you the following message:

```
"INS -- WARNING -- PRIVILEGED TASK NOT MAPPED TO THE I/O PAGE"
```

NOTE

When you use a privileged task, the Executive has dedicated almost all the APRs to the necessary mapping for the Executive, the I/O page, and your task. Your task can issue PLAS directives to remap any number of these APRs to regions. However, such remapping can cause obscure and difficult-to-find system bugs. Also, note that when a directive unmaps a window that formerly mapped the Executive or the I/O page, the Executive restores the former mapping.

6.9 EXAMPLE 6-1: BUILDING A PRIVILEGED TASK TO EXAMINE UNIT CONTROL BLOCKS

The MACRO-11 source program PRIVEX.MAC in Example 6-1 illustrates one possible use of a privileged task.

NOTE

The nature of a privileged task is such that you must have a working knowledge of system concepts to understand its operation or to write one. If this example deals with Executive functions that are unfamiliar to you, you may prefer to skip this section and return to it at a later time.

If you assemble, build, and install PRIVEX into your system, it will scan the system device tables and examine the UCBs of all nonpseudo devices on your system. It will determine whether each device is attached by a task and print on your terminal the names of all attached devices on your system with the name of each attached program.

PRIVILEGED TASKS

PRIVEX accesses two Executive routines: \$SWSTK (switch stack) and \$SCDVT (scan device tables). The routine \$SWSTK switches the processor to system state (kernel mode). This switch to system state is necessary because it inhibits all other processes from modifying the Executive data structures until PRIVEX is finished with them. The double semicolons (;;) indicate the portion of the task that is running in system state.

The routine \$SCDVT performs the actual scanning of the device tables. It returns to PRIVEX each time it accesses a new UCB.

PRIVEX also calls the system library routine \$EDMSG (edit message) to format the data it has retrieved from the device tables. This routine is documented in the IAS/R SX-11 System Library Routines Reference Manual.

Example 6-1, Part 1 Source Code for PRIVEX

```

; MACRO LIBRARY CALLS
  .TITLE PRIVEX
  .MCALL ALUN$C,EXIT$$,QIOW$$
;
; LOCAL DATA
;
  .NLIST BEX

ATTMES: .ASCIZ /%2A%P: IS ATTACHED BY %2R/      ;
BUFMES: .ASCIZ /BUFFER OVERFLOW/                ;
        .LIST BEX
QIOBUF: .BLKB 132.                               ;MESSAGE OUTPUT BUFFER
        .EVEN

;
; BUFFER INTO WHICH INFORMATION IS STORED AT SYSTEM STATE FOR
; PRINTING AT USER STATE. AN ENTRY IS FOUR WORDS LONG:
;
;
; ADDRESS IN DCB OF THE TWO ASCII CHARACTER DEVICE NAME
;
; BINARY UNIT NUMBER
;
; FIRST RAD50 WORD OF NAME OF ATTACHED TASK
;
; SECOND RAD50 WORD OF NAME OF ATTACHED TASK
;
; THE BUFFER IS TERMINATED BY A
;
; 0 = ALL UNITS IN THE SYSTEM HAVE BEEN EXAMINED
; -1 = THE BUFFER WAS FILLED BEFORE ALL UNITS COULD BE EXAMINED
;

BUFFER: .BLKW 4*200.+1                          ;
BUFEND=-2                                       ;ADDRESS OF LAST WORD OF BUFFER

```

(continued on next page)

PRIVILEGED TASKS

Example 6-1, Part 1(Cont.) Source Code for PRIVEX

```

START:  MOV      #BUFFER,R2      ;GET ADDRESS OF INFORMATION BUFFER
        CLR      (R2)           ;ASSUME NO UNITS ARE ATTACHED
        CLR      R1             ;INITIALIZE CURRENT DCB ADDRESS

;
; "CALL $SWSTK,FORMAT" SWITCHES TO SYSTEM STATE. ALL REGISTERS
; ARE PRESERVED ACROSS THE TRANSITION FROM USER MODE TO KERNEL
; MODE. BEING IN SYSTEM STATE LOCKS OTHER PROCESSFS OUT OF THE
; EXECUTIVE (GUARANTEES THAT THE DATA BEING EXAMINED WILL NOT
; CHANGE WHILE IT IS BEING EXAMINED). A "RETURN" WILL GIVE
; CONTROL TO "FORMAT" AND WILL RESTORE THE CONTENTS OF THE
; REGISTERS TO THEIR VALUES BEFORE THE "CALL $SWSTK".
;
        CALL     $SWSTK,FORMAT   ;SWITCH TO SYSTEM STATE
        MOV      #$SCDVT,-(SP)  ;;GET ADDRESS OF SCAN DEVICE TABLES
                                ;;COROUTINE
20$:    CALL     @(SP)+          ;;GET NEXT NONPSEUDO DEVICE UCB
                                ;; ADDRESS
        BCS     100$            ;;IF CS NO MORE UCBS

;
; AT THIS POINT:
; R3 - ADDRESS OF THE DEVICE CONTROL BLOCK
; R4 - ADDRESS OF THE STATUS CONTROL BLOCK
; R5 - ADDRESS OF THE UNIT CONTROL BLOCK
;
        CMP     R1,R3           ;;IS THIS A NEW DCB?
        BEQ     40$             ;;IF EQ NO
        MOV     R3,R1          ;;REMEMBER THIS DCB
        CLR     R0             ;;FORM LOWEST UNIT NUMBER ON
        BISB   D.UNIT(R3),R0   ;; THIS DCB
40$:    MOV     U.ATT(R5),R4    ;;IS A TASK ATTACHED?
        BEQ     60$             ;;IF EQ NO
                                ;;IF NE R4 IS TCB ADDRESS
        CMP     #BUFEND,R2     ;;ANY MORE ROOM IN BUFFER?
        BLOS   80$             ;;IF LOS NO
        ADD    #D.NAM,R3      ;;FORM ADDRESS OF DEVICE NAME
        MOV    R3,(R2)+        ;;SAVE IT IN BUFFER
        MOV    R0,(R2)+        ;;SAVE UNIT NUMBER
        MOV    T.NAM(R4),(R2)+ ;;SAVE NAME OF ATTACHED TASK
        MOV    T.NAM+2(R4),(R2)+ ;;
        CLR    (R2)           ;;ASSUME NO MORE ATTACHED UNITS
60$:    INC    R0              ;;INCREMENT UNIT NUMBER
        BR     20$            ;;
80$:    CALL   @(SP)+         ;;GET $SCDVT TO CLEAN OFF STACK
        BCC   80$            ;;
        COM   (R2)           ;;SHOW BUFFER OVERFLOW
100$:   RETURN                ;;RETURN TO USER STATE AT FORMAT

        .ENABL  LSB
FORMAT: TST    (R2)           ;ANY MORE INFORMATION IN BUFFER?
        BEQ   EXIT          ;IF EQ NO
        CMP   #-1,(R2)      ;OVERFLOWED BUFFER?
        BNE   40$          ;IF NE NO
        MOV   #BUFMES,R1    ;GET ADDRESS OF OVERFLOW MESSAGE
        CALL  PRINT         ;PRINT IT
EXIT:    EXIT$$            ;

```

(continued on next page)

PRIVILEGED TASKS

Example 6-1, Part 1(Cont.) Source Code for PRIVEX

```

40$:  MOV      #ATTMES,R1      ;GET ADDRESS OF TEMPLATE
      CALL    PRINT          ;FORMAT AND PRINT THE INFORMATION
      BR      FORMAT        ;

      .DSABL  LSB

;
; PRINT - FORMAT AND PRINT A MESSAGE
;
; INPUTS:
;   R1 - ADDRESS OF AN $EDMSG INPUT STRING
;   R2 - ADDRESS OF AN $EDMSG PARAMETER BLOCK
;
; OUTPUTS:
;   R2 - ADDRESS OF NEXT PARAMETER IN THE $EDMSG PARAMETER BLOCK
;   R0, R1, R3, R4 ARE DESTROYED
;   R5 IS PRESERVED
;

PRINT: MOV     #QIOBUF,R0     ;GET ADDRESS OF OUTPUT BUFFER
      MOV     R0,R3          ;SAVE FOR QIOW$$
      CALL    $EDMSG        ;FORMAT MESSAGE INTO OUTPUT BUFFER

;
; REMOVE LEADING ZEROS FROM UNIT NUMBER
;

      MOV     R3,R0          ;POINT AT OUTPUT BUFFER
      TST    (R0)+          ;INCREMENT BY TWO (POINT PAST
                          ; DEVICE NAME)
20$:  MOV     R0,R4          ;REMEMBER THIS SPOT
      DEC     R1             ;ASSUME NEXT BYTE IS A LEADING ZERO
                          ; (REDUCE LENGTH OF MESSAGE)
      CMPB   #'0,(R0)+      ;IS IT?
      BEQ    20$            ;IF EQ YES -- IGNORE IT
      INC    R1             ;COUNTERACT TOO MUCH DECREMENTING
      CMPB   #' :,-(R0)     ;WAS THE BYTE A COLON (WAS THE UNIT
                          ; NUMBER ZERO)?
      BNE    40$            ;IF NE NO
      MOVB  #'0,(R4)+       ;ADD A ZERO UNIT NUMBER
      INC    R1             ;INCREASE LENGTH OF MESSAGE
40$:  MOVB   (R0)+,(R4)+     ;TACK ON REST OF MESSAGE
      BNE    40$            ;IF NE NOT DONE

;
; PRINT THE MESSAGE ON LUN "OUTLUN" (DEFINED BY THE TASK BUILD FILE)
; AND WAIT USING EVENT FLAG 1
;

      QIOW$$ #IO.WVB,#OUTLUN,#1,,,,R3,R1,#' >> ;
      RETURN
      .END    START

```

PRIVEX.MAC should be assembled with the following assembler command string:

```
MAC>PRIVEX,PRIVEX/-SP=DR0:[1,1]EXEMC/ML,[11,10]RSXMC/PA:1,DR2:[303,1]PRIVEX
```

The file EXEMC is the Executive macro library and the file RSXMC is the Executive prefix file. The switches used in the command string are described in the IAS/RSX-11 MACRO-11 Programmer's Reference Manual.

PRIVILEGED TASKS

The Task Builder command sequence for PRIVEX is as follows:

```
>TKB
TKB> PRIVEX/PR:5,PRIVEX/--SP=PRIVEX
TKB> DR0:[3,54]RSX11M.STB,DR0:[1,1]EXELIB/LB
TKB> /
Enter Options:
TKB> UNITS=1 ;DEFINE NUMBER OF LUNS
TKB> GBLDEF=OUTLUN:1 ;DEFINE LUN ON WHICH TO PRINT MESSAGES
TKB> ASG=TI0:1 ;ASSIGN LUN TO DEVICE
TKB> //
>
```

This command sequence directs TKB to build PRIVEX as a privileged task and to add a bias of 120000 to all locations within it. APR 5 was chosen in this example because the Executive in the system on which this example was originally built is 20K words long. If the Executive in your system is 16K words or less, you can use /PR:4 when you build the task.

In the options section of the TKB command sequence, the UNITS=1 option specifies that PRIVEX will use only one logical unit. The GBLDEF=OUTLUN:1 option defines the symbol OUTLUN as being equal to 1, and the ASG=TI0:1 option associates device TI0: with logical unit 1.

The TKB map for PRIVEX is shown in Example 6-1, Part 2. The GLOBAL SYMBOL SECTION has been shortened to save space. Note that the task's address limits begin at virtual address 120000. Figure 6-3 illustrates how TKB allocates virtual address space for the program.

Example 6-1, Part 2 Task Builder Map for PRIVEX

```
PRIVEX.TSK;1 Memory allocation map TKB M40.10 Page 1
7-OCT-82 13:26
```

```
Partition name : GEN
Identification : 01
Task UIC : [303,1]
Stack limits: 120230 121227 001000 00512.
PRG xfr address: 124610
Task attributes: PR
Total address windows: 1.
Task image size : 1920. words
Task address limits: 120000 127323
R-W disk blk limits: 000002 000011 000010 00008.
```

```
*** Root segment:PRIVEX
```

```
R/W mem limits: 120000 127323 007324 03796.
Disk blk limits: 000002 000011 000010 00008.
```

(continued on next page)

PRIVILEGED TASKS

Example 6-1, Part 2 (Cont.) Task Builder Map for PRIVEX

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.:(RW,I,LCL,REL,CON) 121230 005746 03046.			
121230 003656 01966.	PRIVEX 01		PRIVEX.OBJ;2
\$\$RESL:(RO,I,LCL,REL,CON) 127176 000124 00084.			


Global symbols:

AS.DEL 000001	BT.UAB 000002	DV.SDI 000020	D.RS81 177657
D.VOUT 000004	F.NWAC 000034	IE.DAA 177770	
AS.EXT 000004	B.DIR 000026	DV.SQD 000040	D.RS83 177655
D.VPWF 000006	F.SCHA 000015	IE.DNA 177771	
	.		
	.		
\$PDVTA 020000	\$REMOV 054044	\$SGFFR 020652	\$TICLR 041032
\$YHCTB 022674	.TT14 023770		

*** Task builder statistics:

Total work file references: 250535.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 13486. words (52. PAGES)
 Size of work file: 12032. words (47. PAGES)

Elapsed time:00:00:51



ZK-433-81

Figure 6-3 Allocation of Virtual Address Space for PRIVEX

CHAPTER 7

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

This chapter discusses the Task Builder's ability to divide a user task into instruction and data space (I- and D-space). A series of figures and text explain task mapping and the use of task windows in an RSX-11M-PLUS system with an I- and D-space task. In the text, comparisons are made between conventional tasks and I- and D-space tasks. A conventional task is one that does not separately map instruction space and data space.

The I- and D-space feature is an RSX-11M-PLUS system generation option. The feature is available only on specific processor hardware. Conventional tasks can be run in an I- and D-space system, but an I- and D-space task cannot run in a system that does not have the option specified.

7.1 USER TASK DATA SPACE DEFINED

User task data space is that space that contains data and which the user task accesses through D-space APRs. The function of I- and D-space allows a total of 16 APRs to map your task: 8 APRs for data space and 8 APRs for instruction space. If your task uses both I- and D-space to its maximum capacity, it can contain 64K words of virtual address space. In addition to both I- and D-space, if your task links to a 32K word supervisor-mode library, it can contain 96K words of virtual address space.

To separate the data and instructions, your task can use PSECTS to contain the data or instructions. Also, your task can use the CRAW\$ and CRRG\$ directives to dynamically create and map to data-space regions. See the RSX-11M/M-PLUS Executive Reference Manual for the use of these directives.

Conventional tasks and tasks that separate instruction space and data space differ in only a few areas of interest. The next sections discuss these areas.

7.2 I- AND D-SPACE TASK IDENTIFICATION

Two fields denote an I- and D-space task. In the task header, the byte that has the offset H.DMAP identifies the task D-space mapping mask. In the Task Control Block (TCB), the T4.DSP bit in the fourth task status word identifies the I- and D-space task to the system.

The system task loader or the VMR FIX command initializes these two fields at the time the task is loaded. Therefore, tasks built on a system other than an I- and D-space system may be run without rebuilding on an RSX-11M-PLUS system that supports I- and D-space.

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

The I- and D-space task is one in which TKB separates the data areas and instructions. In this task, data areas should be defined by the MACRO-11 .PSECT directive that has the data attribute. Similarly, the .PSECT directive with the "I" attribute defines instruction areas.

7.3 COMPARISON OF CONVENTIONAL TASKS AND I- AND D-SPACE TASKS

A conventional task operating in user mode can contain 32Kwords of virtual address space and access approximately 32Kwords of physical memory. However, a task using both I- and D-space APRs can contain 64Kwords of virtual address space and access approximately 64Kwords of memory.

The conventional task in an I- and D-space system uses both sets of APRs. However, the relocation addresses in both I-space and D-space APRs are identical. Also, the task windows refer to I-space APRs in a task that does not use D-space.

An I- and D-space task can use separately both I- and D-space APRs; that is, APRs used in this way are not overmapped. Because of this, the task can use eight D-space APRs to access and use data, and eight I-space APRs to access and execute instructions. Using 16 APRs allows the I- and D-space task to access a total of 64Kwords of physical memory at one time.

Table 7-1 contains a brief mapping summary of the combinations of I- and D-space tasks, I- and D-space systems, and the APR mapping that occurs.

Table 7-1
Mapping Comparison Summary

I/D Task	I/D System	Mapping Summary
No	Yes	I-space APRs and D-space APRs contain the same relocation addresses.
No	No	I-space APRs contain relocation.
Yes	Yes	I-space APRs map instruction space. D-space APRs map data space.
Yes	No	Not possible.

7.4 CONVENTIONAL TASK MAPPING

Conventional tasks map their virtual addresses to their logical addresses through both I-space and D-space APRs. That is, TKB does not separate instruction space or data space nor does the system differentiate the spaces except by the logic inherent in the task. Therefore, the task must map to its logical address space by both sets of APRs, which are overmapped.

Figure 7-1 shows an 8K conventional task linked to an 8K region that maps to its logical address space through both D-space and I-space APRs in an I- and D-space system.

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

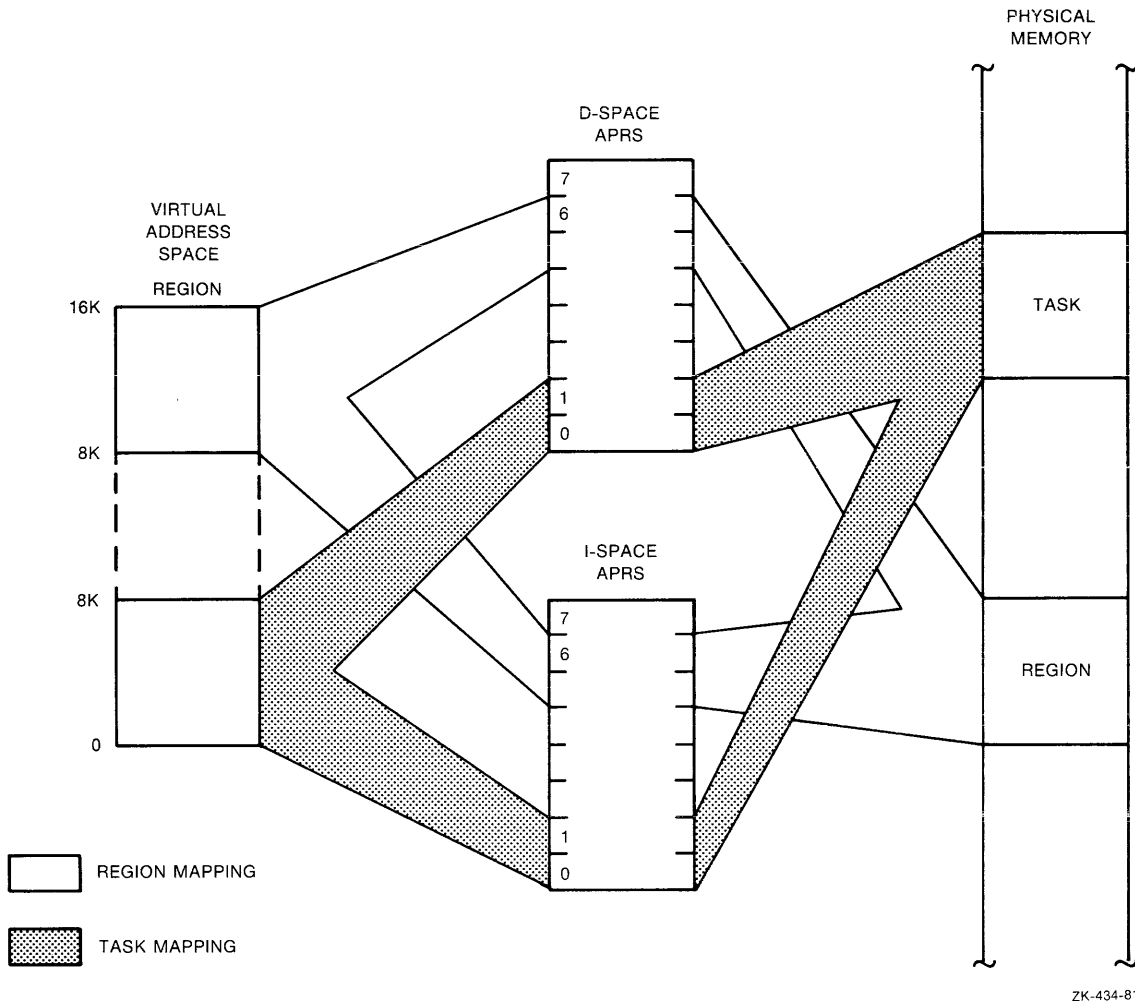


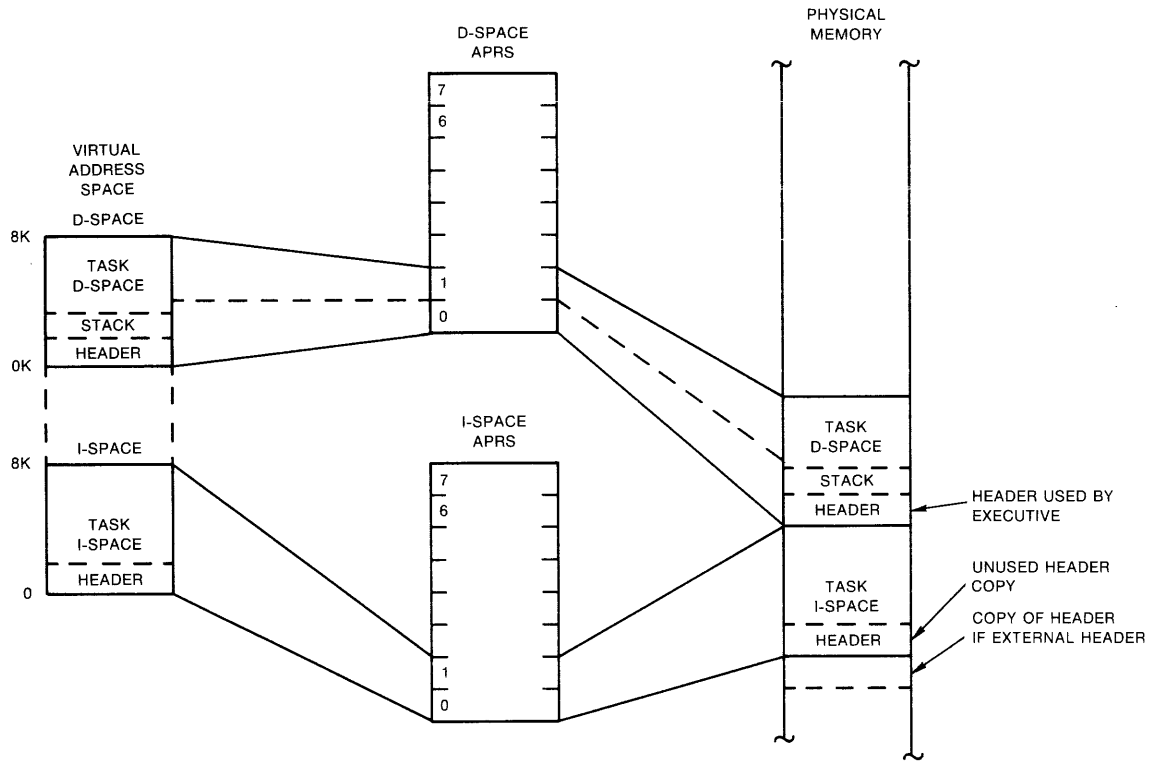
Figure 7-1 Conventional Task Linked to a Region in an I- and D-Space System

7.5 I- AND D-SPACE TASK MAPPING

Figure 7-2 shows an 8K I- and D-space task. TKB separated the data and instructions in this task. Because of the way TKB processes task space, the task header must physically reside at the beginning of the task in I-space. TKB puts the header that the Executive uses for task control in D-space. Also, the task's stack is in D-space. If the task is to have an external header (under control of the /XH switch), the Executive copies the header in D-space and puts it into the contiguous space immediately before the task's I-space in memory. For more details, see Figure B-4, Image on Disk of Overlaid I- and D-Space Task, in Appendix B.

The task shown uses two APRs because of its size (8K). D-space APR 0 maps the task's header and stack and part of D-space.

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)



ZK-435-81

Figure 7-2 I- and D-space Task Mapping in an I- and D-space System

7.6 TASK WINDOWS IN I- AND D-SPACE TASKS

TKB uses different windows to map various portions of an I- and D-space task. Window 0 in an I- and D-space task cannot be used because it maps the root in I-space. Similarly, you cannot use window 1 because it maps the D-space part of the root. The root of the task, which TKB divides into I- and D-space, therefore requires two windows. TKB reserves the use of these two windows. You can specify up to 14 windows for a task that uses I- and D-space.

7.7 SPECIFYING DATA SPACE IN YOUR TASK

You design an I- and D-space task by specifying data space separately from instruction space. Good programming practice suggests that all data areas and buffers should be located in adjacent locations. Similarly, all instructions should be located in adjacent locations. However, TKB will separate and agglomerate instruction and data space when it builds the task. For TKB to do this, you must use a method of informing it about which statements are data and which are instructions.

For the MACRO-11 programmer, the way to separate data and instructions is to use the MACRO-11 .PSECT directive. You can use this directive with the instruction (I) attribute for all the instruction locations in your task's code. Also, you can use .PSECT and the data (D) attribute for all the data locations. You must define a data .PSECT in an I- and D-space task even though no actual data is contained in the task. In this case, the .PSECT can be of 0 length.

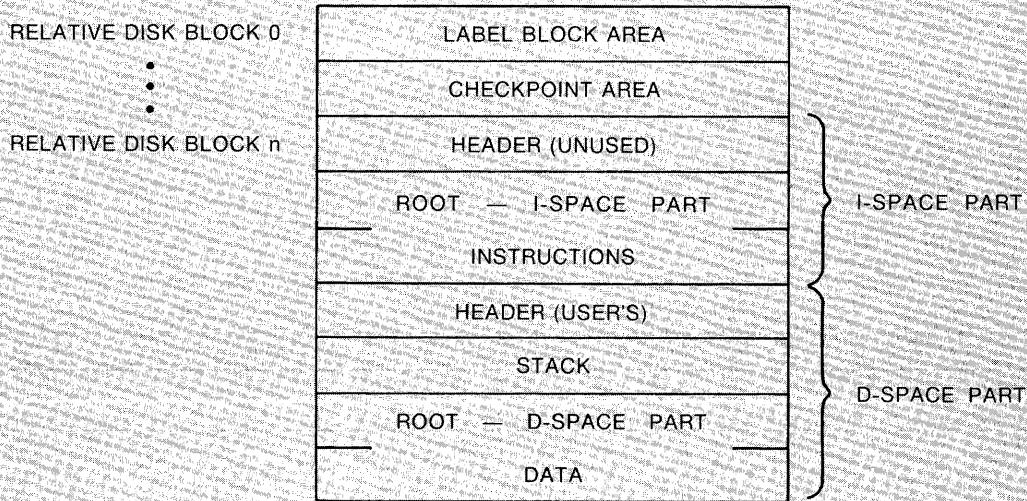
Note that I- and D-space libraries have not been defined and are not a possible configuration.

7.8 OVERLAID I- AND D-SPACE TASKS

Except for the mapping of an I- and D-space task and the location of instructions and data, the I- and D-space task differs little from a conventional task. However, there are structural differences between a non-overlaid and an overlaid I- and D-space task. By comparing the two kinds of tasks, the figures and text in the following sections describe the non-overlaid and overlaid I- and D-space tasks. Also, you may want to refer to the description of overlaid conventional tasks in Chapter 3.

Figure 7-3 shows a simplified disk image of the non-overlaid I- and D-space task. This task contains four I-space PSECTs and four D-space PSECTs. TKB collects all the I-space PSECTs together in one part of the root and all the D-space PSECTs in another part of the root.

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)



ZK-1098-82

Figure 7-3 Simplified Disk Image of a Non-Overlaid I- and D-Space Task

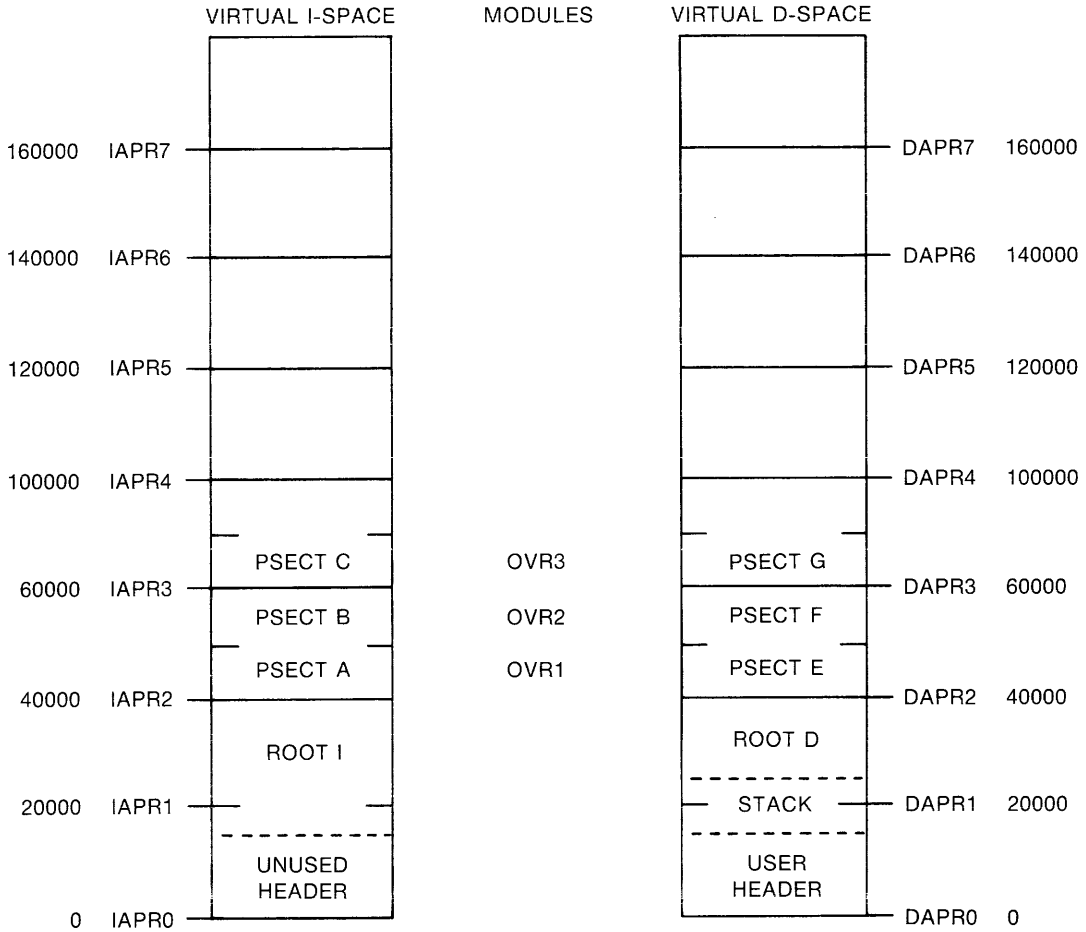
Figure 7-4 shows the virtual address space and physical memory occupied by an overlaid I- and D-space task called IAND. The task has a total physical size of 160000(octal) bytes. (You may want to compare Figure 7-4, which is shown next, with Figure 3-1 in Chapter 3, which shows a conventional overlaid task.) The instructions and data occupy the same virtual address space and are of equal physical size; but because they are mapped through different APRs, they occupy different locations in physical memory. The instructions in IAND occupy four PSECTs that have the instruction (I) attribute, and the data in IAND occupy four PSECTs that have the data (D) attribute.

In Figure 7-4, the virtual instruction space contains PSECTs A, B, and C, which are those that contain instructions, and ROOT I, which is the PSECT in the root that contains instructions. TKB places the unused header in I-space part of the root.

Also, in Figure 7-4, the virtual data space contains PSECTs D, E, and F, which are those that contain data, and ROOT D, which is the PSECT in the root that contains data. TKB places the task's user header in the D-space part of the root.


As an overlaid task, a possible overlay tree may look like the one shown in Figure 7-5.

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)



ZK-1099-82

Figure 7-4 Overlaid I- and D-Space Task Virtual Address Space



ZK-1100-82

Figure 7-5 Example Overlay Tree for Overlaid I- and D-Space Task IAND

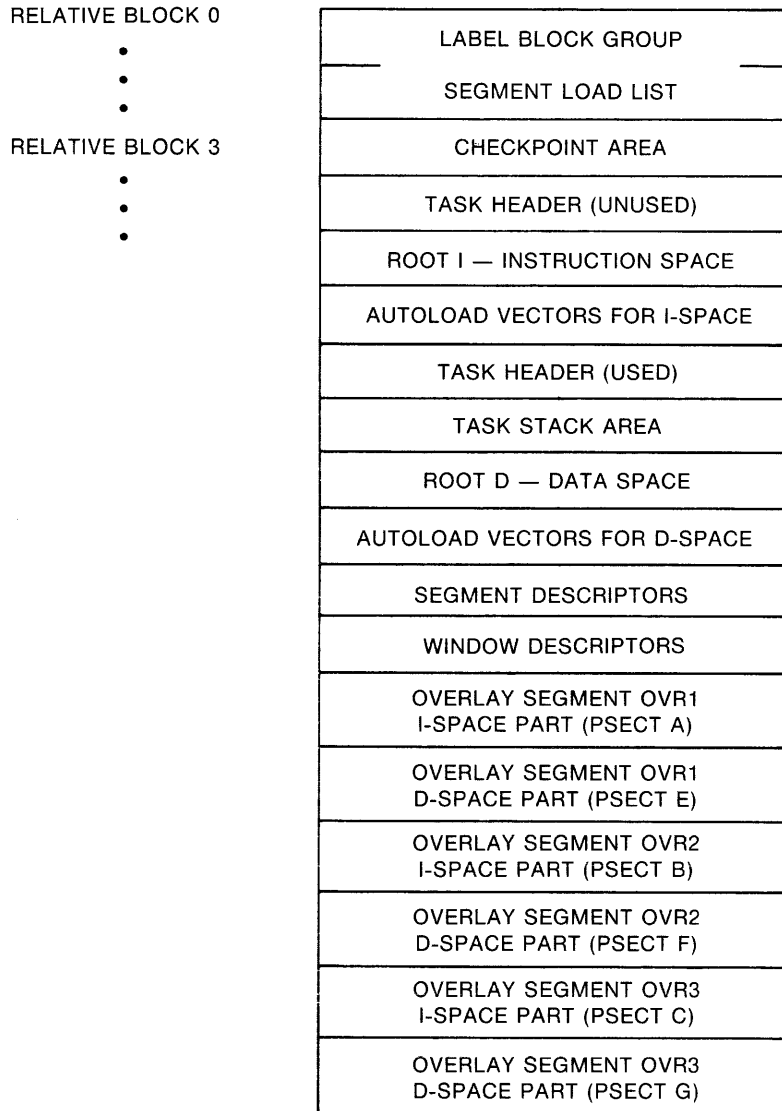
The accompanying ODL statement for this task is:

```
.ROOT ROOT-(OVR1,OVR2,OVR3)
```

Notice that this ODL statement is not different from any overlaid task with this tree structure. In this statement, the module OVR1 contains the instruction PSECT A and the data PSECT E, the module OVR2 contains the instruction PSECT B and the data PSECT F, and the module OVR3 contains the instruction PSECT C and the data PSECT G. Also, the ROOT module contains the instruction PSECT I and the data PSECT D.

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

The disk image of this overlaid task, shown in Figure 7-6, contains the instruction and data PSECTs in separate areas. Figure 7-6 also illustrates the difference between disk images of overlaid and non-overlaid I- and D-space task disk images when you compare it with the disk image shown in Figure 7-3. Notice that TKB separates the segments of the overlaid IAND task into instruction parts and data parts. Any autoloading vectors generated because of calls from these segments are also included in the segment area. The autoloading vectors for I- and D-space tasks contain two parts: an I-space part and a D-space part. TKB places each part with its corresponding segment part as shown in Figure 7-6. Autoloading vectors for I- and D-space tasks are discussed in detail in Chapter 4.



ZK-1101-82

Figure 7-6 Simplified Disk Image of Overlaid I- and D-Space Task IAND

7.8.1 Autoload Vectors and .STB Files

If your I- and D-space task links to an overlaid shared region, that region must have been built with a version of TKB that supports overlaid I- and D-space tasks. The reason for this is that the .STB files for overlaid shared regions built by older versions of TKB do not contain the ISD records that are needed to create the type of autoload vectors that I- and D-space tasks use.

For newer versions of TKB that support overlaid I- and D-space tasks, TKB allocates autoloadable vectors in the root of the task only for those entry points in the library referenced by the task. To create the autoload vectors, TKB uses ISD records in the .STB file when linking the task to the library if the ISD records are present. Therefore, tasks built with newer versions of TKB tend to be smaller because fewer autoload vectors are present.

For the Fast Task Builder (FTB) and older versions of TKB that do not support I- and D-space tasks, each autoload vector in the shared region's .STB file is allocated in the root of the task being linked to the region, whether or not the entry point is referenced by the task.

NOTE

Libraries created with older versions of TKB do not have the ISD records in the .STB file that newer versions of TKB use to include autoload vectors in the task from the .STB file. Therefore, TKB must create autoload vectors for every entry point in the library.

If you are using one of these older libraries and you are linking an I- and D-space task to it, TKB will give you the fatal error message:

```
"Module      module-name      contains
incompatible autoload vectors"
```

This message occurs because the .STB file contains conventional autoload vectors that are not usable by an I- and D-space task.

For more information about linking shared regions to I- and D-space tasks, see the section in Chapter 5 entitled, Autoload Vectors and .STB Files for Overlaid Shared Regions.

7.9 I- AND D-SPACE TASK MEMORY ALLOCATION AND EXAMPLE MAPS


The following section discusses and shows the differences between two versions of a task that is built both as a conventional task and as an I- and D-space task. The conventional task is called MAIN.TSK and the I- and D-space version of MAIN.TSK is called MAINID.TSK. Both of

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

these tasks are similar to but not the same as the task called MAIN.TSK shown in Chapter 5. After MAIN.TSK was coded, built, and the map printed, MAIN.TSK was rebuilt as an I- and D-space task to create MAINID.TSK. To do this, the /ID switch was used in the TKB command line. Both versions of this task are overlaid and link to a library.

7.9.1 Virtual Memory Allocation for MAIN.TSK

MAIN.TSK has a root called MAIN and three overlay segments called INPUT, CALC, and OUTPUT. By comparing the map of this task in Example 7-1 and the memory allocation diagram in Figure 7-7, you will be able to determine the virtual memory space allocation and structure of this task. Note that the overlay segments occupy the same virtual address space, and the root and segments are mapped in both I-space and D-space.



ZK-1107-82

Figure 7-7 Memory Allocation Diagram for MAIN.TSK

7.9.2 Virtual Memory Allocation for MAINID.TSK

MAINID.TSK has a root called MAIN and three overlay segments. In this way MAINID.TSK resembles MAIN.TSK. However, the I-PSECTS and D-PSECTS in MAINID.TSK are separated and they are mapped through their respective I-space or D-space APRs. Therefore, MAINID.TSK has two virtual address spaces: an I-space and a D-space. Figures 7-8 and 7-9 show the memory allocation for the I-space and D-space in MAINID.TSK.

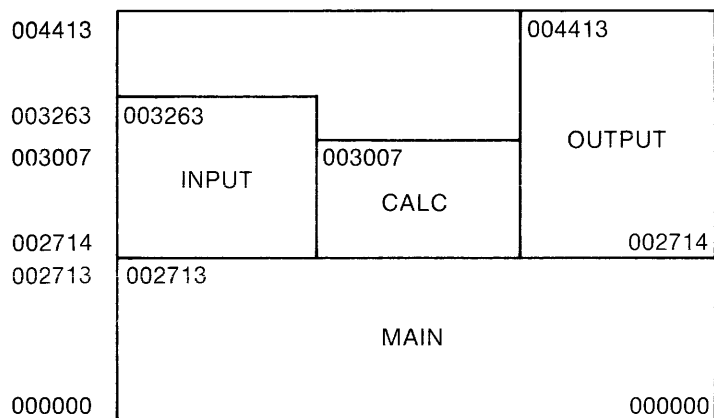
The three segments INPUT, CALC, and OUTPUT occupy the same virtual I-space, because these three segments contain instructions and, therefore, I-PSECTS. However, the overlay segment OUTPUT is the only segment that occupies virtual D-space because the segments INPUT and CALC do not contain data or D-space PSECTS. Note that the two overlay segments INPUT and CALC have no D-space. You can see this in both Figure 7-9 and the map in Example 7-2. The map in Example 7-2 shows the virtual address space allocation for both I-space and D-space.

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

An I- and D-space task uses more virtual memory space than a conventional task. The map in Example 7-2 shows that MAINID.TSK uses 1888. words of space as opposed to the 1664. words used by MAIN.TSK. The reasons for the increase in size of MAINID.TSK over MAIN.TSK are as follows:


- An I- and D-space task contains an unused task header.
- Autoload vectors in an I- and D-space task contain two more words than conventional autoload vectors. PSECTS \$\$ALVD and \$\$ALVI in Example 7-2 contain the autoload vectors. You can see from this map that they use more space than the \$\$ALVC PSECT in MAIN.TSK, which contains conventional autoload vectors.
- The segment descriptors in an overlaid I- and D-space task contain an extension.

In addition to these reasons for the increase in size of an overlaid I- and D-space task, a memory-resident overlaid I- and D-space task would be even larger because of the need for two window descriptors for each memory-resident segment.



ZK-1108-82

Figure 7-8 Memory Allocation Diagram for MAINID.TSK I-Space



ZK-1109-82

Figure 7-9 Memory Allocation Diagram for MAINID.TSK D-Space

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

Example 7-1 Map of Overlaid Task MAIN.TSK

MAIN.TSK;1 Memory allocation map TKB M40.10 Page 1
 20-OCT-82 10:08

Task name : ...CBP
 Partition name : GEN
 Identification : V00.00
 Task UIC : [240,1]
 Stack limits: 000320 001317 001000 00512.
 PRG xfr address: 002350
 Total address windows: 3.
 Task image size : 1664. words
 Task address limits: 000000 006307
 R-W disk blk limits: 000002 000012 000011 00009.

MAIN.TSK;1 Overlay description:

Base	Top	Length	
000000	004513	004514 02380.	MAIN
004514	005063	000350 00232.	INPUT
004514	004607	000074 00060.	CALC
004514	006307	001574 00892.	OUTPUT

MAIN.TSK;1 Memory allocation map TKB M40.10 Page 2
 MAIN 20-OCT-82 10:08

*** Root segment: MAIN

R/W mem limits: 000000 004513 004514 02380.
 Disk blk limits: 000002 000006 000005 00005.

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.: (RW, I, LCL, REL, CON)	001320 000466 00310.		
	001320 000250 00168.	EDDAT 03	SYSLIB.OLB;5
	001570 000216 00142.	CBTA 04.3	SYSLIB.OLB;5
COM1 : (RO, D, GBL, REL, CON)	002006 000024 00020.		
	002006 000024 00020.	MAIN V00.00	MAIN.OBJ;36
COM2 : (RW, D, GBL, REL, CON)	002032 000032 00026.		
	002032 000032 00026.	MAIN V00.00	MAIN.OBJ;36
COM3 : (RW, D, GBL, REL, CON)	002064 000010 00008.		
	002064 000010 00008.	MAIN V00.00	MAIN.OBJ;36
COM4 : (RW, D, GBL, REL, CON)	002074 000234 00156.		
	002074 000234 00156.	MAIN V00.00	MAIN.OBJ;36
COM5 : (RW, D, GBL, REL, CON)	002330 000020 00016.		

(continued on next page)

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

Example 7-1 (Cont.) Map of Overlaid Task MAIN.TSK

```

LIBROT:(RW,I,GBL,REL,CON) 002330 000020 00016. MAIN V00.00 MAIN.OBJ;36
000000 000140 00096.
000000 000140 00096. LIBROT 03.01 LIBFS0.STB;1
MAIN : (RO,I,LCL,REL,CON) 002350 000142 00098.
002350 000142 00098. MAIN V00.00 MAIN.OBJ;36
$$ALER:(RO,I,LCL,REL,CON) 002512 000024 00020.
002512 000000 00000. OVCTR 15.03 SYSLIB.OLB;5
002512 000024 00020. ALERR 02.00 SYSLIB.OLB;5
$$ALVC:(RO,I,LCL,REL,CON) 002536 000070 00056.

```

Global symbols:

```

AADD 002556-R N.DTDS 000020 $CBDAT 001570-R .FSRPT 000050
ARGBLK 002046-R N.FAST 000013 $CBDMG 001576-R .NALER 003354-R

```

```

MAIN.TSK;1 Memory allocation map TKB M40.10 Page 4
INPUT 20-OCT-82 10:08

```

*** Segment: INPUT

```

R/W mem limits: 004514 005063 000350 00232.
Disk blk limits: 000007 000007 000001 00001.

```

Memory allocation synopsis:

Section	Title	Ident	File
. BLK:(RW,I,LCL,REL,CON)	004514 000074 00060.		
	004514 000074 00060.	CATB 03	SYSLIB.OLB;5
INPUT:(RO,I,LCL,REL,CON)	004610 000252 00170.		
	004610 000252 00170.	INPUT 01	INPUT.OBJ;32
\$\$ALVC:(RO,I,LCL,REL,CON)	005062 000000 00000.		

Global symbols:

```

INPUT 004610-R $CDTB 004514-R $COTB 004522-R

```

(continued on next page)

Example 7-1 (Cont.) Map of Overlaid Task MAIN.TSK

MAIN.TSK;1 Memory allocation map TKB M40.10 Page 5
 20-OCT-82 10:08

*** Task builder statistics:

Total work file references: 13626.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 5198. words (20. pages)
 Size of work file: 4096. words (16. pages)

Elapsed time:00:00:19

Example 7-2 Map of Overlaid I- and D-Space Task MAINID.TSK

MAINID.TSK;1 Memory allocation map TKB M40.10 Page 1
 15-OCT-82 11:51

Task name : ...CBP
 Partition name : GEN
 Identification : V00.00
 Task UIC : [240,1]
 Stack limits: 000256 001255 001000 00512.
 PRG xfr address: 000744
 Task attributes: ID
 Total address windows: 4.
 Task image size : 1184. words, I-Space
 704. words, D-Space
 Task Address limits: 000000 004413 I-Space
 000000 002513 D-Space
 R-W disk blk limits: 000002 000014 000013 00011.

MAINID.TSK;1 Overlay description:

Base	Top	Length		
----	---	-----	-----	
000000	002713	002714	01484.	I MAIN
000000	002413	002414	01292.	D
002714	003263	000350	00232.	I INPUT
002414	002413	000000	00000.	D
002714	003007	000074	00060.	I CALC
002414	002413	000000	00000.	D
002714	004413	001500	00832.	I OUTPUT
002414	002513	000100	00064.	D

(continued on next page)

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

Example 7-2 (Cont.) Map of Overlaid I- and D-Space Task MAINID.TSK

MAINID.TSK;1 Memory allocation map TKB M40.10 Page 2
 MAIN 15-OCT-82 11:51

*** Root segment: MAIN

R/W mem limits: 000000 002713 002714 01484. I-Space
 000000 002413 002414 01292. D-Space

Disk blk limits: 000002 000004 000003 00003. I-Space
 000005 000007 000003 00003. D-Space

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.: (RW, I, LCL, REL, CON)	000256 000466 00310.		
	000256 000250 00168. EDDAT	03	SYSLIB.OLB;10
	000526 000216 00142. CBTA	04.3	SYSLIB.OLB;10
COM1 : (RO, D, GBL, REL, CON)	001256 000024 00020.		
	001256 000024 00020. MAIN	V00.00	MAIN.OBJ;34
COM2 : (RW, D, GBL, REL, CON)	001302 000032 00026.		
	001302 000032 00026. MAIN	V00.00	MAIN.OBJ;34
COM3 : (RW, D, GBL, REL, CON)	001334 000010 00008.		
	001334 000010 00008. MAIN	V00.00	MAIN.OBJ;34
COM4 : (RW, D, GBL, REL, CON)	001344 000234 00156.		
	001344 000234 00156. MAIN	V00.00	MAIN.OBJ;34
COM5 : (RW, D, GBL, REL, CON)	001600 000020 00016.		
	001600 000020 00016. MAIN	V00.00	MAIN.OBJ;34
LIBROT: (RW, I, GBL, REL, CON)	000000 000140 00096.		
	000000 000140 00096. LIBROT	03.01	LIBFS0.STB;1
MAIN : (RO, I, LCL, REL, CON)	000744 000142 00098.		
	000744 000142 00098. MAIN	V00.00	MAIN.OBJ;34
\$\$ALER: (RO, I, LCL, REL, CON)	001106 000024 00020.		
	001106 000000 00000. OVIDR	01	SYSLIB.OLB;10
	001106 000024 00020. ALERR	02.00	SYSLIB.OLB;10
\$\$ALVD: (RO, D, LCL, REL, CON)	001620 000034 00028.		
\$\$ALVI: (RO, I, LCL, REL, CON)	001132 000070 00056.		

Global symbols:

AADD 001152-R N.DTDS 000020 \$CBDAT 000526-R \$TIM 000402-R
 ARGBLK 001316-R N.FAST 000013 \$CBDMG 000534-R \$VEXT 000056

MAINID.TSK;1 Memory allocation map TKB M40.10 Page 4
 INPUT 15-OCT-82 11:51

(continued on next page)

USER-MODE I- AND D-SPACE (RSX-11M-PLUS ONLY)

Example 7-2 (Cont.) Map of Overlaid I- and D-Space Task MAINID.TSK

*** Segment: INPUT

R/W mem limits: 002714 003263 000350 00232. I-Space
 002414 002413 000000 00000. D-Space

Disk blk limits: 000010 000010 000001 00001. I-Space
 000011 000011 000000 00000. D-Space

Memory allocation synopsis:

Section		Title	Ident	File
. BLK.: (RW, I, LCL, REL, CON)	002714 000074 00060.			
	002714 000074 00060.	CATB	03	SYSLIB.OLB;10
INPUT : (RO, I, LCL, REL, CON)	003010 000252 00170.			
	003010 000252 00170.	INPUT	01	INPUT.OBJ;32
\$\$ALVD: (RO, D, LCL, REL, CON)	002414 000000 00000.			
\$\$ALVI: (RO, I, LCL, REL, CON)	003262 000000 00000.			
\$\$RTS : (RO, I, GBL, REL, OVR)	002710 000002 00002.			
\$\$SLVC: (RO, I, LCL, REL, CON)	003262 000000 00000.			

Global symbols:

INPUT 003010-R \$CDTB 002714-R \$COTB 002722-R

.

MAINID.TSK;1 Memory allocation map TKB M40.10 Page 5
 CALC 15-OCT-82 11:51

.

*** Task builder statistics:

Total work file references: 13920.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 5010. words (19. pages)
 Size of work file: 4096. words (16. pages)

Elapsed time:00:00:28

CHAPTER 8

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

A supervisor-mode library is a resident library that doubles a user task's virtual address space by mapping the instruction space of the processor's supervisor mode. Supervisor-mode libraries are available only on RSX-11M-PLUS systems running on PDP-11/44s and PDP-11/70s.

8.1 INTRODUCTION

A call from within a user task to a subroutine within a supervisor-mode library causes the processor to switch from user to supervisor mode. The user task transfers control to a mode-switching vector that TKB includes within the task. The mode-switching vector performs the mode switch and then transfers control to the called subroutine within the supervisor-mode library. The library routine executes with the processor in supervisor mode. When the library routine finishes executing, it transfers control to a completion routine within the library. The completion routine mode switches the processor back to user mode. The user task continues executing with the processor in user mode at the return address on the stack. This process recurs whenever the user task calls a subroutine in the supervisor-mode library.

8.2 MODE-SWITCHING VECTORS

In a task that links to a supervisor-mode library, TKB includes a 4-word, mode-switching vector in the user task's address space for each entry point referenced of a subroutine in the library.

The following shows the contents of a mode-switching vector:

```
MOV #COMPLETION-ROUTINE,-(SP)
CSM #SUPERVISOR-MODE-ROUTINE ADDRESS
```

NOTE

When mode switching from user to supervisor mode, all registers of the referencing task are preserved. All condition codes in the PS saved on the stack are cleared and must be restored by the completion routine.

8.3 COMPLETION ROUTINES

After the subroutine finishes executing, its RETURN statement transfers control to a completion routine that mode-switches from the supervisor to user mode. The completion routine returns program control back to the referencing task at the instruction after the call to the subroutine. There are two completion routines in SYSLIB:

- \$CMPCS restores only the carry bit in the user-mode PS.
- \$CMPAL restores all the condition code bits in the user-mode PS.

8.4 RESTRICTIONS ON THE CONTENTS OF SUPERVISOR-MODE LIBRARIES

The following restrictions are placed on the contents of a supervisor-mode library:

- Only subroutines using the form JSR PC, x should be used within the library.
- The library must not contain subroutines that use the stack to pass parameters.
- If both the library and the referencing task link to a subroutine from SYSLIB, then the entry point name of the subroutine must be excluded from the .STB file for the library.
- Unless you include the MSDS\$ directive within the library, the library must not contain data of any kind (even R/O) because the user supervisor D-space APRs map the user task by default. This includes user data, buffers, I/O status blocks, and directive parameter blocks (only the \$\$ directive form can be used, because the DPB for this form is pushed onto the user stack at run time).

Using the Map Supervisor D-space Executive Directive (MSDS\$), the library can map data within the instruction space of the supervisor-mode library by using the supervisor D-space APRs. The directive maps specific supervisor D-space APRs to supervisor instruction space by copying the supervisor I-space APRs that map the data portion of the library. To effectively contain data within a supervisor-mode library, you must know which APRs map the data portions of your task and library.

NOTE

You cannot use MSDS\$ to map supervisor D-space APR 0. Mapping library data and the user task simultaneously should be done with extreme care. Section 5.3.4.2 of the RSX-11M/M-PLUS Executive Reference Manual discusses the MSDS\$ directive in detail.

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

8.5 SUPERVISOR-MODE LIBRARY MAPPING

Supervisor-mode libraries are mapped with the supervisor I-space APRs. Supervisor D-space APRs can map the user task, data within the library, or both the user task and library data simultaneously. They map the user task by default.

The supervisor D-space APRs can be mapped differently according to whether the library contains data.

Supervisor D-space APRs are copies of user I-space APRs, which map the entire user task. This gives the library access to data within the user task. Figure 8-1 illustrates this mapping.

8.5.1 Supervisor Mode Library Data

Libraries that contain data require extremely complicated mapping that may overwrite the user task or cause the task to fail.

Supervisor D-space APRs are copies of user I-space APRs, which map the entire user task. For I- and D-space tasks, the supervisor D-spaces APRs are copies of the user D-space APRs. Including the MSDS\$ Executive directive (see Section 8-4) within the library code enables the library to map data within its own instruction space. The user task may be overmapped. The library has access to data within its instruction space and to data in the user task that is not overmapped by the MSDS\$ directive.

Figure 8-2 illustrates this mapping.

8.5.2 Supervisor Mode Libraries with I- and D-Space Tasks

I- and D-space tasks may link to supervisor-mode libraries. Instead of mapping to the entire user task, the supervisor-mode library's D-space APRs map the task's data space. Because the I- and D-space task maps its data with the D-space APRs, the task's D-space APRs are copied into the supervisor-mode library's D-space APRs. Therefore, the supervisor-mode library maps its own instructions with supervisor-mode I-space APRs and maps the task's data with supervisor-mode D-space APRs.

Figure 8-3 illustrates the mapping of an I- and D-space task linked to a supervisor mode library.

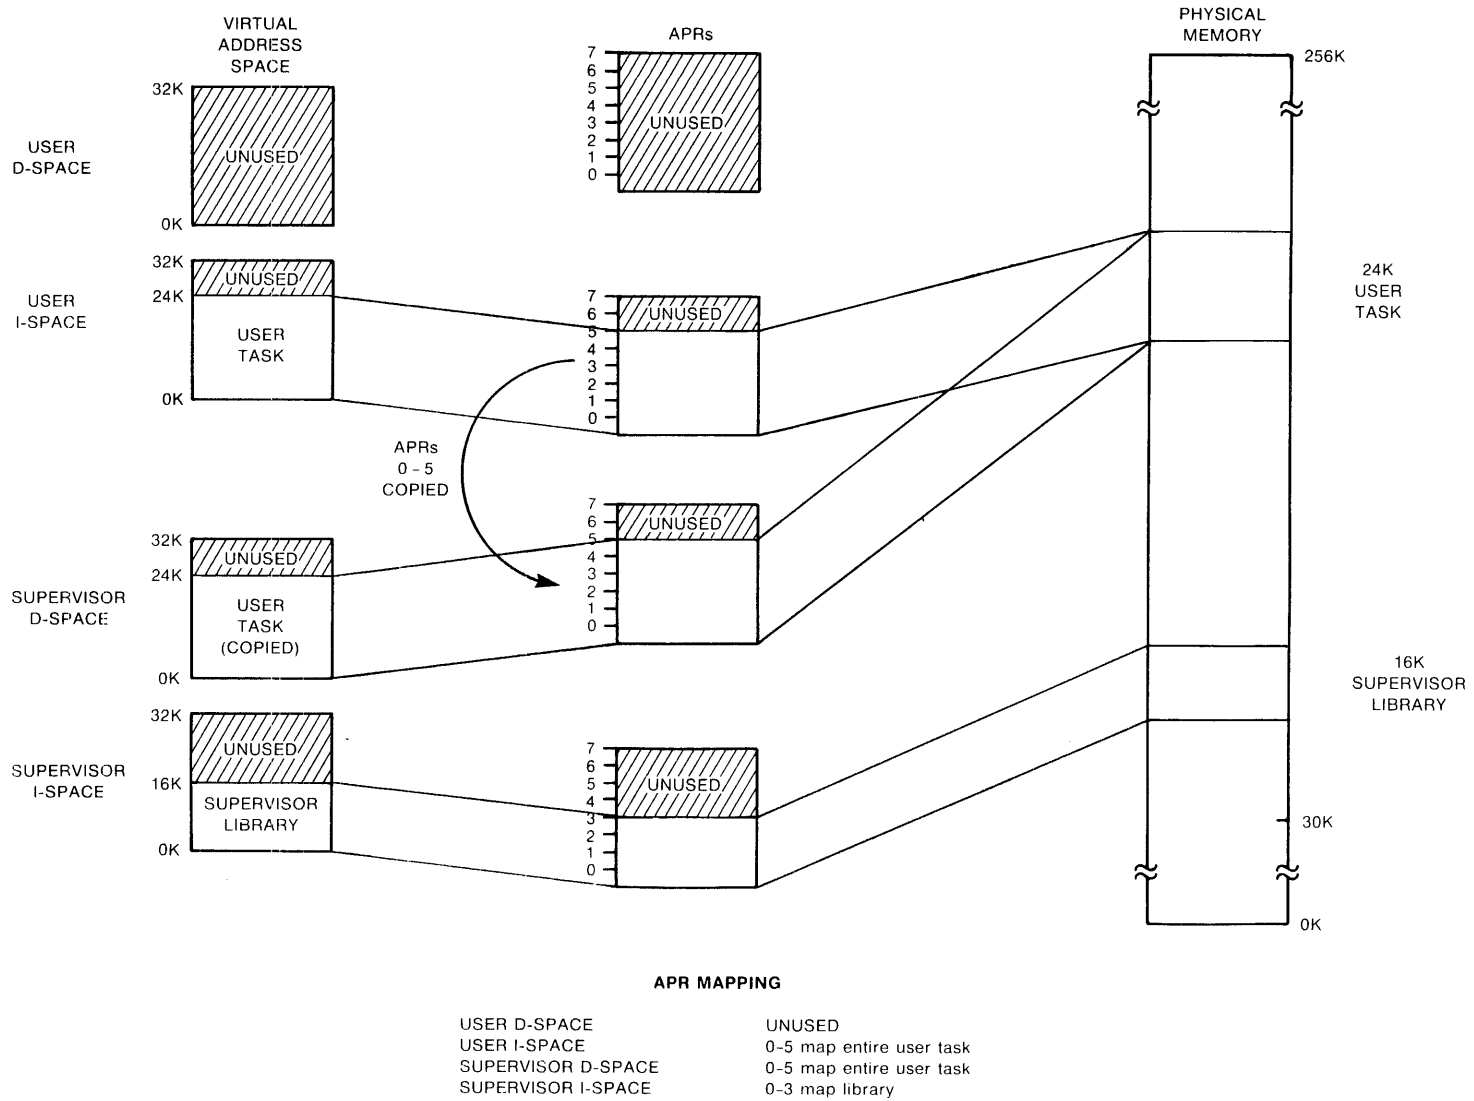
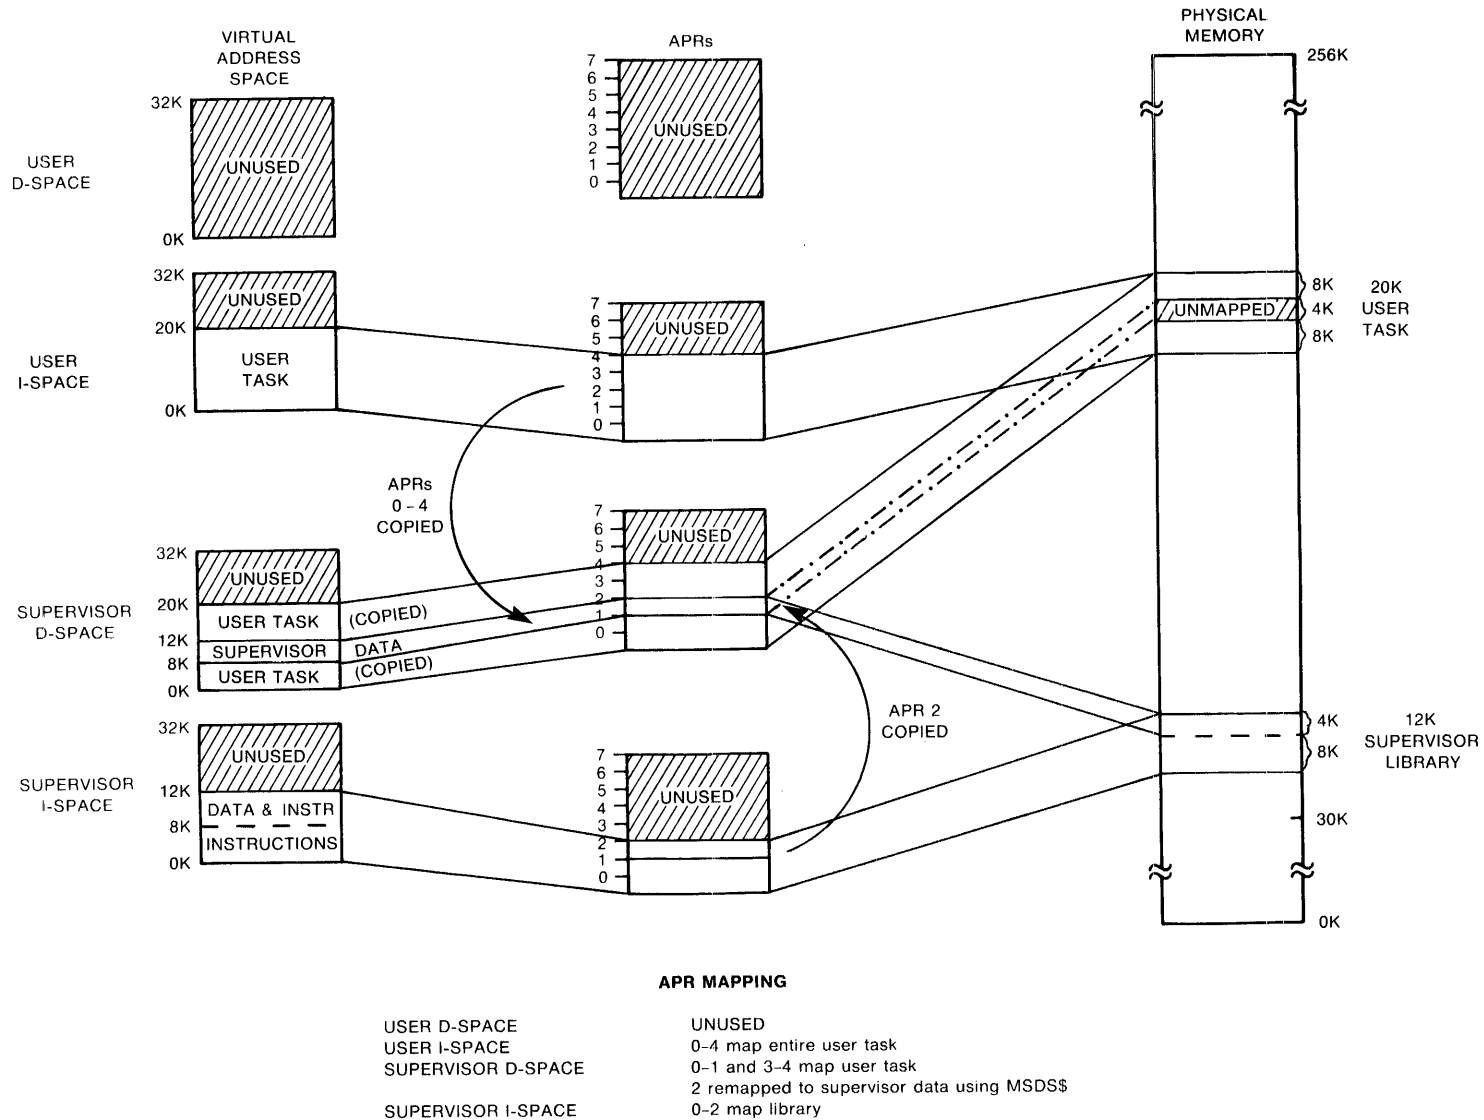
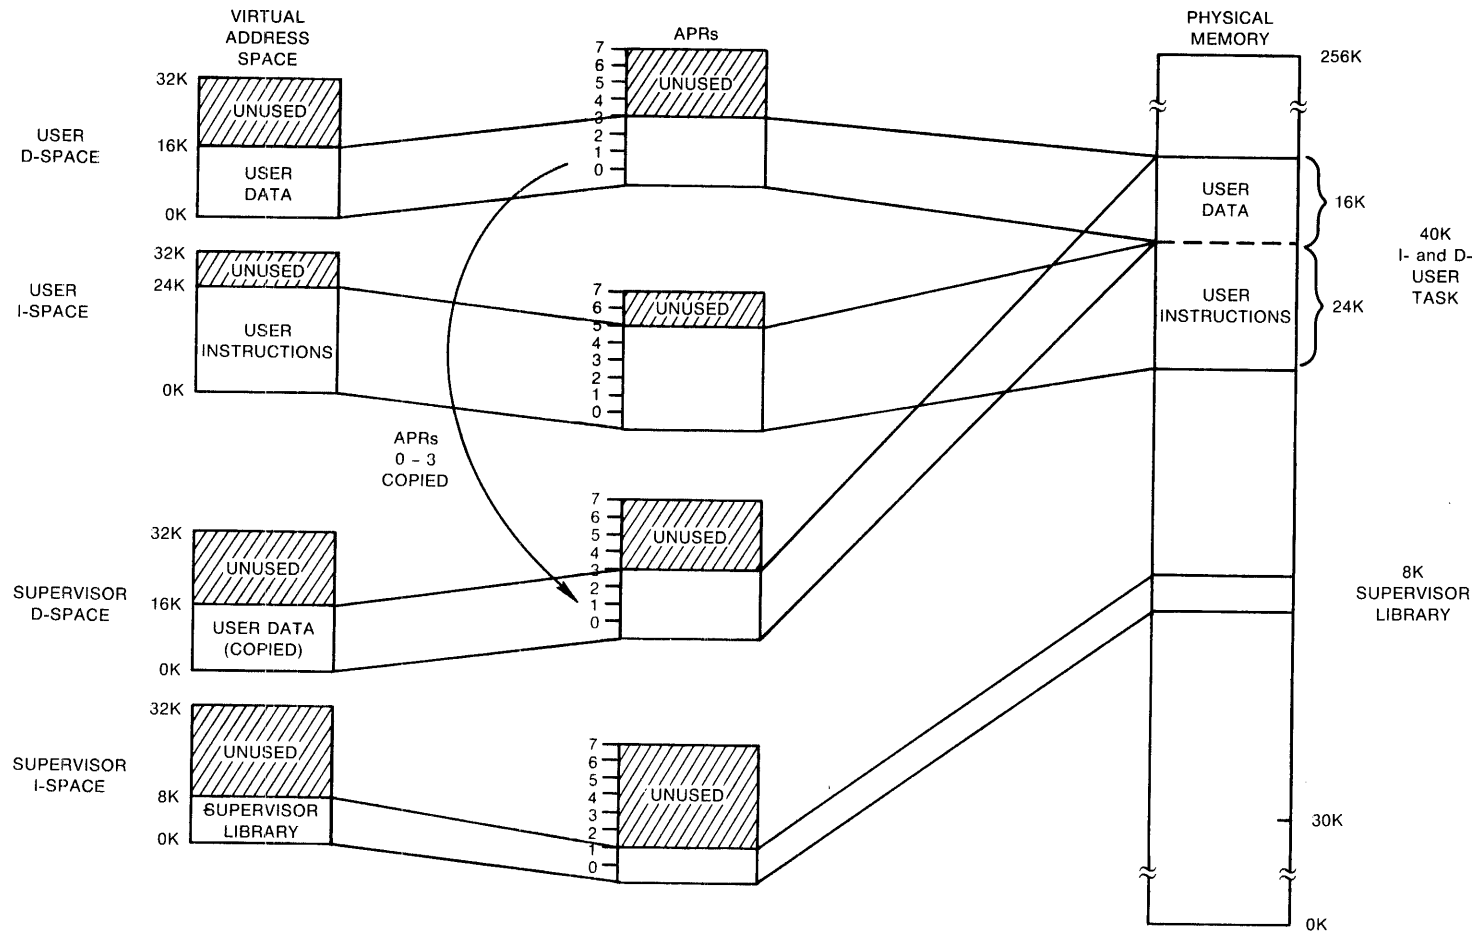


Figure 8-1 Mapping of a 24K Conventional User Task That Links to a 16K Supervisor-Mode Library



SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

Figure 8-2 Mapping of a 20K Conventional User Task That Links to a 12K Supervisor-Mode Library Containing 4K of Data



APR MAPPING

USER D-SPACE	0-3 map user data
USER I-SPACE	0-5 map user instructions
SUPERVISOR D-SPACE	0-3 map user data
SUPERVISOR I-SPACE	0-1 map library

ZK-1105-82

Figure 8-3 Mapping of a 40K I- and D-Space Task That Links to an 8K Supervisor-Mode Library

8.6 BUILDING AND LINKING TO SUPERVISOR-MODE LIBRARIES

Building and linking to a supervisor-mode library is essentially the same as building and linking to a conventional resident library (discussed in Chapter 5). When you build a supervisor-mode library, you suppress the header by attaching `/-HD` to the task image file. During option input, you suppress the stack area by specifying `STACK=0`. You specify the partition in which the library is to reside and, optionally, the base address and length of the library with the `PAR` option.

8.6.1 Relevant TKB Options

Use the following options to build and reference supervisor-mode libraries:

<code>CMVRT</code>	Indicates that you are building supervisor-mode library and specifies the name of the completion routine.
<code>RESSUP (SUPLIB)</code>	Indicates that your task references a supervisor-mode library.
<code>GBLXCL</code>	Excludes a global symbol from the <code>.STB</code> file of the supervisor-mode library.

These options are discussed briefly below and are fully documented in Chapter 11.

8.6.2 Building the Library

You indicate to the Task Builder that you are building a supervisor-mode library with the `CMVRT` option. The argument for this option identifies the entry symbol of the completion routine. When the Task Builder processes this option, it places the completion routine entry point in the library's `STB` file. To exclude a global symbol from the library's `.STB` file, you specify the name of the global symbol as the argument of the `GBLXCL` option. You must exclude from the `.STB` file of a supervisor-mode library any symbol defined in the library that represents the following:

- An entry point to a subroutine that uses the stack to pass parameters
- An entry point to a subroutine mapped in user mode that the referencing user task calls

8.6.3 Building the Referencing Task

When you build a task that references a supervisor-mode library, use the `RESSUP` option if you are referencing a user-owned, supervisor-mode library and `SUPLIB` if you are referencing a system-owned, supervisor-mode library. (Like the `RESLIB` and `LIBR` options for linking to conventional libraries, `RESSUP` and `SUPLIB` are functionally the same.) The arguments for these options are:

- The filespec (`RESSUP` option) or name (`SUPLIB`) of the library to be referenced

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

- A switch that tells TKB whether to use system-supplied vectors to perform mode switching from user to supervisor mode.
- For position-independent libraries, the first available supervisor-mode I-space APR that you want to map the library.

8.6.4 Mode Switching Instruction

Mode switching occurs with a new instruction available on the 11/44 and emulated by the Executive on the 11/70. Throughout the remainder of the chapter, supervisor-mode libraries are referred to as CSM (change supervisor mode) libraries.

8.7 CSM LIBRARIES

This section discusses how you build and link to CSM libraries. It also shows an extended example of building and linking to a CSM library and explains the context-switching vectors and completion routines for CSM libraries.

8.7.1 Building a CSM Library

You indicate to the Task Builder that you are building a CSM library by specifying the name of the completion routine as the argument for the CMPRT option. This option places the name of the completion routine into the library's .STB file. Link the completion routine, either \$CMPAL or \$CMPCS, located in LB:[1,2]SYSLIB.OLB, as the first input file. Although the completion routines are located in SYSLIB (which is ordinarily referenced by default), you must explicitly indicate it and link it as the first input file. You must also specify in the PAR option a 0 base for the partition in which the library will reside. These two steps locate the completion routine at virtual 0 of the library's virtual address space.

You specify the name of any global symbols that you would like to exclude from the library's .STB file as the argument to the GBLXCL option. You must exclude from the .STB file of a supervisor-mode library any symbol defined in the library that represents the following:

- An entry point to a subroutine that uses the stack to pass parameters
- An entry point to a subroutine mapped in user mode that the referencing user task calls

A sample TKB sequence for building a CSM library in UFD [301,55] on SY: follows:

```
TKB>CSM/--HD/LI/PI,CSM/MA,CSM=  
TKB>LB:[[1,2]]SYSLIB/LB:CM PAL,SY:[[301,55]]CSM  
TKB>/  
Enter Options:  
TKB>STACK=0  
TKB>PAR=GEN:0:2000  
TKB>CMPRT=$CMPCS  
TKB>GBLXCL=$SAVAL  
TKB>///
```

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

The library is built without a header or stack, like all shared regions. It is position independent and has only one program section named .ABS. The /LI switch accomplishes this, eliminating program section name conflicts between the library and the referencing task.

The completion routine module of SYSLIB, CMPAL, is specified first in the input line. The library will run in partition GEN at 0 and is not more than 1K. These are two aspects of building supervisor-mode libraries specific to CSM libraries: the completion routine must be linked first, and must reside at virtual 0. Why the CSM library must reside at virtual 0 is discussed in Section 8.8.5.

The CMPRT option specifies the global symbol \$CMPCS, which is the entry point of the completion routine. Note that the SYSLIB module name is "CMPCS" and its corresponding global symbol is "\$CMPCS".

The GBLXCL option excludes \$SAVAL from the library's .STB file because the user task must reference a copy of \$SAVAL that is mapped with user mode APRs.

8.7.2 Linking to a CSM Library

If your task links to a user-owned CSM library, you use the RESSUP option. If your task links to a system-owned CSM library, you use the SUPLIB option. These options tell TKB that the task will link to a supervisor-mode library. The option takes up to three arguments:

- The filespec (RESSUP option) or name (SUPLIB option) of the library
- A switch that tells the Task Builder whether to use system-supplied, mode-switching vectors
- For position-independent libraries, an APR that must be APR 0 so that the library's completion routine is mapped at virtual 0.

This information enables the Task Builder to find the .STB file for the CSM library, include a 4-word, mode-switching vector within the user task for each call to a subroutine within the library, and correctly map the library at virtual 0 in the library image.

The following sample task-build command sequence builds a task named REF, which references the library SUPER that you built in the previous section:

```
TKB>REF,REF=REF
TKB>/
Enter Options:
TKB>RESSUP=SUPER/SV:0
TKB>///
```

This sequence tells TKB to include in the logical address space of REF a user-owned, supervisor-mode library named SUPER. TKB will include a 4-word mode switching vector within the user task for each call to a subroutine within the library. The CSM library is position independent and will be mapped with APR 0.

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

8.7.3 Example CSM Library and Linking Task

This example shows you the code, maps, and TKB command sequences for building and linking to a CSM library that contains no data in a system not having user data space. Example 8-1 shows the code for the library SUPER, and Example 8-2 shows its accompanying map. Example 8-3 shows the code for the completion routine \$CMPCS that is linked in to SUPER from SYSLIB. Example 8-4 shows the code for referencing task TSUP, and Example 8-5 shows its accompanying map.

Example 8-1 Code for SUPER.MAC

```

.TITLE SUPER
.IDENT /01/

SORT::
CALL    $$AVAL          ; SAVE ALL REGISTERS
TST     (R5)+           ; SKIP OVER NUMBER OF ARGUMENTS
MOV     (R5)+,R0        ; GET ADDRESS OF LIST
MOV     (R5)+,R4        ; GET ADDRESS OF LENGTH OF LIST
MOV     (R4),R4         ; GET LENGTH OF LIST
BEQ     40$             ; IF NO ARGUMENTS
MOV     R0,R5           ;
DEC     R4              ;
10$:    MOV     R5,R0     ; COPY
        MOV     R4,R3     ; COPY LENGTH OF LIST
20$:    TST     (R0)+     ; MOVE POINTER TO NEXT ITEM
        CMP     (R5),(R0) ; COMPARE ITEMS
        BLE     30$      ; IF LE IN CORRECT ORDER
        MOV     (R5),R2   ; SWAP ITEMS
        MOV     (R0),(R5) ;
        MOV     R2,(R0)   ;
30$:    DEC     R3        ; DECREMENT LOOP COUNT
        BGE     20$      ; IF NE LOOP
        DEC     R4        ; DECREMENT
        BLE     40$      ; IF EQ SORT COMPLETED
        TST     (R5)+    ; GET POINTER TO NEXT ITEM TO BE COMPARED
        BR     10$
40$:    RETURN

SEARCH::
CALL    $$AVAL          ; SAVE ALL THE REGISTERS
CMP     #4,(R5)+        ; FOUR ARGUMENTS?
BNE     20$             ; IF NE NO
MOV     (R5)+,R0        ; GET ADDRESS OF NUMBER TO LOCATE
MOV     (R5)+,R1        ; ADDRESS OF LIST SEARCHING
MOV     (R5)+,R2        ; GET ADDRESS OF LENGTH OF LIST
MOV     (R2),R2         ; GET LENGTH OF LIST
BEQ     20$             ; IF NO ARGUMENTS
MOV     (R5),R5         ; ADDRESS OF RETURNED VALUE
MOV     R2,R3          ; COPY LENGTH
10$:    CMP     (R0),(R1)+ ; IS THIS THE NUMBER?
        BEQ     30$      ; IF EQ YES
        BMI     20$      ; IF MI NUMBER NOT THERE
        DEC     R2        ; DECREMENT LOOP COUNT
        BNE     10$      ; IF NE NOT AT END OF LIST
    
```

(continued on next page)

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

Example 8-1 (Cont.) Code for SUPER.MAC

```

20$:      MOV      #-1,(R5)          ; END OF LIST PASS BACK ERROR
          RETURN
30$:      SUB      R2,R3             ; NUMBER FOUND - GET INDEX INTO LIST
          INC      R3               ;
          MOV      R3,(R5)          ; RETURN INDEX
          RETURN
          .END
    
```

Example 8-2 Memory Allocation Map for SUPER

SUPER.TSK;1 Memory allocation map TKB M40.10 Page 1
 29-DEC-82 15:04

Partition name : GEN
 Identification : 0203
 Task UIC : [301,55]
 Task attributes: -HD,PI
 Total address windows: 1.
 Task image size : 160. words
 Task address limits: 000000 000473
 R-W disk blk limits: 000002 000002 000001 00001.
 *** Root segment: CMPAL

R/W mem limits: 000000 000473 000474 00316.
 Disk blk limits: 000002 000002 000001 00001.

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.:(RW,I,LCL,REL,CON)	000000 000474 00316.		
	000000 000136 00094.	CMPAL 0203	SYSLIB.OLB;6
	000136 000136 00094.	CMPAL 0203	SYSLIB.OLB;6
	000274 000136 00094.	SUPER 01	SUPER.OBJ;3
	000432 000042 00034.	SAVAL 00	SYSLIB.OLB;6

Global symbols:

SEARCH 000352-R SORT 000274-R \$CMPAL 000022-R \$CMPCS 000110-R \$SAVAL 000432-R

*** Task builder statistics:

Total work file references: 320.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 6988. words (27. pages)
 Size of work file: 1024. words (4. pages)
 Elapsed time:00:00:04

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

Example 8-3 Completion Routine, \$CMPCS, from SYSLIB.OLD

```
.TITLE CMPAL
.IDENT /0204/
```

```
;
;          COPYRIGHT (c) 1983 BY
;          DIGITAL EQUIPMENT CORPORATION, MAYNARD
;          MASSACHUSETTS. ALL RIGHTS RESERVED.
;
; THIS SOFTWARE IS FURNISHED UNDER A LICENSE AND MAY BE USED
; AND COPIED ONLY IN ACCORDANCE WITH THE TERMS OF SUCH LICENSE
; AND WITH THE INCLUSION OF THE ABOVE COPYRIGHT NOTICE. THIS
; SOFTWARE OR ANY OTHER COPIES THEREOF, MAY NOT BE PROVIDED OR
; OTHERWISE MADE AVAILABLE TO ANY OTHER PERSON. NO TITLE TO AND
; OWNERSHIP OF THE SOFTWARE IS HEREBY TRANSFERED.
;
; THE INFORMATION IN THIS DOCUMENT IS SUBJECT TO CHANGE WITHOUT
; NOTICE AND SHOULD NOT BE CONSTRUED AS A COMMITMENT BY DIGITAL
; EQUIPMENT CORPORATION.
;
; DIGITAL ASSUMES NO RESPONSIBILITY FOR THE USE OR RELIABILITY OF
; ITS SOFTWARE ON EQUIPMENT THAT IS NOT SUPPLIED BY DIGITAL.
;
;          .ENABL LC
;
;
; This module supports the "new" transfer vector format generated by
; the taskbuilder for entering super mode libraries. This format
; optimized for speed and size and supports user data space tasks.
;
; The CSM dispatcher routine and the standard completion routines,
; $CMPAL and $CMPCS are included in this module due to the close
; interaction between them.
;
;
; **--CSM Dispatcher-Dispatch CSM entry
;
; This module must be linked at virtual zero in the supervisor mode
; library. It is entered via a four word transfer vector of the form:
;
;          MOV      #completion-routine,-(SP)
;          CSM      #routine
;
; Note: Immediate mode emulation of the CSM instruction is required
; in the executive.
;
; The CSM instruction transfers control to the address contained in
; supervisor mode virtual 10. At this point the stack is the following:
;
;          (SP) routine address
;          2(SP) PC (past end of transfer vector)
;          4(SP) PS with condition codes cleared
;          6(SP) Completion-routine address
;          10(SP) Return address
;
```

(continued on next page)

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

Example 8-3 (Cont.) Completion Routine, \$CMPCS, from SYSLIB.OLD

```

; A routine address of 0 is special cased to support return to
; supervisor mode from a user mode debugging aid (ODT). In this case
; stack is the following:
;
;      (SP) zero
;      2(SP) PC from CSM to be discarded
;      4(SP) PS from CSM to be discarded
;      6(SP) Super mode PC supplied by debugger
;      10(SP) Super mode PS supplied by debugger
;
; To allow positioning at virtual zero, this code must be in the blank
; PSECT which is first in the TKBs PSECT ordering.

.PSECT
.ENABL LSB

; Debugger return to super mode entry. Must start at virtual zero

CMP      (SP)+,(SP)+      ; Clean off PS and PC from CSM

;
; **-$SRTI-SUPER mode RTI
;
; This entry point performs the necessary stack management to allow
; an RTI from super mode to either super mode or user mode.
; The is as required for an RTI:
;
;      (SP) Super mode PC
;      2(SP) Super mode PS

$SRTI:: TST      2(SP)      ; Returning to user mode?
BR      70$             ; Join common code

; CSM transfer address, this word must be at virtual 10 in super mode

.WORD   CSMSVR          ; CSM dispatcher entry

; Dispatch CSM entry

CSMSVR: MOV      6(SP),2(SP) ; Set completion routine address for RETURN
JMP     @(SP)+         ; Transfer to super mode library routine

;
; **-$CMPAL-Completion routine which sets up NZVC in the PS
;
; Copy all condition codes to stacked PS. Current stack:
;
;      (SP) PS with condtion codes cleared
;      2(SP) Completion routine address (to be discarded)
;      4(SP) Return address

$CMPAL::BPL      40$      ;
BNE     20$          ;
BVC     10$          ;
BIS     #16,(SP)     ; Set NZV
BR      $CMPCS      ;
10$:   BIS     #14,(SP) ; Set NZ
BR      $CMPCS      ;

```

(continued on next page)

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

Example 8-3 (Cont.) Completion Routine, \$CMPCS, from SYSLIB.OLD

```

20$:   BVC     30$           ;
      BIS     #12,(SP)      ; Set NV
      BR      $CMPCS       ;
30$:   BIS     #10,(SP)     ; Set N
      BR      $CMPCS       ;
40$:   BNE     60$         ;
      BVC     50$         ;
      BIS     #6,(SP)      ; Set ZV
      BR      $CMPCS       ;
50$:   BIS     #4,(SP)     ; Set Z
      BR      $CMPCS       ;
60$:   BVC     $CMPCS       ;
      BIS     #2,(SP)     ; Set V

;
; **-$CMPCS-Completion routine which sets up only C in the PS
;
; Copy only carry to stacked PS. Current stack:
;
;      (SP)   PS with condition codes cleared
;      2(SP)  Completion routine address (to be discarded)
;      4(SP)  Return address
;
$CMPCS: ADC     (SP)         ; Set up carry
      MOV     4(SP),2(SP)   ; Setup return address for RTT
      MOV     (SP)+,2(SP)   ; And PS. Returning to super mode?
70$:   BPL     80$         ; If PL yes
      MOV     #6,-(SP)     ; Number of bytes for (SP), PS, and PC
      ADD     SP,(SP)      ; Compute clean stack value
      MTPI   SP           ; Set up previous stack pointer
80$:   RTT                    ; Return to previous mode and caller

      .DSABL  LSB
      .END

```

Example 8-4 Code for TSUP.MAC

```

      .TITLE  TSUP
      .IDENT  /01/

      .MCALL  QIOW$,DIR$,QIOW$$
WRITE: QIOW$  IO.WVB,5,1,,,,<OUT,,40>
READIN: QIOW$ IO.RVB,5,1,,,,<OUT,5>

IARRAY: .BLKW 12.
LEN:    .BLKW 1
IART:   .BLKW 1
INDEX:  .WORD 0
OUT:    .BLKW 100.
ARGBLK:
EDBUF:  .BLKW 10.

```

(continued on next page)

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

Example 8-4 (Cont.) Code for TSUP.MAC

```

FMT1: .ASCIZ  /%2SARRAY(%D)=/
FMT2: .ASCIZ  /%N%2SNUMBER TO SEARCH FOR?/
FMT3: .ASCIZ  /%N%2S%D WAS FOUND IN ARRAY(%D)/
FMT4: .ASCIZ  /%N%2S%D WAS NOT IN ARRAY/
FMT5: .ASCIZ  /%2SARRAY(%D)=%D/

      .EVEN
START:
      MOV     #IARRAY,R0      ; GET ADDRESS OF ARRAY
      MOV     #10,R1         ; SET LENGTH OF ARRAY
5$:   CLR     (R0)+           ; INITIALIZE ARRAY
      DEC     R1             ; LOOP
      BNE    5$
      MOV     #IARRAY,R0
      MOV     #INDEX,R2
10$:  MOV     #FMT1,R1       ; FORMAT SPECIFICATION (ADDRESS
      ; OF INPUT STRING)
      MOV     (R2),EDBUF     ; GET INDEX
      INC     EDBUF         ;
      CALL   PRINT          ; PRINT MESSAGE
      CALL   READ           ; READ INPUT
      MOV     IART,(R0)+    ; PUT BINARY KEYBOARD INPUT INTO ARRAY
      BEQ    20$           ; ZERO MARKS END OF INPUT
      INC     (R2)         ;
      CMP     (R2),#10.    ;
      BNE    10$           ; IF NE YES
20$:  MOV     (R2),LEN      ; CALCULATE LENGTH OF ARRAY
      MOV     #ARGBLK,R5    ; GET ADDRESS OF ARGUMENT BLOCK
      MOV     #2,(R5)+     ; NUMBER OF ARGUMENTS
      MOV     #IARRAY,(R5)+ ; PUT ADDRESS OF ARRAY
      MOV     #LEN,(R5)    ;
      MOV     #ARGBLK,R5   ;
      CALL   SORT          ; SORT ARRAY
;+
;Task Builder replaced call to SORT subroutine in SUPLIB with 4-word
;context switching vector. Flow of control switches to SUPLIB via
;the vector and back via the completion routine $CMPCS. TSUP
;continues excuting at the next instruction.
;-
      CLR     R2            ;
      MOV     #IARRAY,R0   ; GET ARRAY ADDRESS
30$:  INC     R2            ; INCREMENT INDEX
      MOV     R2,EDBUF     ; GET INDEX FOR PRINT
      MOV     (R0)+,EDBUF+2 ; GET CONTENTS OF ARRAY
      MOV     #FMT5,R1     ; GET ADDRESS OF FORMAT SPECIFICATION
      CALL   PRINT        ;
      CMP     R2,LEN      ; MORE TO PRINT?
      BLT    30$         ; IF LE YES
      MOV     #FMT2,R1    ; GET ADDRESS OF FORMAT SPECIFICATION

```

(continued on next page)

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

Example 8-4 (Cont.) Code for TSUP.MAC

```

CALL PRINT           ; OUTPUT MESSAGE
CALL READ           ; READ RESPONSE
MOV #ARGBLK,R5      ;
MOV #4,(R5)+        ; SET NUMBER OF ARGUMENTS
MOV #IART,(R5)+     ; SET ADDRESS OF NUMBER LOOKING FOR
MOV #IARRAY,(R5)+   ; SET ADDRESS OF ARRAY
MOV #LEN,(R5)+      ; SET ADDRESS OF LEN OF ARRAY
MOV #INDEX,(R5)     ; ADDRESS OF RESULT
MOV #ARGBLK,R5      ;
CALL SEARCH         ; SEARCH FOR NUMBER IN IART
;
;Call to SUPLIB for SEARCH subroutine.
;
TST INDEX           ; WAS NUMBER FOUND?
BLT 40$             ; IF LT NO
MOV IART,EDBUF      ; GET NUMBER LOOKING FOR
MOV INDEX,EDBUF+2   ; GET ARRAY NUMBER
MOV #FMT3,R1        ; GET FORMAT ADDRESS
CALL PRINT          ;
BR 100$             ; DONE
40$:
MOV #FMT4,R1        ; GET FORMAT ADDRESS
MOV IART,EDBUF      ; GET NUMBER
CALL PRINT
100$:
CALL $EXST          ; EXIT WITH STATUS
PRINT:
CALL $SAVAL         ; SAVE ALL REGISTERS
MOV #OUT,R0         ; ADDRESS OF OUTPUT BLOCK
MOV #EDBUF,R2       ; START ADDRESS OF ARGUMENT BLOCK
CALL $EDMSG         ; FORMAT MESSAGE
MOV R1,WRITE+Q.IOPL+2 ; PUT LENGTH OF OUTPUT
; BLOCK INTO PARAMETER BLOCK
DIR$ #WRITE         ; WRITE OUTPUT BLOCK
RETURN
READ:
CALL $SAVAL         ; SAVE ALL REGISTERS
DIR$ #READIN ; READ REQUEST
MOV #OUT,R0 ; GET KEYBOARD INPUT
CALL $CDTB         ; CONVERT KEYBOARD INPUT TO BINARY
MOV R1,IART        ; PUT INPUT INTO BUFFER
RETURN
.END START

```

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

Example 8-5 Memory Allocation Map for TSUP

TSUP.TSK;1 Memory allocation map TKB M40.10 Page 1
 29-DEC-82 15:01

Partition name : GEN
 Identification : 01
 Task OIC : [301,55]
 Stack limits: 000274 001273 001000 00512.
 PRG xfr address: 002130
 Total address windows: 2.
 Task image size : 1344. words
 Task address limits: 000000 005133
 R-W disk blk limits: 000002 000007 000006 00006.

*** Root segment: TSUP

R/W mem limits: 000000 005133 005134 02652.
 Disk blk limits: 000002 000007 000006 00006.

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.:(RW,I,LCL,REL,CON)	001274 002334 01244.		
	001274 001234 00668.	TSUP	01 TSUP.OBJ;22
CMPAL : (RW,I,LCL,REL,CON)	000000 000474 00316.		
PURSD : (RO,I,LCL,REL,CON)	003630 000076 00062.		
PURSI : (RO,I,LCL,REL,CON)	003726 000752 00490.		
\$\$RESL:(RO,I,LCL,REL,CON)	004700 000212 00138.		
\$\$SLVC:(RO,I,LCL,REL,CON)	005112 000020 00016.		

TSUP.TSK;1 Memory allocation map TKB M40.10 Page 2
 29-DEC-82 15:01

*** Task builder statistics:

Total work file references: 2477.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 6988. words (27. pages)
 Size of work file: 1024. words (4. pages)

Elapsed time:00:00:05

TSUP prompts you to enter numbers at your terminal. It calls a subroutine in SUPER to sort the numbers. Then it displays the numbers you entered as array entries and prompts you to request a number to search for. TSUP calls a subroutine in SUPERLIB to search for the number. Finally, TSUP indicates at your terminal either that the number was not found or the array location in which the number is stored.

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

8.7.3.1 Building SUPER - To build SUPER in UFD [301,55] on SY:, use the following task-build command sequence:

```
TKB>SUPER/-HD/LI/PI,SUPER/MA,SUPER=  
TKB>LB: [[1,2]]SYSLIB/LB:CMPAL,SY: [[301,55]]SUPER  
TKB>/  
Enter Options:  
TKB>STACK=0  
TKB>PAR=GEN:0:2000  
TKB>CMPRT=$CMPCS  
TKB>GBLXCL=$SAVAL  
TKB>//
```

SUPER is build without a header or stack. It is position independent and has only one program section, named .BLK. The /LI switch accomplishes this, eliminating program section name conflicts between the library and the referencing task.

The completion routine module of SYSLIB, CMPAL, is specified first in the input line. The library will run in partition GEN at 0 and is not more than lk.

The GBLXCL option excludes \$SAVAL from the library's .STB file. You exclude \$SAVAL from the .STB file because the referencing task, TSUP, also calls \$SAVAL. If TSUP finds \$SAVAL in the .STB file of SUPER, it will not link a separate copy of \$SAVAL into its task image from SYSLIB. If TSUP cannot link to a copy of \$SAVAL that is mapped through user APRs, the TSUP would call \$SAVAL as a subroutine residing within the supervisor-mode library, but without the necessary mode-switching vector and completion routine support. This option forces TKB to link \$SAVAL from SYSLIB into the task image for TSUP.

The memory allocation map in Example 8-2 shows the following:

- SUPER begins at virtual 0.
- The completion routine, \$CMPAL, is linked into the library from SYSLIB at virtual 0.
- The entry point \$CMPAL is located at virtual 22, SEARCH is located at 35, and sort is located at 274. All of these entry points are relocatable.

8.7.3.2 Building TSUP - Use the following task-build command sequence to build a task, TSUP, that links to SUPER:

```
TKB>TSUP,TSUP=TSUP  
TKB>/  
Enter Options:  
TKB>RESSUP=SUPER/SV:0  
TKB>//
```

This command sequence tells TKB to include in the logical address space of TSUP a user-owned, supervisor-mode library named SUPER. TKB includes a 4-word, mode-switching vector within the task image for each call to a subroutine within the library. The library is position independent and is mapped with supervisor I-space APR0. This is a requirement for CSM libraries because the CSM expects to find the entry point of the completion routine at location 10.

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

The memory allocation map for TSUP (Example 8-5) shows:

- \$CMPAL is linked from the .STB file of the library and begins at location 0.
- The mode-switching vectors begin at 005136 and are 16 bytes. That means that TSUP calls subroutines within the library 2 times (4 words per vector).
- The initiation routine \$SUPL is located at 4700.
- The SEARCH and SORT subroutines that were located at virtual 112 and 32, respectively, in the virtual address space of SUPER have been relocated to the mode-switching vectors residing at 5136 and 5146 respectively, in TSUP.
- The SAVAL module from SYSLIB containing \$SAVAL has been linked into the task image instead of including \$SAVAL from the library's .STB file.

8.7.3.3 Running TSUP - After building SUPER and TSUP as indicated in the task-build command sequence discussed previously, you install SUPER and run TSUP. TSUP prompts you for a number:

```
ARRAY (x)
```

```
x
```

The position in which to store the number in the array. You enter a number. TSUP stores the number in the array and prompts you again for a number. This continues until you either have entered a 0, an illegal number, or 10 numbers. Then TSUP calls the SORT routine in SUPER.

You enter a number. TSUP calls the SEARCH routine in SUPER. Then TSUP outputs a message indicating whether the number was in the array.

8.7.4 The CSM Library Dispatching Process

When you build the referencing task, if you specify the SV argument to the RESSUP or SUPLIB option, then TKB includes a 4-word context-switching vector for each call to a subroutine in the library. This has been very generally discussed in Section 8.2. This section discusses the CSM library vector in detail.

CSM mode switching occurs as follows:

1. The vector is entered with the return address on top of the stack (TOS).
2. The vector pushes the completion routine address on the stack.
3. A CSM instruction is executed with the supervisor-mode entry point as the immediate addressing mode parameter. The CSM instruction:
 - a. Evaluates the source parameter and stores the entry point address in a temporary register

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

- b. Copies the user stack pointer to the supervisor stack pointer
- c. Places the current PS and PC on the supervisor stack clearing the condition codes in the PS
- d. Pushes the entry point address on the supervisor stack
- e. Places the contents of location 10 in supervisor I-space into the PC

The stack looks like this when the processor begins to execute at the contents of virtual 10 in supervisor mode:

```
user sp ----> return address
               completion routine address
               PS
               PC

super sp ----> entry point address
```

The most important aspect of how the CSM library mode-switching vector works is that the processor begins executing at the contents of virtual 10 in supervisor mode. This is why the completion routine must be located at virtual 0, so that virtual location 10 is within the completion routine.

8.8 CONVERTING SCAL LIBRARIES TO CSM LIBRARIES

You can easily convert your SCAL libraries to CSM libraries. Rebuilding a task on an RSX-11M-PLUS V2.0 system or later that linked to a library on a V1.0 system requires that you rebuild the library also. Rebuild the library specifying the completion routine as the first input module. If the library was not built to run at a starting address of 0 in its partition, rebuild it so that it does begin at 0 to enable TKB to find the completion routine.

8.9 USING SUPERVISOR-MODE LIBRARIES AS RESIDENT LIBRARIES

Supervisor-mode libraries can double as conventional resident libraries. For position-independent, supervisor-mode libraries, you rebuild the referencing task using the RESLIB option instead of the RESSUP option. Indicate the first available user-mode APR that you want to map the library. For CSM libraries this will always change, because you cannot map a shared region with APR 0. You do not have to rebuild the library.

For absolute supervisor-mode libraries, rebuild the referencing task using the RESLIB option instead of the RESSUP option. Rebuild the library only if the beginning partition address in the PAR option is incompatible with the address limits of your referencing task.

8.10 MULTIPLE SUPERVISOR-MODE LIBRARIES

A user task can reference multiple supervisor-mode CSM libraries. However, all the CSM libraries must use the completion routine that begins at virtual zero in supervisor-mode instruction space.

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)

8.11 LINKING A RESIDENT LIBRARY TO A SUPERVISOR-MODE LIBRARY

You can link a conventional resident library to a supervisor-mode library using the following command sequence:

```
TKB>F4PRES/--HD,F4PRES,F4PRES,LB:[1,1]F4PRES=  
TKB>F4PRES/LB  
TKB>/  
Enter Options:  
TKB>STACK=0  
TKB>SUPLIB=FCSFSL:SV  
TKB>PAR=F4PRES:140000:20000  
TKB>//
```

This command sequence shows you how to link F4PRES to FCSFSL.

8.12 LINKING SUPERVISOR-MODE LIBRARIES

You cannot link supervisor-mode libraries together, and you cannot link a supervisor-mode library to a resident user-mode library. Calling a user-mode library is not possible because its code is not mapped through the I-space APRs while in the supervisor-mode library. However, you can link user-mode libraries to a supervisor-mode library.

8.13 WRITING YOUR OWN VECTORS AND COMPLETION ROUTINES


You can write your own mode-switching vectors and completion routines. This may be necessary for threaded code. If you use your own vectors, build them into the task and use the `-SV` switch on the `RESSUP` or `RESLIB` option when you build the referencing task. If you create your own completion routines, write your completion routine to resemble the system-supplied completion routines (see Example 8-3) as much as possible. If you do not retain the last three lines of code as indicated in Example 8-3, then if the Executive processes an interrupt before the mode switch back to user mode has completed, your task may crash.

8.14 OVERLAID SUPERVISOR-MODE LIBRARIES

It is possible to use overlaid supervisor-mode libraries. Three restrictions must be noted when building these libraries:

- The completion routine for the library must be in the root.
- Only one level of overlay is allowed. This is illustrated in Figure 8-4.
- Although the Fast Task Builder (FTB) can link to supervisor-mode libraries, it cannot link to overlaid supervisor-mode libraries.

SUPERVISOR-MODE LIBRARIES (RSX-11M-PLUS ONLY)



ZK-1102-82

Figure 8-4 Overlay Configuration Allowed for Supervisor-Mode Libraries

CHAPTER 9

MULTIUSER TASKS

9.1 INTRODUCTION

A multiuser task is a task that shares the pure (read-only) portion of its code with two or more copies of the impure (read/write) portion of its code. When the system receives an initial run request for a multiuser task, a copy of both the read-only and read/write portions of the task are read into physical memory. As long as the task is running, all subsequent run requests for it result in the system duplicating only the read/write portion of the task in physical memory. Thus, multiuser tasks are memory efficient.

When you build a task, you designate it as multiuser by applying the /MU switch to the task image file. This switch directs the Task Builder to create two regions for the task. One region (region 0) contains the read-write portion of the task; the other region (region 1) contains the read-only portion of the task.

As with all other tasks, TKB uses a program section's access code to determine its placement within a multiuser task's image. It divides address space into read/write and read-only sections. Unlike in a single user task, however, the read-only portion of the task is hardware protected. In addition, TKB separates the read/write portions of a multiuser task from the read-only portions and places them in separate regions at opposite ends of the task's address space. It allocates the low-address APRs to the read/write portion (which includes the task's header and stack area) and the highest available APRs to the read-only portion. Figure 9-1 illustrates this allocation.

For I- and D-space multiuser tasks, in addition to having the multiuser task divided into regions of read-only PSECTs and read/write PSECTs, these regions themselves are divided into I-space areas and D-space areas. All of the following combinations must be present in an I- and D-space multiuser task:

- .PSECT psectnamew, RO, I, ...
- .PSECT psectnamex, RW, I, ...
- .PSECT psectnamey, RO, D, ...
- .PSECT psectnamez, RW, D, ...


If neither the read-only nor the read/write portion of the task contains memory-resident overlays, TKB allocates two window blocks in the header of the task. When the task is installed, the INSTALL processor will initialize these window blocks as follows:

- Window block 0 describes the range of virtual addresses (the window) for the read/write portion of the task. This region always contains the task's header.

MULTIUSER TASKS

- Window block 1 describes the range of virtual addresses for the read-only portion.

Figure 9-2 below shows the window-to-region relationship of a multiuser task.



ZK-441-81

Figure 9-1 Allocation of Program Sections in a Multiuser Task

9.1.1 Overlaid Multiuser Task

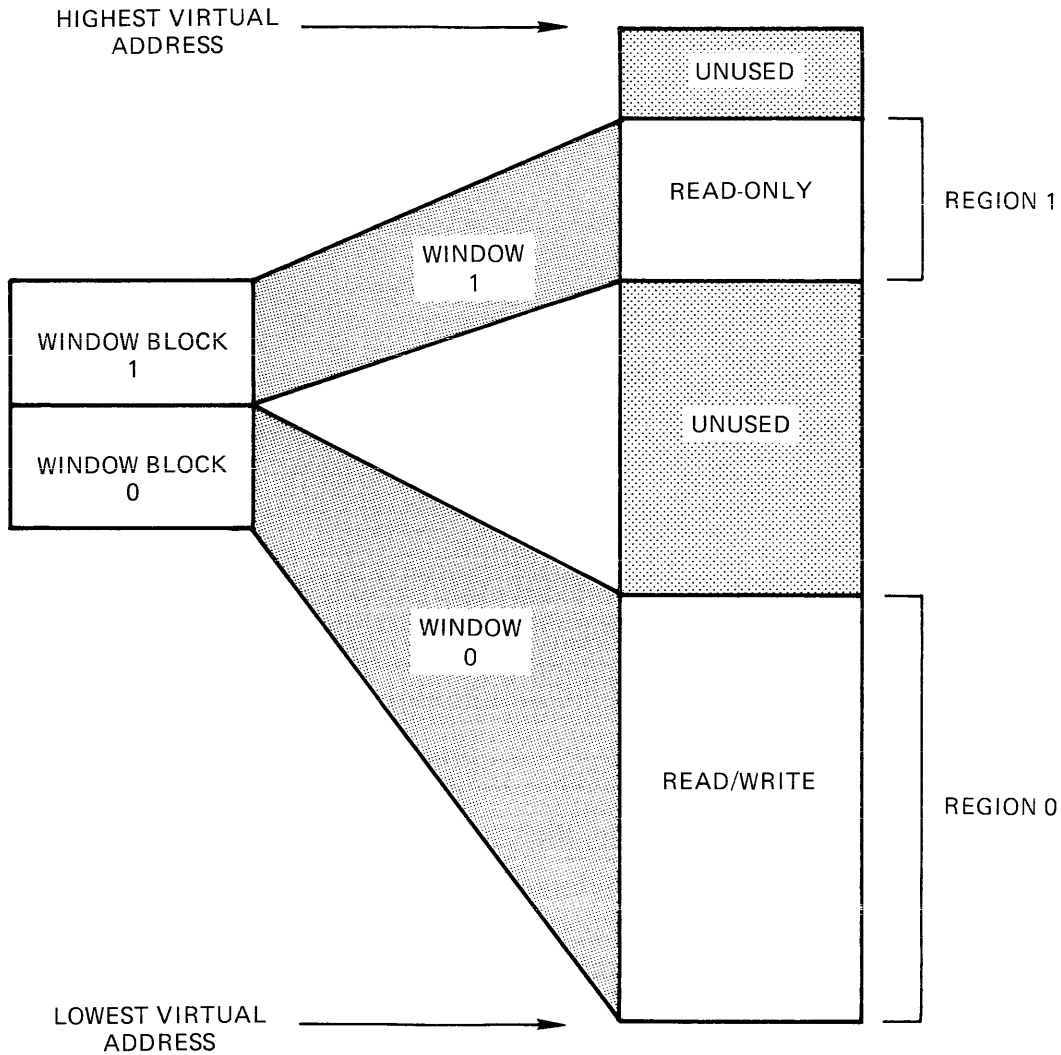
If a multiuser task is an overlaid task (described in Chapter 3), the read-only portion of the task can be made up of the following:

- The read-only program sections of the root segment
- Branches of an overlay structure if the complete branch is memory resident and read-only
- A co-tree structure if the entire co-tree is memory resident and read-only.

9.1.2 Disk Image of a Multiuser Task

The disk image of a multiuser task is somewhat different from that of a single-user task. The read-only portion of the task is placed at the end of the disk image. The relative block number of the read-only portion and the number of blocks it occupies appears in the label block. The read-only portion of the image is described in the first library descriptor of the LIBRARY REQUEST section of the label block. (Refer to Appendix B for more information on the task image data structures.)

MULTIUSER TASKS



ZK-442-81

Figure 9-2 Windows for a Multiuser Task

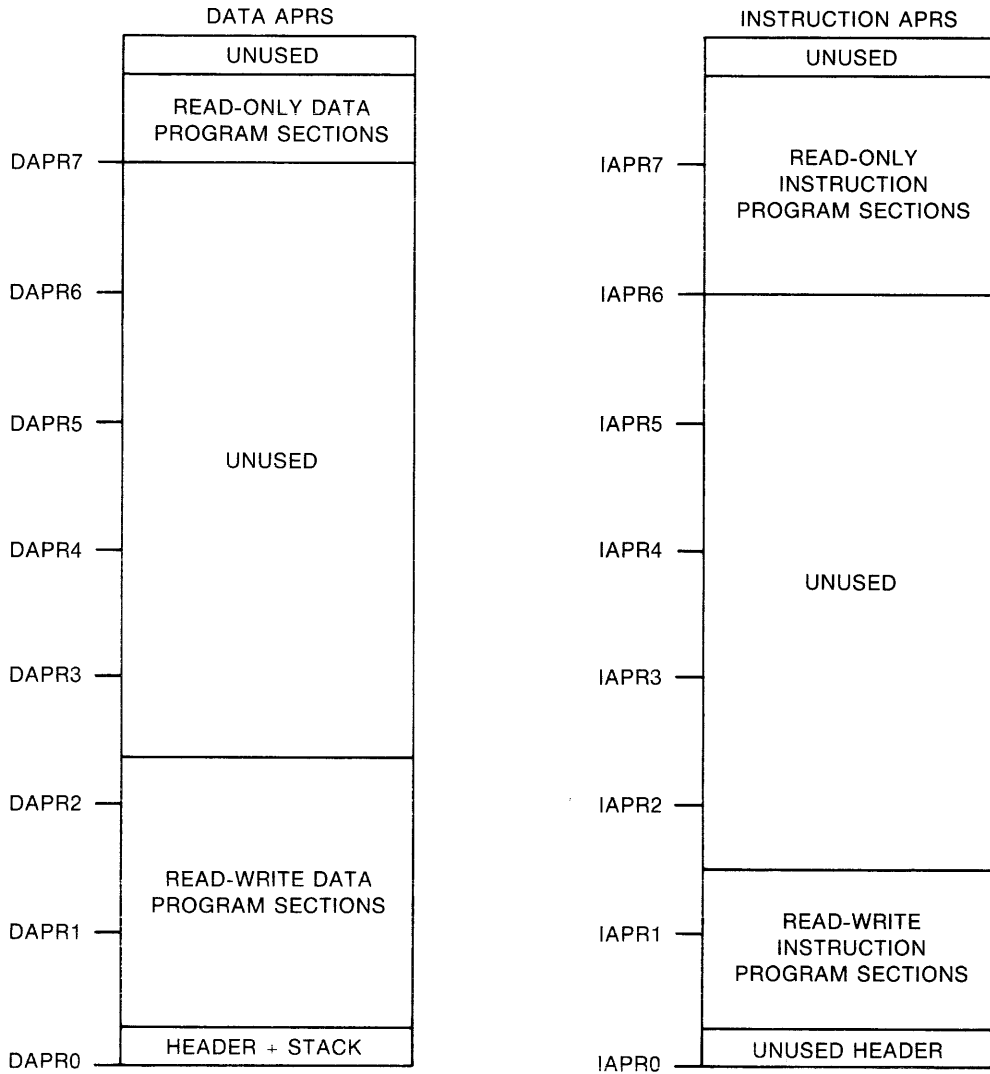
9.1.3 I- and D-Space Multiuser Tasks

The APR and window block assignment in an I- and D-space multiuser task differs from that in a conventional multiuser task.

D-space APRs map the read-write and read-only PSECTs that have the data attribute. Similarly, I-space APRs map the read-write and read-only PSECTs that have the instruction attribute. Figure 9-3 shows the APR mapping for both kinds of PSECTs in an I- and D-space multiuser task.

TKB needs four window blocks to map an I- and D-space multiuser task. Window blocks 0 and 1 map region 0, which contains the read-write instruction and data PSECTs. Window blocks 2 and 3 map region 1, which contains the read-only instruction and data PSECTs. Figure 9-4 illustrates the mapping and assignment of these window blocks for an I- and D-space multiuser task.

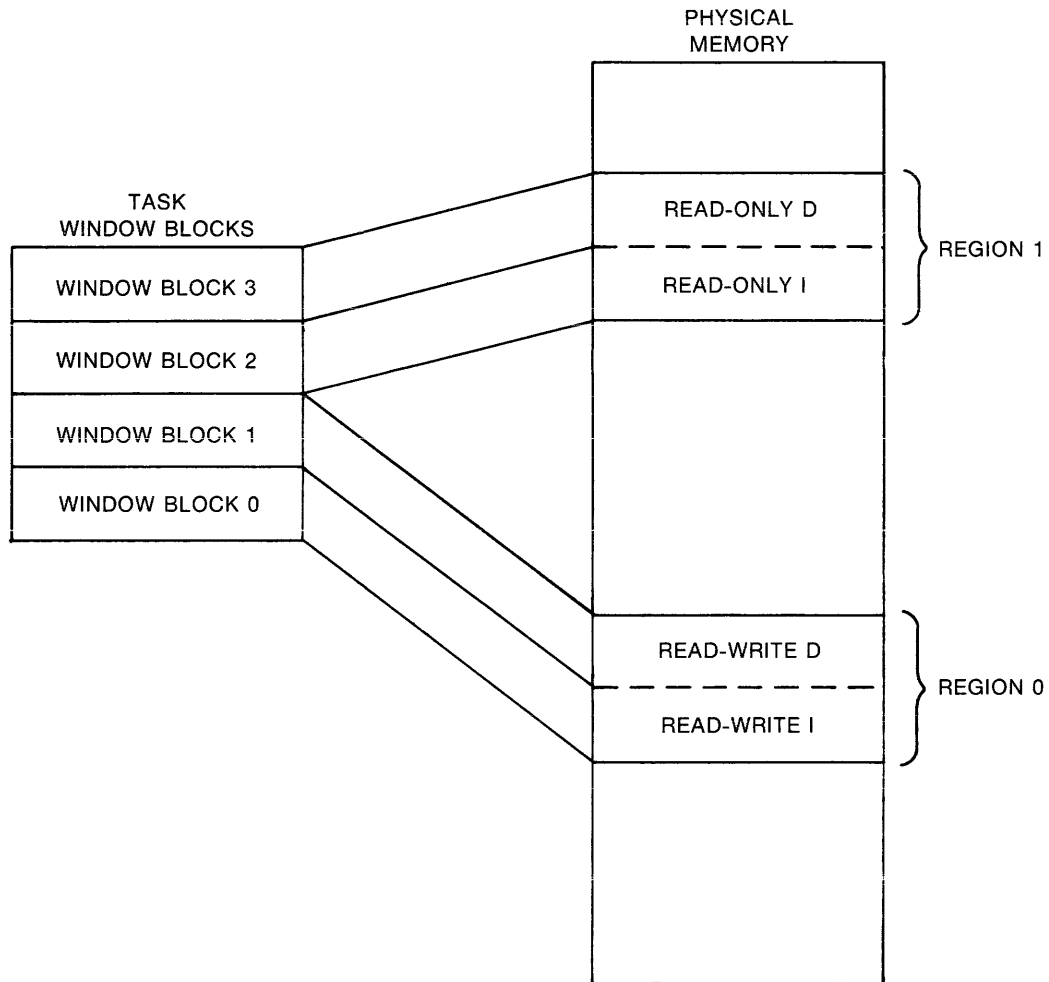
MULTIUSER TASKS



ZK-1103-82

Figure 9-3 Example Allocation of Program Sections in an I- and D-Space Multiuser Task

MULTIUSER TASKS



ZK-1104-82

Figure 9-4 Windows for an I- and D-Space Multiuser Task

9.2 EXAMPLE 9-1: BUILDING A MULTIUSER TASK

The text in this section and the figures associated with it illustrate the development of a multiuser task. This example was created by concatenating into a single file the resident library file (LIB.MAC) and the task that links to it (MAIN.MAC) from Example 5-3. It is not intended to represent a typical multiuser task application. However, it does illustrate the Task Builder's allocation of program sections in a multiuser task and that is its primary value. The concatenated source file, named ROTASK.MAC, for this example is shown in Example 9-1.

MULTIUSER TASKS

Example 9-1, Part 1 Source Listing for ROTASK.MAC

```

.TITLE ROTASK
.IDENT /01/
.MCALL QIOWSS,EXITSS

OP1:  .WORD 1 ; OPERAND 1
OP2:  .WORD 1 ; OPERAND 2
ANS:  .BLKW 1 ; RESULT

OUT:  .BLKW 100. ; FORMAT MESSAGE
FORMAT: .ASCIZ /THE ANSWER = %D,/
.EVEN

START:
MOV #ANS,-(SP) ; TO CONTAIN RESULT
MOV #OP2,-(SP) ; OPERAND 2
MOV #OP1,-(SP) ; OPERAND 1
MOV #3,-(SP) ; PASSING 3 ARGUMENTS
MOV SP,R5 ; ADDRESS OF ARGUMENT BLOCK
CALL AADD ; ADD TWO OPERANDS
CALL PRINT ; PRINT RESULTS
MOV SP,R5 ; ADDRESS OF ARGUMENT BLOCK
CALL SUBB ; SUBTRACT SUBROUTINE
CALL PRINT ; PRINT RESULTS
MOV SP,R5 ; ADDRESS OF ARGUMENT BLOCK
CALL MULL ; MULTIPLY SUBROUTINE
CALL PRINT ; PRINT RESULTS
MOV SP,R5 ; ADDRESS OF ARGUMENT BLOCK
CALL DIVV ; DIVIDE SUBROUTINE
CALL PRINT ; PRINT RESULTS
EXITSS

;+
;** PRINT - PRINT RESULT OF OPERATION.
;
PRINT: MOV #OUT,R0 ; ADDRESS OF SCRATCH AREA
MOV #FORMAT,R1 ; FORMAT SPECIFICATION
MOV #ANS,R2 ; ARGUMENT TO CONVERT
CALL $EDMSG ; FORMAT MESSAGE
QIOWSS #IO.WVB,#5,#1,,,,<#OUT,R1,#40>
RETURN ; RETURN FROM SUBROUTINE
;** FORTRAN CALLABLE SUBROUTINE TO ADD TWO INTEGERS
.PSECT AADD,RO,I,GBL,REL,CON

AADD:: CALL $SAVAL ; SAVE R0-R5
MOV @2(R5),R0 ; FIRST OPERAND
MOV @4(R5),R1 ; SECOND OPERAND
ADD R0,R1 ; SUM THEM
MOV R1,@6(R5) ; STORE RESULT
RETURN ; RESTORE REGISTERS AND RETURN

```

(continued on next page)

MULTIUSER TASKS

Example 9-1, Part 1 Source Listing for ROTASK.MAC

```

** FORTRAN CALLABLE SUBROUTINE TO SUBTRACT TWO INTEGERS
.PSECT SUBB,RO,I,GBL,REL,CON
SUBB:: CALL  $$AVAL          ; SAVE R0-R5
        MOV   @2(R5),R0      ; FIRST OPERAND
        MOV   @4(R5),R1      ; SECOND OPERAND
        SUB   R1,R0          ; SUBTRACT SECOND FROM FIRST
        MOV   R0,@6(R5)      ; STORE RESULT
        RETURN                ; RESTORE REGISTERS AND RETURN

** FORTRAN CALLABLE SUBROUTINE TO DIVIDE TWO INTEGERS
.PSECT DIVV,RO,I,GBL,REL,CON
DIVV:: CALL  $$AVAL          ; SAVE REGS R0-R5
        MOV   @2(R5),R3      ; FIRST OPERAND
        MOV   @4(R5),R1      ; SECOND OPERAND
        CLR   R2              ; LOW ORDER 16 BITS
        DIV   R1,R2          ; DIVIDE
        MOV   R2,@6(R5)      ; STORE RESULT
        RETURN                ; RESTORE REGISTERS AND RETURN

** FORTRAN CALLABLE SUBROUTINE TO MULTIPLY TWO INTEGERS
.PSECT MULL,RO,I,GBL,REL,CON
MULL:: CALL  $$AVAL          ; SAVE R0-R5
        MOV   @2(R5),R0      ; FIRST OPERAND
        MOV   @4(R5),R1      ; SECOND OPERAND
        MUL   R0,R1          ; MULTIPLY
        MOV   R1,@6(R5)      ; STORE RESULT
        RETURN                ; RESTORE REGISTERS AND RETURN
.END      START

```

Once you have assembled ROTASK, you can build it with the following command sequence:

```

TKB>ROTASK/MU,ROTASK/-WI/-SP=ROTASK
TKB>/
Enter Options:
TKB>ROPAR=RDONLY
TKB>//

```

This command sequence directs TKB to build a multiuser (/MU) task image named ROTASK.TSK and to create an 80-column (/WI) map file named ROTASK.MAP. Because /-SP is attached to the map file, TKB does not output a map to the line printer.

The ROPAR option specifies that the system is to load the read-only portion of the task into a partition named RDONLY. Specifying a separate partition for the task's read-only region is not a system requirement. The system will load the read/write portion into partition GEN. The system will not load either region until it receives a run request for the task.

The map that results from this command sequence is shown in Example 9-1, Part 2. Note that TKB has added one field to the task attributes section of this map describing the disk block limits of the read-only portion of the task. It has also added a field to the root segment portion of the map that describes the memory limits of the read-only portion of the task.

Finally, note that TKB has allocated space for all the program sections with the read-only attribute, beginning with the highest available APR (in this case, APR 7).

MULTIUSER TASKS

Example 9-1, Part 2 Task Builder Map for ROTASK.TSK

ROTASK.TSK;1 Memory allocation map TKB M40.10 PAGE 1
 10-DEC-82 14:42

Partition name : GEN
 Identification : 01
 Task UIC : [7,62]
 Stack Limits: 000274 001273 001000 00512.
 PRG xfr address: 001634
 Task attributes: MU
 Total address windows: 2.
 Task image size : 1088. words
 Task address limits: 000000 004157
 R-W disk blk limits: 000002 000006 000005 00005.
 R-O disk blk limits: 000007 000007 000001 00001.

*** Root segment: ROTASK

R/W mem limits: 000000 004157 004160 02160.
 R-O mem limits: 160000 160377 000400 00256.
 Disk blk limits: 000002 000006 000005 00005.

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.: (RW,I,LCL,REL,CON) 001274 002662 01458.			
	001274 000530 00344.	ROTASK 01	ROTASK.OBJ;1
AADD : (RO,I,LCL,REL,CON) 160000 000024 00020.			
	160000 000024 00020.	ROTASK 01	ROTASK.OBJ;1
DDIV : (RO,I,LCL,REL,CON) 160024 000026 00022.			
	160024 000026 00022.	ROTASK 01	ROTASK.OBJ;1
MMUL : (RO,I,LCL,REL,CON) 160052 000024 00020.			
	160052 000024 00020.	ROTASK 01	ROTASK.OBJ;1
SSUB : (RO,I,LCL,REL,CON) 160076 000024 00020.			
	160076 000024 00020.	ROTASK 01	ROTASK.OBJ;1
\$\$RESL: (RO,I,LCL,REL,CON) 160122 000212 00138.			

Global symbols:

AADD 160000-R DIVV 160024-R MULL 160052-R SUBB 160076-R

*** Task builder statistics:

Total work file references: 2145.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 7086. words (27. PAGES)
 Size of work file: 1024. words (4. PAGES)

Elapsed time:00:00:07

CHAPTER 10

SWITCHES

You use switches and options to control the construction of your task image. This chapter provides detailed reference information on all the TKB switches. Chapter 11 describes the TKB options.

10.1 SWITCHES

The syntax for a file specification, as given in Chapter 1, is:

```
dev:[group,member]filename.type;version/sw1/sw2.../swn
```

Optionally, you can conclude a file specification with one or more switches (sw1,sw2,...swn). When you do not specify a switch, the Task Builder establishes a default setting for it.

You designate a switch by a 2- to 4-character code preceded by a slash (/). If you precede the 2- to 4-character code with a minus sign (-) or the letters NO, TKB negates the function of the two characters. For example, TKB recognizes the following settings for the switch CP (checkpointable):

```
/CP      The task is checkpointable.  
/-CP     The task is not checkpointable.  
/NOCP    The task is not checkpointable.
```

In some cases, two particular switches cannot both be used in a file specification. When such a conflict occurs, TKB selects the overriding switch according to the following table:

Switch	Switch	Overriding Switch
/AC (Ancillary Control Processor)	/PR (Privileged)	/AC
/EA (Extended Arithmetic Element)	/FP (Floating Point Processor)	/FP
/CC (Concatenated object file)	/LB (Library file)	/LB

For example:

```
MCR>TKB IMG5=IN6,IN5/LB/CC
```

TKB assumes that the input file IN5 is a library file. It searches the file for undefined global references. It does not include in the task image all of the modules in IN5.

SWITCHES

The switches that TKB recognizes are given in alphabetical order in Table 10-1. Sections 10.1.1 through 10.1.30 give detailed descriptions of each switch, in alphabetical order, including:

- The switch format
- The file(s) to which the switch can be applied
- A description of the effect of the switch on the Task Builder
- The default assumption made if the switch is not present

Table 10-1
Task Builder Switches

Format	Meaning	Applies to File	Default
/AC[:n]	Task is an ancillary control processor.	.TSK	/-AC
/AL	Task can be checkpointed to space allocated in the task image file.	.TSK	/-AL
/CC	Input file consists of concatenated object modules.	.OBJ	/CC
/CM	Memory-resident overlays are aligned on 256-word physical boundaries.	.TSK	/-CM
/CO	Causes TKB to build a shared common.	.TSK .STB	/CO
/CP	Task is checkpointable.	.TSK	/-CP
/CR	A global cross-reference listing is appended to the memory allocation file.	.MAP	/-CR
/DA	Task contains a debugging aid.	.TSK, .OBJ	/-DA
/DL	Specified library file is a replacement for the system object module library.	.OLB	/-DL
/EA	Task uses extended arithmetic element.	.TSK	/-EA
/EL	Specifies library size according to partition size.	.TSK	/-EL
/FP	Task uses the Floating Point Processor.	.TSK	/-FP on RSX-11M /FP on RSX-11M-PLUS

(continued on next page)

SWITCHES

Table 10-1 (Cont.)
Task Builder Switches

Format	Meaning	Applies to File	Default
/FU	All co-tree overlay segments are searched for matching definition or reference when modules from the default object module library are being processed.	.TSK	/-FU
/HD	Task image includes a header.	.TSK, .STB	/HD
/ID	Task will use I- and D-space.	.TSK	/-ID
/IP	Allows TKB to inform INS that the task purposely overmaps the I/O page.	.TSK	/-IP
/LB	Input file is a library file.	.OLB	/-LB
/LI	Informs TKB to build a shared library.	.TSK .STB	/-LI
/MA	Map file includes information from the file.	.MAP, .OBJ	/MA or /-MA ¹
/MM	System on which the task is to run has memory management.	.TSK	/MM or /-MM ²
/MP	Input file contains an overlay description.	.ODL	/-MP
/MU	Task is a multiuser task.	.TSK	/-MU
/NM	Tells TKB to inhibit two diagnostic messages.	.TSK	/-NM
/PI	Task is position independent.	.TSK, .STB	/-PI
/PM	Postmortem Dump is requested.	.TSK	/-PM
/PR[:n]	Task has privileged access rights.	.TSK	/-PR
/RO	Memory-resident overlay operator (!) is enabled.	.TSK	/RO
/SE	Messages can be directed to the task by means of the Executive SEND directive.	.TSK	/SE

1. The default is /MA for an input file, and /-MA for system library and resident library .STB files.

2. The default for the memory management switch is /MM if the host system has memory management hardware, and /-MM if the host system does not have memory management hardware.

(continued on next page)

SWITCHES

Table 10-1 (Cont.)
Task Builder Switches

Format	Meaning	Applies to File	Default
/SG	Allocates task program sections alphabetically by access code (RW followed by R0).	.TSK	/-SG
/SH	Short memory allocation file is requested.	.MAP	/SH
/SL	Task is slaved to an initiating task.	.TSK	/-SL
/SP	Spool map output.	.MAP	/SP
/SQ	Allocates task program sections in input order by access code.	.TSK	/-SQ
/SS	Selective search for global symbols.	.OBJ	/-SS
/TR	Task is to be traced.	.TSK	/-TR
/WI	Memory allocation file is printed at a width of 132 characters.	.MAP	/WI
/XH	RSX-11M-PLUS only switch. Task is to have an external header.	.TSK	3
/XT[:n]	TKB exits after n diagnostic.	.TSK	/-XT

3. The default is ultimately determined by the /XHR switch in the INSTALL command, which overrides the TKB setting except for /-XH.

10.1.1 /AC[:n] -- Ancillary Control Processor

File

Task image

Syntax

file.TSK/AC:0=file.OBJ

or

file.TSK/AC:4=file.OBJ

or

file.TSK/AC:5=file.OBJ

Description

The /AC switch informs TKB that your task is an ancillary control processor; that is, it is a privileged task that extends certain Executive functions. For example, the system task FllACP is an ancillary control processor that receives and processes FILES-11 related input and output requests on behalf of the Executive.

Effect

This switch also informs TKB that your task is privileged. TKB sets the AC attribute flag and the privileged attribute flag in your task's label block flag word.

The value of n is an octal number that specifies the first KT-11 Active Page Register (APR) that you want the Executive to use to map your task's image when your task is running in user mode. Legal values are 0, 4, and 5. If you do not specify n, the Task Builder assumes a value of 5.

If you do not explicitly specify that your task is to run on a mapped system (through the /MM switch) and it is not otherwise implied (TKB is not running in a system with KT-11 hardware), TKB merely tests the value of n for validity, but otherwise ignores it.

Default

/-AC

NOTE

You should not use /AC and /PR on the same command line.

SWITCHES

AL

10.1.2 /AL -- Allocate Checkpoint Space

File

Task image

Syntax

file.TSK/AL=file.0BJ

Description

The /AL switch informs TKB that your task is checkpointable. The system will checkpoint it to a space in your task's image file. However, the system uses the system checkpoint file first if you specified dynamic checkpointing.

Effect

As well as making your task checkpointable, this switch directs TKB to allocate additional space in your task image file to contain the checkpointed task image.

Default

/-AL

NOTES

Do not use /CP in the same command line in which you use /AL.

Also, the /AL switch should not be used with the /-HD switch to build tasks. Examples of tasks that use the /-HD switch are: the Executive, device drivers, and commons.

10.1.3 /CC -- Concatenated Object Modules

File

Input

Syntax

file.TSK=file.OBJ/-CC

Description

/CC controls the way TKB extracts modules from your input file.

Effect

By default, TKB includes in your task's image all the modules of your input file. If you negate this switch (as in the Syntax section above), TKB includes only the first module of your input file.

Default

/CC

CM

10.1.4 /CM -- Compatibility Mode Overlay Structure

File

Task image

Syntax

file.TSK/CM=file.OBJ

Description

/CM causes the Task Builder to build your task in compatibility mode.

Effect

TKB aligns memory-resident overlay segments on 256-word boundaries for compatibility with other implementations of the mapping directives.

Default

/-CM

10.1.5 /CO -- Build a Common Block Shared Region

File

Task image
 .STB file

Syntax

```
file.TSK/CO=file.OBJ
    or
,,file.STB/CO=file.OBJ
```

Description

The /CO switch informs TKB that a shared common is being built. If you build a shared common, you should use the /CO switch and the /-HD switch.

If you use the /-PI switch for an absolute shared common, all the program sections in the common are marked absolute. Using the /-PI/-HD switches without the /CO switch causes TKB to build a shared library.

If you use the /PI switch for a relocatable shared common, all program sections in the common are marked relocatable.

In either case, the .STB file contains all the program section names, attributes, length, and symbols. TKB links common blocks by means of program sections. Therefore, the .STB file of a shared region built with the /CO switch contains all defined program sections.

Using the /PI/-HD switches without the /CO switch causes TKB to build a shared common.

The /CO switch does not have a /-CO form.

Effect

This switch causes TKB to include all program section declarations in the .STB file.

Defaults

/CO

CP

10.1.6 /CP -- Checkpointable

File

Task image

Syntax

file.TSK/CP=file.OBJ

Description

/CP causes TKB to mark your task as checkpointable. The system will checkpoint it to space that you have allocated in the system checkpoint file on the system disk. This switch assumes that you have allocated the checkpoint space through the MCR command ACS. (Refer to the RSX-11M/M-PLUS MCR Operations Manual.)

Effect

The system writes your task to the system checkpoint file on secondary storage when its physical memory is required by a task of higher priority.

Default

/-CP

NOTE

Using /AL also makes your task checkpointable.

10.1.7 /CR -- Cross-Reference

File

Memory allocation (map)

Syntax

file.TSK,file.MAP/CR=file.OBJ

Description

The /CR switch directs TKB to add a cross-reference listing to the map file of your task.

Effect

TKB creates a special work file (file.CRF) that contains segment, module, and global symbol information. The Task Builder then calls the Cross-Reference Processor (CRF) to process the file. CRF creates a cross-reference listing from the information contained in the file, and then deletes file.CRF. (Refer to the RSX-11 Utilities Manual for more information on CRF.)

The Example section below describes the cross-reference listing and its contents.

NOTE

For this switch to be effective, CRF must be installed in your system.

Default

/-CR

Example

Example 10-1 shows a cross-reference listing for task OVR. The numbered items in the notes correspond to the numbers in Example 10-1.

CR (Cont.)

Example 10-1 Cross-Reference Listing for OVR.TSK

```

CREF          CREATED BY  TKB      ON 27-JUL-82 AT 09:46      PAGE 1
GLOBAL CROSS REFERENCE                                CREF  V01
SYMBOL  VALUE      REFERENCES...
AADD    020000-R  * AADD    @  CALC
ADDEXI  020060-R  * AADD
ARGBLK  001340-R  CALC    #  MAIN
BUFF    001366-R  #  MAIN    OUTPUT
CALC    003270-R  *  CALC    @  MAIN
DIFR    001360-R  CALC    #  MAIN
DIVEXI  020062-R  *  DIVV
DIVR    001364-R  CALC    #  MAIN
DIVV    020000-R  @  CALC    *  DIVV
I       001350-R  INPUT   #  MAIN
IE.EOF  177766    INPUT   #  QIOSYM
INITL   005664-R  #  INITL  ^  MAIN
INPUT   003364-R  *  INPUT  @  MAIN
IOSB    001334-R  INPUT   #  MAIN

```

```

CREF          CREATED BY  TKB      ON 27-JUL-82 AT 09:46      PAGE 2
SEGMENT CROSS REFERENCE                                CREF  V01
SEGMENT NAME  RESIDENT MODULES
AADD          AADD
CALC          CALC
DIVV          DIVV
INPUT        ARITH  CATB  INPUT  QIOSYM  SAVRG
LIBROT       INITL  SAVAL
MAIN         ALERR  AUTO  MAIN   OVCTR  OVDAT  OVRES  SAVR1
              VCTDF
MULL         MULL
OUTPUT       ARITH  CATB  CBTA   CDDMG  C5TA   DARITH  EDDAT
              EDTMG  OUTPUT QIOSYM SAVRG
SUBB         SUBB

```

NOTES

- ① The cross-reference page header gives the name of the memory allocation file, the originating task (TKB), the date and time the memory allocation file was created, and the cross-reference page number.
- ② The cross-reference list contains an alphabetic listing of each global symbol along with its value and the name of each referencing module. When a symbol is defined in several segments within an overlay structure, the last defined value is printed. Similarly, if a module is loaded in several segments within the structure, the module name is displayed more than once within each entry.

CR (Cont.)

The suffix -R appears next to the value if the symbol is relocatable.

Prefix symbols accompanying each module name define the type of reference as follows:

Prefix Symbol	Reference Type
blank	Module contains a reference that is resolved in the same segment or in a segment toward the root.
^	Module contains a reference that is resolved directly in a segment away from the root or in a co-tree.
@	Module contains a reference that is resolved through an autoloading vector.
#	Module contains a nonautoloadable definition.
*	Module contains an autoloadable definition.

- ③ The segment cross-reference lists the name of each overlay segment and the modules that compose it. If the task is a single-segment task, this section does not appear.

DA

10.1.8 /DA -- Debugging Aid

File

Task image or input

Syntax

file.TSK/DA=file.OBJ

or

file.TSK=file.OBJ,file.OBJ/DA

Description

/DA causes TKB to include a debugging aid in your task. The debugging aid controls the task's execution.

Effect

If you apply this switch to your task image file, TKB includes the system debugging aid LB0:[1,1]ODT.OBJ into your task image. If you use the /DA switch with the /ID switch, TKB includes LB:[1,1]ODTID.OBJ in the task.

TKB passes control to the debugging program when you or the system starts task execution.

If you apply this switch to one of your input files, TKB assumes that the file is a debugging aid that you have written. Such debugging programs can trace a task, printing out relevant debugging information, or monitor the task's performance for analysis. The default file type for the debugging aid is .OBJ.

In either case, /DA has the following effects on your task image:

- The transfer address of the debugging aid overrides the task transfer address.
- TKB initializes the header of your task so that, on initial task load, registers R0 through R4 contain the following values:

R0 - Transfer address of task.

R1 - Task name in Radix-50 format (word #1).

R2 - Task name (word #2).

R3 - The first three of six RAD50 characters representing the version of your task. TKB derives the version from the first .IDENT directive it encounters in your task. If no .IDENT directive is in your task, this value is 01.

R4 - The second three RAD50 characters representing the version of your task.

Default

/-DA

10.1.9 /DL -- Default Library

File

Input

Syntax

file.TSK=file.OBJ,file.OLB/DL

Description

This switch causes the input file to be a replacement for the system object module library. The default file type for the input file is .OBJ.

Effect

The library file you have specified replaces the file LBO:[1,1]SYSLIB.OLB as the library file that the Task Builder searches to resolve undefined global references. The default device for the replacement file is SY0:. TKB refers to it only when undefined symbols remain after it has processed all the files you have specified. You can apply the /DL switch to only one input file.

Default

/-DL

EA

10.1.10 /EA -- Extended Arithmetic Element

File

Task image

Syntax

file.TSK/EA=file.OBJ

Description

/EA informs TKB that your task uses the Kell-A Extended Arithmetic Element.

Effect

TKB allocates three words in your task's header for saving the state of the extended arithmetic element.

Default

/-EA

NOTE

You should not use /EA and /FP on the same command line.

SWITCHES

EL

10.1.11 /EL -- Extend Library

File

Task image

Syntax

file.TSK/LI/-HD/EL=file.OBJ

Description

/EL places the upper address limit as determined by the PAR option in the library's label block, though the actual size of the library may be smaller. This switch is useful when you build vectored libraries such as RMS, which are subject to size changes.

Effect

This switch specifies the maximum possible size for the library according to the size specified in the PAR option. The switch specifies a larger library virtual address range than is actually present in the library to allow RMS to map its vectored library segments.

Default

/-EL

FP

10.1.12 /FP -- Floating Point

File

Task image

Syntax

file.TSK/FP=file.OBJ

Description

/FP informs TKB that your task uses the Floating Point Processor.

Effect

TKB allocates 25 words in your task's header for saving the state of the Floating Point Processor.

Default

~~/-FP on RSX-11M systems~~
~~/FP on RSX-11M-PLUS systems~~

NOTES

1. You should not use /FP and /EA on the same command line.
2. The /FP switch allocates space in the task header to save the floating point status if your task is context switched. Therefore, in an RSX-11M system, if a task that uses the Floating Point Processor is built without the /FP switch, the task will run correctly until a second task that uses the Floating Point Processor is run. Then both tasks will either crash or produce incorrect results. For information on changing the Task Builder's defaults, refer to Appendix F.

10.1.13 /FU -- Full Search

File

Task image

Syntax

file.TSK/FU=file.ODL/MP

Description

This switch controls the Task Builder's search for undefined symbols when it is processing modules from the default library.

Effect

When TKB processes modules from the default object module library, and it encounters undefined symbols within those modules, it normally limits its search for definitions to the root of the main tree and to the current tree. Thus, unintended global references between co-tree overlay segments are eliminated. When the /FU switch is appended to the task image file of an overlaid task, TKB searches all co-tree segments for a matching definition or reference. See Sections 3.2.2 and 3.2.3 in Chapter 3 for more details.

Default

/-FU

HD

10.1.14 /HD -- Header

File

Task image or symbol definition

Syntax

file.TSK/-HD,,file.STB=file.OBJ

or

file.TSK,,file.STB/-HD=file.OBJ

Description

The /-HD form of this switch directs TKB to exclude a header from your task image.

Effect

TKB does not construct a header in your task image. You use the negated form of this switch when you are building commons, resident libraries, and loadable drivers.

Default

/HD

10.1.15 /ID -- I- and D-space Task (RSX-11M-PLUS Only)

File

Task image

Syntax

file.TSK/ID=file.OBJ

Description

This switch directs TKB to mark your task as one that uses I-space APRs and D-space APRs in user mode. TKB separates I-PSECTS from D-PSECTS.

Effect

TKB includes the data structures in the task label block that informs INSTALL that the task has separate I-space and D-space.

Default

/-ID

IP

10.1.16 /IP -- Task Maps I/O Page

File

Task image

Syntax

file.TSK/PR/-IP=file.OBJ

Description

You use the /-IP switch to inform TKB that the task is purposely over 12K and does not need to be mapped to the I/O page.

Effect

TKB sets a bit in the task's label block that informs INSTALL (INS) that the task intentionally does not map the I/O page. When this bit is set, INS does not display an error message when it detects that the privileged task extends into APR 7.

Default

/IP

10.1.17 /LB -- Library File

File

Input

Syntax

file.TSK=file.OBJ,file.OLB/LB

or

file.TSK=file.OBJ,file.OLB/LB:mod-1:mod-2...:mod-8

Description

The file to which this switch is attached is an object module library file. The Task Builder's interpretation of this switch depends upon which of the following forms you use:

- Without arguments (the first syntax given above)
- With arguments (the second syntax given above)

The default file type is .OLB.

Effect

If you apply this switch without arguments, TKB assumes that your input file is a library file of relocatable object modules. TKB searches the file immediately to resolve undefined references in any modules preceding the library specification. It also extracts from the library, for inclusion in the task image, any modules that contain definitions for such references.

If you apply the switch with arguments, TKB extracts from the library the modules named as arguments of the switch regardless of whether the modules contain definitions for unresolved references.

If you want TKB to search an object module library file both to resolve global references and to select named modules for inclusion in your task image, you must name the library file twice: once, with the modules you want included in your task image listed as arguments of the /LB switch; and a second time, with the /LB switch and no arguments. For example:

```
file.TSK=file.OLB/LB:mod-1:mod-2,file.OLB/LB
```

SWITCHES

LB (Cont.)

The position of the library file within TKB command sequence is important. The following rules apply:

- The library file must follow to the right of the input file(s) that contain references to be defined in the library. For example:

```
TKB>file.TSK=infile1.OBJ,lib.OLB/LB
```

The command above illustrates the correct usage of the /LB switch; the following command illustrates incorrect usage:

```
TKB>file.TSK=lib.OLB/LB,file1.OBJ
```

- If you are using the Task Builder's multiline input, and you specify a given library more than once during the command sequence, you must attach the /LB switch to the library file each time you specify the library. For example:

```
>TKB
TKB>file.TSK=file1.OBJ,file2.OBJ,lib.OLB/LB
TKB>file3.OBJ,file4.OBJ,lib.OLB/LB
//
```

- When you are building an overlay structure, you must specify object module libraries for an overlay structure within the Overlay Description Language (ODL) file for the structure. To do this, you must use the .FCTR directive to specify the library. For example:

```
                .ROOT CNTRL-LIB-(AFCTR,BFCTR,C)
AFCTR:          .FCTR A0-LIB-(A1,A2-(A21,A22))
BFCTR:          .FCTR B0-LIB-(B1,B2)
LIB:            .FCTR LB:[303,3]LIBOBJ.OLB/LB
                .END
```

The technique used in the ODL file above allows you to control the placement of object module library routines into the segments of your overlay structure. (For more information on overlaid tasks, see Chapter 3.)

NOTES

1. You should not use the /LB switch and the /CC switch in the same command sequence.
2. You can use the /SS switch in conjunction with the /LB switch (with or without arguments) to perform a selective search for global definitions.

Default

/-LB

10.1.18 /LI -- Build a Library Shared Region

File

Task image
.STB file

Syntax

file.TSK/LI=file.OBJ

or

.,file.STB/LI=file.OBJ

Description

The /LI switch makes TKB build a shared library. However, you must use the /-HD switch with the /LI switch to build the shared library. The /LI switch does not have a /-LI form.

Effect

TKB includes only one program section declaration in the .STB file.

If you use the /-PI switch for an absolute library, TKB names the program section .ABS, makes the library position dependent, and defines all symbols as absolute. Also, if you use the /-PI switch without the /LI switch, TKB assumes /LI to be the default.

If you use the /PI switch for a relocatable library, TKB names the program section the same as the root segment of the library. TKB forces this name to be the first and only declared program section in the library. TKB declares all global symbols in the .STB file relative to that program section. Also, if you use the /PI switch without the /LI switch, TKB assumes that a shared common is to be built (/CO is the default).

Default

/-LI

MA

10.1.19 /MA -- Map Contents of File

File

Input or memory allocation

Syntax

file.TSK,file.MAP=file.OBJ,file.OBJ/-MA

or

file.TSK,file.MAP/MA=file.OBJ

Description

TKB is to include information from your input file in the memory allocation output file.

Effect

If you negate this switch and apply it to an input file, TKB excludes from the map and cross-reference listings all global symbols defined or referred to in the file. In addition, TKB does not list the file in the "file contents" section of the map.

If you apply this switch to the map file, TKB includes in the map file the names of routines it has added to your task from SYSLIB. It also includes in the map file information contained in the symbol definition file of any shared region to which the task refers.

Default

/MA for input files

/-MA for system library and resident library STB files

10.1.1.20 /MM[:n] -- Memory Management

File

Task image

Syntax

```
file.TSK/MM[:n]=file.OBJ
```

or

```
file.TSK/--MM[:n]=file.OBJ
```

Description

The /MM switch informs TKB whether the system on which your task is to run has memory management hardware. Specify n as the decimal numbers 28 or 30.

Effect

If you use n with the /-MM switch (for an unmapped system), n specifies the highest physical address in K-words of the task or system being built. If you do not specify n with /-MM, the default highest address of the task or system is 28K.

If you specify n with /MM, n is ignored.

Default

When you do not apply /MM or /-MM to your task image file, TKB allocates memory according to the mapping status of the system on which your task is being built. The maximum task size for a mapped system is always 32K. The default highest address for a task or system in an unmapped system is 28K.

NOTE

When you use /-MM, TKB does not recognize the memory-resident overlay operator(!). TKB checks the operator for correct syntax, but it does not create any resident overlay segments.

MP

10.1.21 /MP -- Overlay Description

File

Input

Syntax

```
file.TSK=file.ODL/MP
```

Description

The /MP switch specifies that the input file is an Overlay Description Language (ODL) file.

Effect

TKB receives all the input file specifications from this file. It allocates virtual address space as directed by the overlay description. If you use the Task Builder's multiline command format (see Section 1.3), TKB requests option information at the console terminal by displaying:

```
ENTER OPTIONS:.
```

NOTES

1. If you use the multiline command format when you specify an ODL file, TKB automatically prompts for option input. Therefore, you must not use the single slash (/) to direct TKB to switch to option input mode when you have specified /MP on your input file.
2. When you specify /MP on the input file for your task, it must be the only input file that you specify. The default file type is .ODL.

Default

```
/-MP
```

10.1.22 /MU -- Multiuser (RSX-11M-PLUS Only)**File**

Input

Syntax

file.TSK/MU=file.OBJ

Description

The /MU switch specifies to TKB that the task is a multiuser task.

Effect

TKB separates your task's read-only and read/write program sections. It then places the read-only program sections in your task's upper virtual address space and the read/write program sections in your task's lower virtual address space.

Default

/-MU

NM

10.1.23 /NM -- No Diagnostic Messages

File

Task image

Syntax

file.TSK/NM=file.OBJ

Description

The /NM switch controls the printing of diagnostic messages.

Effect

This switch eliminates two messages:

n Undefined symbols segment seg-name

and

Module module-name multiply defines P-section p-sect-name

Default

/-NM

10.1.24 /PI -- Position Independent**File**

Task image or symbol definition

Syntax

file.TSK/PI=file.OBJ

or

file.TSK,,file.STB/PI=file.OBJ

Description

/PI informs TKB that the task's shared region contains only position-independent code or data. Use this switch with /-HD and either /CO or /LI.

Effect

TKB sets the position-independent code (PIC) attribute flag in the label block flag word of the shared region.

Be aware that if you specify /PI without using the /CO or /LI switches, TKB builds a shared common (/CO default). Also, if you specify /-PI without using the /CO or /LI switch, TKB builds a shared library (/LI default).

Default

/-PI

PM

10.1.25 /PM -- Postmortem Dump

File

Task image

Syntax

file.TSK/PM=file.OBJ

Description

If you use /PM and your task terminates abnormally, the system automatically lists the contents of the memory image.

Effect

TKB sets the Postmortem Dump flag in your task's label flag word.

NOTES

1. If your task issues an ABRT\$ (abort task) directive, the system will not dump the task image even though TKB has set the Postmortem Dump flag in your task's label flag word. In this case, the system assumes that a Postmortem Dump is not necessary since you know why your task was aborted.
2. The PMD utility must be installed in your system and be able to get into physical memory for this switch to be effective.

Default

/-PM

10.1.1.26 /PR[:n] -- Privileged

File

Task image

Syntax

file.TSK/PR:0=file.OBJ

or

file.TSK/PR:4=file.OBJ

or

file.TSK/PR:5=file.OBJ

Description

The /PR switch informs TKB that your task is privileged with respect to memory and device access rights. If you specify PR:0, your task does not have access to the I/O page or the Executive. However, if you specify PR:4 or PR:5, your task does have access to the I/O page and the Executive, in addition to its own partition.

Effect

TKB sets the Privileged Attribute flag in your task's label block flag word.

The value of n is an octal number that specifies the first Active Page register that you want the Executive to use to map your task image when your task is running in user mode. Legal values are 0, 4, and 5. If you do not specify one of these values, TKB assumes a value of 5.

If you do not explicitly specify that your task is to run on a mapped system, (through the /MM switch) and it is not implied (by the presence of KT-11 hardware on the system upon which TKB is running), TKB merely tests the value (:n) of the switch for validity; otherwise, TKB ignores it. Privileged tasks are described in Chapter 9.

Default

/-PR

NOTE

You should not use /PR and /AC on the same command line.

RO

10.1.27 /RO -- Resident Overlay

File

Task image

Syntax

file.TSK/-RO=file.ODL/MP

Description

The Task Builder's recognition of the memory-resident overlay operator (!) is enabled.

Effect

The memory-resident overlay operator (!), when present in the overlay description file, indicates to TKB that it is to construct a task image that contains one or more memory-resident overlay segments. If you negate this switch (as in the Syntax section above), TKB checks the operator for correct syntactical usage, but otherwise ignores it. With the memory-resident overlay operator thus disabled, TKB builds a disk-resident overlay from the overlay description file.

Default

/RO

10.1.28 /SE -- Send

File

Task image

Syntax

file.TSK/-SE=file.OBJ

Description

This switch determines whether messages can be directed to your task by means of the Executive Send directive. (Refer to the RSX-11M/M-PLUS Executive Reference Manual for information on the Send directive)

Effect

By default, messages can be directed to your task by means of the Executive Send directive. If you negate this switch (as in the Syntax section above), the system inhibits the queuing of messages to your task.

Default

/SE

SG

10.1.29 /SG -- Segregate Program Sections

File

Task image

Syntax

file.TSK/SG=file.OBJ

Description

The /SG switch allocates virtual address space to all (RW) program sections and then to all read-only (RO) program sections.

Effect

The /SG switch gives you control over the ordering of program sections. By using the /SG switch, you cause TKB to order program sections alphabetically by name within access code (RW followed by RO). If you specify the /SQ switch with the /SG switch, TKB orders program sections in their input order by access code. See the description of the /SQ switch.

You use the negated switch, /-SG, to make TKB interleave the RW and RO program sections. Thus, the combination /-SG/SQ results in a task with its program sections allocated in input order and its RW and RO sections interleaved. Additionally, you can use /-SQ/-SG to make TKB order program sections alphabetically with RW and RO sections interleaved. However, /-SG is the default.

When taskbuilding multiuser tasks, the /MU switch causes TKB to default to /SG. Therefore, to correctly build read-only tasks, you can use the /MU switch only.

Default

/-SG

SWITCHES

SH

10.1.30 /SH -- Short Map

File

Memory allocation (map)

Syntax

file.TSK,file.MAP/SH=file.OBJ

Description

If you specify /SH, TKB produces the short version of the memory allocation file.

Effect

TKB does not produce the "file contents" section of the memory allocation file.

Default

/SH

Example

The memory allocation file consists of the following items:

- Page header
- Task attributes section
- Overlay description (if applicable)
- Root segment allocation
- Tree segment description (if applicable)
- Undefined references (if applicable)
- Task Builder statistics

An example of the memory allocation file (map) is shown in Example 10-2. The numbered and lettered items in the notes correspond to the numbers and letters in Example 10-2.

Example 10-2 Memory Allocation File (Map) Example

```
OVR.TSK;1    Memory Allocation Map  TKB M40.10    Page 1  
             15-DEC-82    11:28
```

```
Partition name : GEN  (a) (b)  
Identification : 01  (c)  
Task  UIC      : [303,3] (d) (e)  
Stack  limits: 000260 001257 001000 00512. (f)
```

1

2

(continued on next page)

SH (Cont.)

Example 10-2 (Cont.) Memory Allocation File (Map) Example

Prg xfr address: 001264 (g) (h) (i)
 Total address windows: 1. (i) (k) (l)
 Task image size : 7488. words (m) (n)
 Task address limits: 000000 035107 (o)
 R-W disk blk limits: 000002 000073 000072 00058. (p) (q)

OVR.TSK Overlay description:

Base	Top	Length	
000000	005033	005034 02588.	ROOTM
005034	021057	014024 06164.	MULOV
005034	021057	014024 06164.	ADDOV
021060	035103	014024 06164.	SUBOV
021060	035107	014030 06168.	DIVOV

OVR.TSK Memory allocation map TKB M40.02 Page 2
 ROOTM 28-DEC-81 09:10

*** Root segment: ROOTM (a)

R/W mem limits: 000000 005033 005034 02588. (b)
 Disk blk limits: 000002 000007 000006 00006. (c)

Memory allocation synopsis:

Section	Title	Ident	File
. BLK.: (RW,I,LCL,REL,CON) (d)			
001260 002514 01356.	ROOTM	01 (e)	ROOTM.OBJ;1
001260 000102 00066.	PRNOV	01	PRNOV.OBJ;1
001362 000260 00176.	SAVOV	01	SAVOV.OBJ;1
001642 000042 00034.			
003774 000002 00002.			
ANS : (RW,D,GBL,REL,OVR)			
003774 000002 000002.	ROOTM	01	ROOTM.OBJ;1
003774 000002 00002.	PRNOV	01	PRNOV.OBJ;1
.			
.			
.			

Global Symbols:

AADD 004032-R DIVV 004052-R PRINT 001550-R SUBB 004042-R (f)
 MULL 004022-R SAVAL 001642-R

OVR.TSK Memory allocation map TKB M40.02 Page 3
 MULOV 28-DEC-81 09:10 (g) (h) (i) (j) (k)

*** Segment: MULOV

R/W mem limits: 005034 021057 014024 06164.
 Disk blk limits: 000010 000024 000015 00013.

(continued on next page)

Example 10-2 (Cont.) Memory Allocation File (Map) Example

Memory allocation synopsis:

Section	Title	Ident	File
-----	-----	-----	-----
. BLK.:(RW,I,LCL,REL,CON) 005034 014024 06164.			
	MULOV	01	MULOV.OBJ;1
\$\$ALVC:(RW,D,LCL,REL,CON) 021060 000000 00000.			
\$\$RTS :(RW,I,BGL,REL,OVR) 004250 000002 00002.			

Global symbols:

MULL 021034-R

.

*** Task builder statistics:

Total work file references: 7156. ^a
 Work file reads: 0. }
 Work file writes: 0. } ^b ^c
 Size of core pool: 7086. words (27. pages) ^d
 Size of work file: 3072. words (12. pages) ^e

ELAPSED TIME:00:00:14

Notes:

- ① The page header shows the name of the task image file and the overlay segment name (if applicable), along with the date, time, and version of TKB that created the map.
- ② The task attributes section contains the following information:
 - ③ Task name -- The name specified in the TASK option. If you do not use the TASK option, TKB suppresses this field.
 - ④ Partition name -- The partition specified in the PAR option. If you do not specify a partition, the default is partition GEN.
 - ⑤ Identification -- The task version as specified in the .IDENT assembler directive. If you do not specify the task identification, the default is 01.
 - ⑥ Task UIC -- The task UIC as specified in the UIC option. If you do not specify the UIC, the default is the terminal UIC.

SH (Cont.)

- ⓔ Task priority -- The priority of the task as specified in the PRI option. If you do not specify PRI, the default is 50, and is not shown on the map.
- ⓕ Stack limits -- The low and high octal addresses of the stack, followed by its length in octal and decimal bytes.
- ⓖ ODT transfer address -- The starting address of the ODT debugging aid. If you do not specify the ODT debugging aid, this field is suppressed.
- ⓗ Program transfer address -- The address of the symbol specified in the .END directive of the source code of your task. If you do not specify a transfer address for your task, TKB automatically establishes a transfer address of 000001 for it. TKB also suppresses this field in the map if you do not specify a transfer address.
- ⓓ Task attributes -- These attributes are listed only if they differ from the defaults. One or more of the following may be displayed:
 - AC Ancillary control processor.
 - AL Task is checkpointable, and task image file contains checkpoint space allocation.
 - CP Task is checkpointable, and task image file will be checkpointed to system checkpoint file.
 - DA Task contains debugging aid.
 - EA Task uses Kell-A extended arithmetic element.
 - FP Task uses Floating Point Processor.
 - HD Task image does not contain header.
 - PI Task contains position-independent code and data.
 - PM Postmortem Dump requested in the event of abnormal task termination.
 - PR Task is privileged.
 - SE Messages addressed to the task through the SEND directive will be rejected by the Executive.
 - SL Task can be slaved.
 - TR Task initial PS word has T-bit enabled.
 - ID Task is I- and D-space task.
- ⓓ Total address windows -- The number of window blocks allocated to the task.

SWITCHES

SH (Cont.)

- Ⓚ Mapped array -- The amount of physical memory (decimal words) allocated through the VSECT option or Mapped Array Declaration (GSD type 7, described in Appendix B); mapped array is not shown if it does not apply.
 - Ⓛ Task extension -- The increment of physical memory (decimal words) allocated through the EXTTSK or PAR option. Without these options, task extension is not shown.
 - Ⓜ Task image size -- The amount of memory (decimal words) required to contain your task's code. This number does not include physical memory allocated through the EXTTSK option.
 - Ⓝ Total task size -- The amount of physical memory (decimal words) allocated, including mapped array area and task extension area. Total task size is not shown in this example.
 - Ⓞ Task address limits -- The lowest and highest virtual addresses allocated to the task, exclusive of virtual addresses allocated to VSECTs and shared regions.
 - Ⓟ Read/write disk block limits -- From left to right: the first octal relative disk block number of the task's header; the last octal relative disk block number of the task image; and the total contiguous disk blocks required to accommodate the read/write portion of the task image in octal and decimal.
 - Ⓠ Read-only disk block limits -- From left to right: the first octal relative disk block of the multiuser task's read-only region; the last octal relative disk block number of the read-only region; and the total contiguous disk blocks required to accommodate the read-only region in octal and decimal. This field appears only when you are building a multiuser task.
- ③ The overlay description shows, for each overlay segment in the tree structure of an overlaid task, the beginning virtual address (the base), the highest virtual address (the top), the length of the segment in octal and decimal bytes, and the segment name. Indenting is used to illustrate the ascending levels in the overlay structure. TKB prints the overlay description only when an overlaid task is created.
- ④ The root segment allocation -- This section has the following elements:
- ⓐ Root segment -- The name of the root segment. If your task is a single-segment task, the entire task is considered to be the root segment.
 - ⓑ Read/write memory limits -- From left to right: the beginning virtual address of the root segment (the base); the virtual address of the last byte in the segment (the top); and the length of the segment in octal and decimal bytes.

SH (Cont.)

- Ⓒ Disk block limits -- From left to right: the first relative block number of the beginning of the root segment; the last relative block number of the root segment; total number of disk blocks in octal; and the total number of disk blocks in decimal.
- Ⓓ Memory allocation synopsis -- From left to right: the program section name; the program section attributes; starting virtual address of the program section; and total length of the program section in octal and decimal bytes.

The program section shown as .BLK. in this field is the unnamed relocatable program section. Notice in this example that there are 636(octal) bytes allocated to it (2034 bytes - 1176 bytes = 636 bytes). This allocation is the result of calls to routines that reside within the unnamed program section in SYSLIB. (For more information, see the description of the /MA switch in Section 10.1.14.)

- Ⓔ Module contributor -- This field lists the modules that have contributed to each program section. In this example, the program section ANS was defined in module ROOTM. The module version is 01 (as a result of the .IDENT assembler directive) and the file name from which the module was extracted is ROOTM.OBJ;1. If the program section ANS had been defined in more than one module, each contributing module and the file from which it was extracted would have been listed here.

NOTE

The absolute section .ABS. is not shown because it appears in every module and always has a length of 0.

- Ⓕ The global symbols section lists the global symbols defined in the segment. Each symbol is listed along with its octal value. A -R is appended to the value if the symbol is relocatable. The list is alphabetized in columns.

The file contents section (which is composed of the four fields listed below) is printed only if you specify /-SH in the TKB command sequence. TKB creates this section for each segment in an overlay structure. It lists the following information:

- Ⓖ Input file -- File name, module name as established by the .TITLE assembler directive, and module version as established by the .IDENT assembler directive.
- Ⓗ Program section -- Program section name, starting virtual address of the program section, ending virtual address of the program section, and length in octal and decimal bytes.

SH (Cont.)

- ① Global symbol -- Global symbol names within each program section and their octal values. If the segment is autoloading (see Chapter 3), this value is the address of an autoloading vector. The autoloading vector in turn contains the actual address of the symbol.

A -R is appended to the value if the symbol is relocatable.

- ② Program section -- The contents of this field is described in note g above.
 - ③ Undefined References -- This field lists the undefined global symbols in the segment.
- 5 The tree segment description is printed for every overlay segment in an overlay structure. Its contents are the same for each overlay segment as the root segment allocation is for the root segment.
 - 6 Task builder statistics lists the following information, which can be used to evaluate TKB performance:
 - a Work file references -- The number of times that TKB accessed data stored in its work file.
 - b Work file reads -- The number of times that the work file device was accessed to read work file data.
 - c Work file writes -- The number of times that the work file device was accessed to write work file data.
 - d Size of pool -- The amount of memory that was available for work file data and table storage.
 - e Size of work file -- The amount of device storage that was required to contain the work file.
 - f Elapsed time -- The amount of wall-clock time required to construct the task image and produce the memory allocation (map) file. Elapsed time is measured from the completion of option input to the completion of map output. This value excludes the time required to process the overlay description, parse the list of input file names, and create the cross-reference listing (if specified).

See Appendix F for a more detailed discussion of the work file.

SL

10.1.31 /SL -- Slave

File

Task image

Syntax

file.TSK/SL=file.OBJ

Description

This switch directs TKB to mark your task as a slave to an initiating task.

Effect

TKB attaches the slave attribute to your task. When your task successfully executes a Receive Data directive, the system gives the UIC and TI: device of the sending task to it. The slave task then assumes the identity and privileges of the sending task.

This switch only applies to your task if the system that you are using has multiuser protection. (Refer to your system generation manual for more information on multiuser protection and slave tasks.)

Default

/-SL

10.1.32 /SP -- Spool Map Output

File

Memory allocation (map)

Syntax

file.TSK,file.MAP/-SP=file.OBJ

Description

This switch determines whether TKB calls the print spooler to spool your memory allocation (map) file after task build.

Effect

By default, when you specify a map file in a TKB command sequence, TKB creates a map file on device SY0: and then has the file queued for listing on LP0:.

If you negate this switch (as shown in the Syntax section above), TKB creates the map file on device SY0: but does not call the print spooler to output it to LP0:

Default

/SP

NOTE

The PRT task must be installed to process the request to print the map.

SQ

10.1.33 /SQ -- Sequential

File

Task image

Syntax

file.TSK/SQ=file.OBJ

Description

This switch causes TKB to construct your task image from the program sections you specified, in the order that you input them.

Effect

If you use this switch, TKB collects all the references to a given program section from your input object modules, groups them according to their access code (RW followed by RO) and, within these groups, allocates memory for them in the order that you input them. However, the /SG switch affects program section ordering and can be used with the /SQ switch. See the /SG switch for further details.

Without the /SQ switch, TKB reorders the program sections alphabetically.

You use this switch to satisfy any adjacency requirements that existing code may have when you are converting it to run under RSX-11. Using this feature is otherwise discouraged for the following reasons:

- Standard library routines (such as FORTRAN I/O handling routines and FCS modules from SYSLIB) do not work properly.
- Sequential allocation can result in errors if you alter the order in which modules are linked.

Alternatively, you can achieve physical adjacency of program sections by selecting names alphabetically to correspond to the desired order.

Default

/-SQ

10.1.34 /SS -- Selective Search

File

Input

Syntax

file.TSK=file.OBJ,file.OBJ/SS

or

file.TSK=file.OBJ,file.STB/SS

or

file.TSK=file.OBJ,file.OLB/LB/SS

Description

The /SS switch directs TKB to include in its internal symbol table only those global symbols for which there is a previously undefined reference.

Effect

When processing an input file, TKB normally includes in its internal symbol table each global symbol it encounters within the file whether or not there are references to it. With the /SS switch attached to an input file, TKB checks each global symbol it encounters within that file against its list of undefined references. If TKB finds a match, it includes the symbol in its symbol table.

Default

/-SS

Example

Assume that you are building a task named SEL.TSK. The task is composed of input files containing global entry points and references (calls) to them as shown in Table 10-2.

Table 10-2
Input Files for SEL.TSK

Input File Name	Global Definition	Global Reference
IN1		A
IN2	A B C	
IN3		C
IN4	A B C	

SS (Cont.)

File IN2 and IN4 contain global symbols of the same name that represent entry points to different routines within their respective files. Assume that you want TKB to resolve the reference to global symbol A in IN1 to the definition for A in IN2. Assume further that you want TKB to resolve the reference to global symbol C in IN3 to the definition for C in IN4. By selecting the sequence of the input files properly and applying the /SS switch to files IN2 and IN4, TKB resolves the references correctly. The following command sequence illustrates the correct sequence:

```
TKB>SEL.TSK=IN1.OBJ,IN2.OBJ/SS,IN3.OBJ,IN4.OBJ/SS
```

TKB processes input files from left to right; therefore, in processing the above command sequence, TKB processes file IN1 first and encounters the reference to symbol A. There is no definition for A within IN1; therefore, TKB marks A as undefined and moves on to process file IN2. Because the /SS switch is attached to IN2, TKB limits its search of IN2 to symbols it has previously listed as undefined, in this case, symbol A. TKB finds a definition for A and places A in its symbol table. Because there are no undefined references to symbols B or C, TKB does not place either of these symbols in its symbol table.

NOTE

It is important to realize that the /SS switch affects only the way the Task Builder constructs its internal symbol table. The routines for which symbols B and C are entry points is included in the task image even though there are no references to them.

TKB moves on to IN3. It encounters the references to symbol C. Because TKB did not include symbol C from IN2 in its symbol table, it cannot resolve the reference to C in IN3. TKB marks symbol C as undefined and moves on to IN4.

When TKB processes IN4, it encounters the definition for C in that file and includes it in the table. Again, since the /SS switch is attached to IN4, TKB includes only C in its symbol table.

When TKB has completed its processing of the above command sequence, it has constructed a task image composed of all of the code from all of the modules, IN1 through IN4. However, only symbols A from IN2 and C from IN4 will appear in its internal symbol table.

SWITCHES

SS (Cont.)

NOTE

The example above does not represent good programming practice. It is included here to illustrate the effect of the /SS switch on TKB during a search sequence.

The /SS switch is particularly valuable when used to limit the size of the Task Builder's internal symbol table during the building of a privileged task that references the Executive's routines and data structures. By specifying the Executive's Symbol Definition File (STB) as an input file and applying the SS switch to it, TKB includes in its internal symbol table only those symbols in the Executive that the task references. An example of a TKB command sequence that illustrates this is shown below:

```
TKB>OUTFILE.TSK/PR:5=INFILE.OBJ,RSX11M.STB/SS
```

The above command sequence directs TKB to build a privileged task named OUTFILE.TSK from the input file INFILE.OBJ. The specification of the Executive's STB file as an input file with the SS switch applied to it directs TKB to extract from RSX11M.STB only those symbols for which there are references within OUTFILE.TSK.

TR

10.1.35 /TR -- Traceable

File

Task image

Syntax

file.TSK/TR=file.OBJ

Description

This switch directs TKB to make your task traceable.

Effect

TKB sets the T-bit in the initial PS word of your task. When your task is executed, a trace trap occurs when each instruction is completed.

Default

/-TR

SWITCHES

WI

10.1.36 /WI -- Wide Listing Format

File

Memory allocation (map)

Syntax

file.TSK,file.MAP/-WI=file.OBJ

Description

This switch controls the width of your map file.

Effect

By default, TKB formats a map file 132 columns wide. When you negate this switch (as in the Syntax section above), TKB formats the map file 80 columns wide.

Default

/WI

XH**10.1.37 /XH -- External Header (RSX-11M-PLUS only)****File**

Task image

Syntax

file.TSK/-XH= file.OBJ

Description

The /XH switch informs TKB that the task is not to have an external header.

Effect

In an RSX-11M-PLUS system, the effect of the /XH switch is two-fold: the header space in the task image is not destroyed when the task is checkpointed; and Executive pool space is conserved. A task built with the /XH switch does not have a header in Executive pool, but has a copy of its header, which the Executive uses, in space allocated physically contiguous to and below the task image. When the task is checkpointed, the system writes the entire task image and the header copy below the task into a checkpoint file. The header in the task image is left unchanged.

Note that if the task is also built with the /FP switch, the floating-point save area is not included in the task image but is in the header copy found below the task image.

Interaction with the INSTALL command:

On RSX-11M-PLUS, the INSTALL switch /XHR interacts with the TKB switch /XH. If you use /-XH, the task will have a pool-resident header always unless you rebuild the task to have an external header. If you use /XH, the task will have an external header, but the INSTALL switch /XHR can override this. The default use of /XH by TKB is /XH (external header) unless this is changed by the INSTALL command.

Default

/XH for RSX-11M-PLUS; overridden by /XHR on INSTALL

10.1.38 /XT[:n] -- Exit on Diagnostic

File

Task image

Syntax

file.TSK/XT:4=file.OBJ

Description

This switch specifies the number of acceptable errors. More than n errors are not acceptable.

Effect

TKB exits after encountering n errors. The number of errors can be specified as a decimal or octal number, using the convention:

n. indicates a decimal number (the decimal point must be included).

#n or n indicates an octal number.

If you do not specify n, TKB assumes that n is 1.

Default

/-XT

CHAPTER 11

OPTIONS

11.1 OPTIONS

Task Builder options provide you with the means to give TKB information about the characteristics of your task.

These options, which are listed in Table 11-1, can be divided into seven categories. The identifying abbreviation and a brief description of each category are listed below:

- **contr** You use control options to affect TKB execution. ABORT is the only member of this category. You can direct the Task Builder to abort the task build with this option.
- **ident** You use identification options to identify your task's characteristics. You can specify the name of your task, its priority, user identification code, and partition with options in this category.
- **alloc** You use allocation options to modify your task's memory allocation. With the options in this category, you can change the size of your task's stack and program sections. When you write programs in a high-level language, you can change the size of your work areas and buffers and declare the virtual base address and size of program sections. Finally, you can declare the number of additional window blocks (if any) that your task requires.
- **share** You use storage-sharing options to indicate to TKB that your task intends to access a shared region.
- **device** You use device-specifying options to specify the number of units required by your task, and the assignment of logical unit numbers to physical devices.
- **alter** You use the content-altering options to define a global symbol and value, or to introduce patches in your task image.
- **synch** You use synchronous trap options to define synchronous trap vectors.

Some TKB options are of interest to all users of the system; others are of interest only to high-level language programmers; and still others are of interest only to MACRO-11 programmers. Table 11-1 lists all the options alphabetically, and gives a brief description of each.

OPTIONS

Table 11-1
Task Builder Options

Option	Meaning	Interest ¹	Category
ABORT	Directs TKB to terminate a task build	H,M	contr
ABSPAT	Declares absolute patch values for conventional tasks or I-space in I-and D-space tasks	M	alter
ACTFIL	Declares number of files open simultaneously	H	alloc
ASG	Declares device assignment to logical units	H,M	device
CLSTR	Declares a group of shared regions accessed by the task and residing in the same virtual address space in the task	H,M	share
CMPRT ²	Declares completion routine for supervisor-mode library	H,M	ident
COMMON LIBR	Declare task's intention to access a memory-resident shared region	H,M	share
DSPPAT	Declares absolute patch values for conventional tasks or D-space in I-and D-space tasks	M	alter
EXTSCT	Declares extension of a program section	H,M	alloc
EXTTSK	Declares extension of the amount of memory owned by a task	H,M	alloc
FMTBUF	Declares extension of buffer used for processing format strings at run time	H	alloc
GBLDEF	Declares a global symbol definition	M	alter
GBLINC	Includes symbols in the .STB file	M	alter
GBLPAT	Declares a series of patch values relative to a global symbol	M	alter
GBLREF	Declares a global symbol reference	H,M	alter

(continued on next page)

1. The user interest range is indicated as follows:

- H indicates options of interest to high-level language (such as FORTRAN) programmers.
- M indicates options of interest to MACRO-11 programmers.

2. These options are applicable to RSX-11M-PLUS systems only.

OPTIONS

Table 11-1 (Cont.)
Task Builder Options

Option	Meaning	Interest ¹	Category
GBLXCL ²	Declares global symbols to be excluded from the .STB file	H,M	alter
LIBR	Declares task's intention to access a memory-resident shared region	H,M	share
MAXBUF	Declares an extension to the FORTRAN record buffer	H	alloc
ODTV	Declares the address and size of the debugging aid SST vector	M	synch
PAR	Declares partition name and dimensions	H,M	ident
PRI	Declares priority	H,M	ident
RESCOM RESLIB	Declare task's intention to access a memory-resident shared region	H,M	share
RESSUP ²	Declares task's intention to access a resident supervisor-mode library	H,M	share
ROPAR ²	Declares partition in which read-only portion of multiuser task is to reside	H,M	ident
STACK	Declares the size of the stack	H,M	alloc
SUPLIB ²	Declares task's intention to access a system-owned supervisor-mode library	H,M	share
TASK	Declares the name of the task	H,M	ident
TSKV	Declares the address of the task SST vector	M	synch
UIC	Declares the user identification code under which the task runs	H,M	ident
UNITS	Declares the maximum number of units	H,M	device
VSECT	Declares the virtual base address and size of a program section	H,M	alloc
WNDWS	Declares the number of additional address windows required by the task.	H,M	alloc

1. The user interest range is indicated as follows:

- H indicates options of interest to high-level language (such as FORTRAN) programmers.
- M indicates options of interest to MACRO-11 programmers.

2. These options are applicable to RSX-11M-PLUS systems only.

ABORT

11.1.1 ABORT -- Abort the Task-Build

You use the ABORT option when you discover that an earlier error in the terminal sequence causes TKB to produce an unusable task image.

The Task Builder, on recognizing the keyword ABORT, stops accepting input and restarts for another task build.

Syntax

ABORT=n

n

An integer value. The integer is required to satisfy the general form of an option; however, the value is ignored in this case.

Default

None

NOTE

If you type a CTRL/Z at any time, it causes TKB to stop accepting input and begin building the task.

The ABORT option is the only correct way for you to restart TKB if you discover an error and decide you do not want the Task Builder output.

ABSPAT

11.1.2 ABSPAT -- Absolute Patch

You use the ABSPAT option to declare a series of object-level patches starting at a specified base address. You may use this option for conventional or I-and D-space tasks. If used for an I-and D-space task, this option patches I-space. See the DSPPAT option. You can specify up to eight patch values.

Syntax

```
ABSPAT=seg-name:address:val1:val2...:val8
```

seg-name

The 1- to 6-character Radix-50 name of the segment.

address

The octal address of the first patch. The address can be on a byte boundary; however, two bytes are always modified for each patch: the addressed byte and the following byte.

val1

An octal number in the range of 0 through 177777 to be stored at address.

val2

An octal number in the range of 0 through 177777 to be stored at address+2

·
·
·

val8

An octal number in the range of 0 through 177777 to be stored at address+16.

NOTE

All patches must be within the segment address limits or TKB generates the following error message:

```
TKB--*DIAG*--Load address out of range in module name
```

ACTFIL

11.1.3 ACTFIL -- Number of Active Files

You use the ACTFIL option to declare the number of files that your task can have open simultaneously. For each active file that you specify, TKB allocates approximately 512 bytes.

If you specify less than four active files (the default), the ACTFIL option saves space. If you want your task to have more than four active files, you must use the ACTFIL option to make the additional allocation.

You must include a language Object Time System (OTS), such as FORTRAN, and record I/O service routines (FCS or RMS-11) in your task image for the extension to take place. The program section that is extended has the reserved name \$\$FSR1.

Syntax

```
ACTFIL=file-max
```

file-max

A decimal integer indicating the maximum number of files that can be open at the same time.

Default

```
ACTFIL=4
```


OPTIONS

ASG

11.1.4 ASG -- Device Assignment

The ASG option declares the physical device that is assigned to one or more logical units.

Syntax

```
ASG=device-name:unit-num1:unit-num2...:unit-num8
```

device-name

A 2-character alphabetic device name followed by a 1- or 2-digit octal unit number. If your task uses more than six logical units, you must use the UNITS option to specify the number of logical units that your task will use.

```
unit-num1  
unit-num2  
.  
.  
.  
unit-num8
```

Octal integers indicating the logical unit numbers.

Default

```
ASG=SY0:1:2:3:4,TI0:5,CL0:6
```

CLSTR

11.1.5 CLSTR -- System-Owned Cluster of Resident Libraries or Commons

The CLSTR option allows you to link your program to one to six shared regions, such as FMS, RMS, FORTRAN or BASIC+2, with a minimum of lost virtual address space for your task. CLSTR allows two to six shared regions in an `RSX-11M` system or an `RSX-11M-PLUS` system to reside in the same virtual address space in your task.

You use CLSTR to declare a cluster or group of system-owned, resident libraries or commons that your task intends to access and have reside at the same virtual address in the address space of your task.

The term "system-owned" means that TKB expects to find the commons or libraries named in the option and the symbol table associated with them under UFD [1,1] on device LB:.

Syntax

```
CLSTR=library_1,library_2,...library_n:switch:apr
```

library_n

The library names must be 1- to 6- character Radix-50 names. TKB expects to find a symbol definition file of the same name for each specified shared region under UFD [1,1] on device LB:. The first specification denotes the first or the default library, which is the library to which the task is mapped when the task starts up and remaps after any call to another library.

In an `RSX-11M` or `RSX-11M-PLUS` system, the total number of libraries to which a task may map is seven. The number of the component libraries in clusters is limited to a maximum of six. A cluster must contain a minimum of two libraries. It is possible to have two clusters of three libraries each or three clusters of two libraries each; any combination of number of clusters and libraries must equal at least two or a maximum of six. If six libraries are used in clusters, the task may map to only one other, separate library.

:switch

The switch `:RW` (read/write) or `:RO` (read-only) indicates the type of access the task requires. All shared regions in the cluster have the same type of access.

:apr

The apr is an integer in the range of 1 through 7 that specifies the first Active Page Register (APR) that you want TKB to reserve for the cluster of shared regions. You can specify it for a cluster made up of only position-independent shared regions. If you omit the APR parameter and all shared regions are position independent, TKB selects the highest available APR to map the cluster. A cluster can be made up of both position-independent and absolute shared regions. If one absolute shared region is present with position-independent shared regions, the position-independent shared regions assume the same base address as that of the absolute shared region. However, if you specify more than one absolute shared region, all must be built with the same base address.

OPTIONS

CLSTR (Cont.)

Default

None

NOTE

All but the first shared region in a cluster must be memory-resident overlaid libraries. The first shared region specified in the cluster option can be a single-segment structure (nonoverlaid) or an overlaid library.

CMPRT**11.1.6 CMPRT -- Completion Routine -- RSX-11M-PLUS Only**

The CMPRT option is available on RSX-11M-PLUS systems only. You use this option to identify a shared region as a supervisor-mode library. The CMPRT option requires an argument that specifies the entry point of the completion routine in the library. The completion routine switches the processor from supervisor to user mode and returns program control to the user task after the supervisor-mode library subroutine that was called from the user task has executed.

Two completion routines are available in SYSLIB:

- \$CMPCS restores only the carry bit in the user-mode PS.
- \$CMPAL restores all the condition code bits in the user-mode PS.

These routines perform all the necessary overhead to switch the processor from supervisor to user mode and return program control to the user task at the instruction following the call to a supervisor-mode library subroutine.

Although you can write your own completion routines, it is best to use either \$CMPCS or \$CMPAL whenever possible. Chapter 8 discusses completion routines in detail.

Syntax

CMPRT=name

name

A 1- to 6-character Radix-50 name identifying the completion routine.

Default

None

COMMON or LIBR

11.1.7 COMMON or LIBR -- System-Owned Resident Common or System-Owned Resident Library

The COMMON and LIBR options are functionally identical; they both declare that your task intends to access a system-owned shared region. However, by convention, the COMMON option identifies a shared region that contains only data, and the LIBR option identifies a shared region that contains only code.

If you use the COMMON option with an I- and D-space task, the common is mapped with D-space APRs only and therefore must contain data only.

If you use the LIBR option with an I- and D-space task, the library is overmapped with both I-space and D-space APRs.

The term "system-owned" means that TKB expects to find the common or library named in the keyword and the symbol definition file associated with it under UFD [1,1] on device LB:.

Syntax

COMMON=name:access-code[:apr]

or

LIBR=name:access-code[:apr]

name

The 1- to 6-character Radix-50 name specifying the common or library. TKB expects to find a symbol definition file having the same name as that of the common or library with an extension of .STB under [1,1] of device LB:.

access-code

The code RW (read/write) or the code RO (read-only) indicating the type of access the task requires.

NOTE

A privileged task can change data in or move data to a resident common even though the task has been linked to the common with read-only access.

COMMON or LIBR (Cont.)**apr**

An integer in the range of 1 through 7 that specifies the first Active Page Register (APR) that you want TKB to reserve for the shared region. TKB recognizes the APR only for a mapped system; you can specify it only for position-independent shared regions. If you omit the APR parameter and the shared region is position independent, TKB selects the highest available APR to map the region.

When a shared region is absolute, the base address of the region -- and therefore the APR that maps it -- is determined by the arguments in the PAR option when the region is built. Refer to PAR in Section 11.1.18.

Default

None

OPTIONS

DSPPAT

11.1.8 DSPPAT -- Absolute Patch for D-space

You use the DSPPAT option to declare a series of object-level patches starting at the specified base address. This option is for making patches to the D-space of an I- and D-space task. You may also use this option to patch a conventional task at any location. You can specify up to eight patch values.

Syntax

```
DSPPAT=seq-name:address:val1:val2:...:val8
```

seqname

The 1- to 6-character Radix-50 name of the segment.

address

The octal address of the first patch. The address can be on a byte boundary; however, two bytes are always modified for each patch: the addressed byte and the following byte.

val1

An octal number in the range of 0 through 177777 to be stored at address.

val2

An octal number in the range of 0 through 177777 to be stored at address+2.

.
.
.

val8

an octal number in the range of 0 through 177777 to be stored at address+16.

NOTE

All patches must be within the segment address limits or TKB generates the following error message:

```
TKB--*DIAG*--Load address out of range in module-name
```

EXTSCT

11.1.9 EXTSCT -- Program Section Extension

You use the EXTSCT option to extend a program section.

If the program section to be extended has the attribute CON (concatenated), TKB extends the section by the number of bytes you specify in the EXTSCT option. If the program section has the attribute OVR (overlay), TKB extends the section only if the length you specify in the EXTSCT option is greater than the length of the program section.

Syntax

```
EXTSCT=p-sect-name:extension
```

p-sect-name

A 1- to 6-character radix-50 name specifying the program section to be extended.

extension

An octal integer that specifies the number of bytes by which to extend the program section.

Example

In the following example, the program section BUFF is 200 bytes long.

```
EXTSCT=BUFF:250
```

The number of bytes by which TKB extends the program section BUFF depends on the CON/OVR attribute:

- For CON, the extension is 250 bytes.
- For OVR, the extension is 50 bytes.

TKB extends the program section if it encounters the program section name in an input object file or in the overlay description file.

Default

None

EXTTSK**11.1.10 EXTTSK -- Extend Task Memory**

The EXTTSK option directs the system to allocate additional memory for your task when it is installed in a system-controlled partition.

The amount of memory that the system allocates for your task is the sum of the task size plus the increment you specify (rounded up to the nearest 32-word boundary). If the task is built for a user-controlled partition, the allocation of task memory reverts to the partition size.

This option extends only the D- space of an I- and D-space task.

In an unmapped system, TKB ignores the EXTTSK keyword.

NOTES

1. You should not use the EXTTSK option to extend a task containing memory-resident overlays because the system does not map the extended area.
2. When you use the EXTTSK option to extend a checkpointable task that has been declared checkpointable with the /AL switch, the check point file within the task image is the size of the task plus the size of the extended task area.
3. Be careful when extending an I- and D-space task that is linked to a library which contains both data and instructions. Normally, libraries are mapped in both I-space and D-space allowing data and instructions to be intermixed. The extension length must not extend into the area mapped for the library or the library will be mapped in I-space only.

Syntax

EXTTSK=length

length

A decimal number in the range $0 < n < 65,535$. specifying the increase in task memory allocation (in words).

Default

The task is extended to the size specified in the PAR option (see Section 11.1.18).

FMTBUF

11.1.11 FMTBUF -- Format Buffer Size

The FMTBUF option declares the length of the internal working storage that you want TKB to allocate within your task for compiling format specifications at run time. The length of this area must equal or exceed the number of bytes in the longest format string to be processed.

Run-time compilation occurs whenever an array is referred to as the source of formatting information within a FORTRAN I/O statement. The program section that TKB extends has the reserved name \$SOBF1.

Syntax

```
FMTBUF=max-format
```

max-format

A decimal integer, larger than the default, that specifies the number of characters in the longest format specification.

Default

```
FMTBUF=132
```

GBLDEF

11.1.12 GBLDEF -- Global Symbol Definition

You use the GBLDEF option to declare the definition of a global symbol.

TKB considers this symbol definition to be absolute. It overrides any definition in your input object modules.

Syntax

```
GBLDEF=symbol-name:symbol-value
```

symbol-name

A 1- to 6-character Radix-50 name of the defined symbol.

symbol-value

An octal number in the range of 0 through 177777 assigned to the defined symbol.

Default

None

GBLINC

11.1.13 GBLINC -- Include Global Symbols

The GBLINC option directs the Task Builder to include the symbol or symbols specified in this option in the .STB file being generated by the link operation in which this option appears. This option is intended for use when creating shared regions, in particular shared libraries, when you want to force particular modules to be linked to your task that references this library. The global symbol references specified by this option must be satisfied by some module or GBLDEF specification when you build the task.

Syntax

```
GBLINC=symbol-name,symbol-name,....,symbol-name
```

symbol-name

The symbol to be included.

Default

None

GBLPAT

11.1.14 GBLPAT -- Global Relative Patch

The GBLPAT option declares a series of object-level patch values starting at an offset relative to a global symbol. You can specify up to eight patch values.

Syntax

```
GBLPAT=seg-name:sym-name[+/-offset]:val1:val2...:val8
```

seg-name

The 1- to 6-character Radix-50 name of the segment.

sym-name

A 1- to 6-character Radix-50 name specifying the global symbol.

offset

An octal number specifying the offset from the global symbol.

val1

An octal number in the range of 0 through 177777 to be stored at the octal address of the first patch.

val2

An octal number in the range of 0 through 177777 to be stored at the first address+2.

.

.

.

val8

An octal number in the range of 0 through 177777 to be stored at the first address+14.

Default

None

NOTE

All patches must be within the segment address limits or TKB generates a fatal error.

GBLREF

11.1.15 GBLREF -- Global Symbol Reference

The GBLREF option declares a global symbol reference. The reference originates in the root segment of the task. This keyword is used for memory-resident overlays of shared regions.

Syntax

```
GBLREF=symbol-name,symbol-name...,symbol-name
```

symbol-name

A 1- to 6-character name of a global symbol reference.

Default

None

GBLXCL**11.1.16 GBLXCL -- Exclude Global Symbols**

The GBLXCL option keyword directs TKB to exclude from the symbol definition file of a shared region the symbol(s) specified in the option.

Syntax

GBLXCL=symbol-name,symbol-name...,symbol-name

symbol-name

The symbol(s) to be excluded.

Default

None

LIBR

11.1.17 LIBR -- System-Owned Library

Refer to COMMON in Section 11.1.7.

MAXBUF**11.1.18 MAXBUF -- Maximum Record Buffer Size**

The MAXBUF option declares the maximum record buffer size required for any file used by the task.

If your task requires a maximum record size that exceeds the default buffer length, you must use this option to extend the buffer.

You must also include a language Object Time System (OTS), such as FORTRAN, in your task image for the extension to take place. The program section that is extended has the reserved name \$\$IOB1.

Syntax

MAXBUF=max-record

max-record

A decimal integer, larger than the default, that specifies the maximum record size in bytes.

Default

MAXBUF=133

ODTV

11.1.19 ODTV -- ODT SST Vector

The ODTV option declares that a global symbol is the address of the ODT Synchronous System Trap vector. You must define the global symbol in the main root segment of your task.

Syntax

ODTV=symbol-name:vector-length

symbol-name

A 1- to 6-character Radix-50 name of a global symbol.

vector-length

A decimal integer in the range of 1 through 32 specifying the length of the SST vector in words.

Default

None

PAR

11.1.20 PAR -- Partition

The PAR option identifies the partition for which your task is built.

In a mapped system, you can install your task in any system partition or user partition large enough to contain it. In an unmapped system, your task is bound to physical memory. Therefore, you must install your task in a partition starting at the same memory address as that of the partition for which it was built.

Syntax

```
PAR=pname[:base:length]
```

pname

The name of the partition.

base

The octal byte address defining the start of the partition. On an unmapped system, the physical address must be specified. On a mapped system, the base must be 0 for a task or a 4K boundary for a shared region.

length

The octal number of bytes contained in the partition.

In a mapped system, a length of 0 implies a system-controlled partition.

If the target system is mapped and you specify a partition length that is greater than the length of your task, the Task Builder automatically extends the length of your task to match the length of the partition. This procedure is equivalent to using the EXTTSK keyword to increase the task memory. If your task size is greater than the partition size that you specify, TKB generates the following error message:

```
TKB--*DIAG*-Task has illegal memory limits
```

Whether or not the target system is mapped, the Task Builder does not extend the length of a shared region, or any task built without a header, to match the specified partition length.

If you do not specify the base and length, TKB tries to obtain that information from the system on which you are building your task. If you have specified a partition that resides in that system, TKB can obtain the base and length.

PAR (Cont.)

TKB binds the task to the addresses defined by the partition base. If the partition is user controlled, TKB verifies that the task does not exceed the length specification.

On RSX-11M-PLUS, a TKB command sequence or build file for a memory-resident overlaid library must contain the statement PAR=xxx, where xxx is the same name as that of the region being built.

Default

PAR=GEN

11.1.21 PRI -- Priority

The PRI option declares your task's execution priority.

On systems with multiuser protection, you cannot run a task at a priority that is greater than the system priority (50) unless it is installed or run from a privileged terminal. If you are working from a privileged terminal, and you do not override this option by specifying a different priority when you install your task, the system uses this priority.

Syntax

PRI=priority-number

priority-number

A decimal integer in the range of 1 through 250

Default

Established by Install; refer to the RSX-11M/M-PLUS MCR
Operations Reference Manual.

RESCOM or RESLIB

11.1.22 RESCOM or RESLIB -- Resident Common or Resident Library

The RESCOM and RESLIB options are functionally identical; they both declare that your task intends to access a user-owned, shared region. However, by convention the RESCOM option identifies a shared region that contains only data and the RESLIB option identifies a shared region that contains only code.

If you use the RESCOM option with an I- and D-space task, the common is mapped with D-space APRs only and therefore must contain data only.

If you use the RESLIB option with an I- and D-space task, the library is overmapped with both I-space and D-space APRs.

The term "user-owned" means that the resident common or library and the symbol definition file associated with it can reside under any UFD that you choose. You can specify the UFD and remaining portions of the file specification for both options. You must not place comments on the same line with either option.

Syntax

```
RESCOM=file-specification/access-code[:apr]
```

or

```
RESLIB=file-specification/access-code[:apr]
```

file-specification

The memory image file of the resident common or resident library. The file specification format is discussed in Chapter 1.

access-code

The code RW (read/write) or the code RO (read-only), indicating the type of access required by the task.

NOTE

A privileged task can change data in or move data into a resident common even though the task has been linked to the common with read-only access.

apr

An integer in the range of 1 through 7 that specifies the first Active Page Register (APR) that you want TKB to reserve for the common or library. TKB recognizes the APR argument only for a mapped system. You can specify it only for position-independent shared regions. If the APR parameter is omitted and the shared region is position independent, TKB selects the highest available APR to map the region.

When a shared region is absolute, the base address of the region -- and therefore the APR that maps it -- is determined by the arguments in the PAR option when the region is built. Refer to PAR in Section 11.1.18.

OPTIONS

RESCOM or RESLIB (Cont.)

NOTES

1. The Task Builder expects to find a symbol definition file having the same name as that of the memory image file but with a file type of .STB, on the same device and under the same UFD as that of the memory image file.
2. Regardless of the version number you give in the file specification, TKB uses the latest version of the .STB file.

Default

When you omit portions of the file-specification, the following defaults apply:

- UFD - Taken from current terminal UIC
- Device - SY0:
- File type - .TSK
- File version - Latest

RESLIB

11.1.23 RESLIB -- Resident Library

Refer to RESCOM in Section 11.1.21.

RESSUP**11.1.24 RESSUP -- Resident Supervisor-Mode Library -- RSX-11M-PLUS only**

The RESSUP option declares that your task intends to access a user-owned, supervisor-mode library. The term "user-owned" means that the library and the symbol definition file associated with it can reside under any UFD that you choose. You can specify the UFD and remaining portions of the file specification. You must not place comments on the line with RESSUP.

Syntax

```
RESSUP=file-specification/[-]SV[:apr]
```

file-specification

The memory image file of the supervisor-mode library. The file specification has the standard RSX-11M/RSX-11M-PLUS format discussed in Chapter 1.

/-**]SV**

The code /SV or /-SV to indicate whether TKB includes mode-switching vectors within the user task. If you specify /SV, TKB includes a 4-word, mode-switching vector within the address space of the user task for each call to a supervisor-mode library subroutine. If you specify /-SV, you must provide your own mode-switching vector. Providing your own mode-switching vectors is useful if your library contains threaded code. It is best to use the system-supplied vectors whenever possible.

apr

An integer in the range of 0 through 7 that specifies the first Supervisor Active Page Register that you want TKB to reserve for your supervisor-mode library. You can specify an APR only for position-independent, supervisor-mode libraries. The default is the lowest available APR.

The library at virtual 0 must have the CSM dispatcher present in the system-supplied completion routine described in Chapter 8.

NOTES

1. The Task Builder expects to find a symbol definition file having the same name as that of the memory image file but with a file type of .STB, on the same device and under the same UFD as that of the memory image file.
2. Regardless of the version number you give in the file specification, TKB uses the latest version of the .STB file.

RESSUP (Cont.)**Default**

When you omit portions of the file specification, the following defaults apply:

- UFD - Taken from the current terminal UIC
- Device - SY0:
- File type - .TSK
- File version - Latest

ROPAR**11.1.25 ROPAR -- Read-Only Partition -- RSX-11M-PLUS Only**

You use this option to declare the partition in which the read-only portion of your multiuser task is to reside.

Syntax

ROPAR=parname

parname

The partition name in which your multiuser task is to reside.

Default

The partition in which the read/write portion of the task resides.

STACK

11.1.26 STACK -- Stack Size

The STACK option declares the maximum size of the stack required by your task.

The stack is an area of memory that the MACRO-11 programmer uses for temporary storage, subroutine calls, and synchronous trap service linkage. The stack is referred to by hardware register 6 (SP, the stack pointer).

Syntax

```
, STACK=stack-size
```

stack-size

A decimal integer specifying the number of words required for the stack.

Default

```
STACK=256
```

SUPLIB

11.1.27 SUPLIB -- Supervisor-Mode Library -- RSX-11M-PLUS Only

This option declares that your task intends to access a system-owned, supervisor-mode library. The term "system-owned" means that TKB expects to find the supervisor-mode library and the symbol definition file associated with it in UFD [1,1] on device LB:.

Syntax

`SUPLIB=name:[-]SV[:apr]`

name

The 1- to 6-character Radix-50 name specifying the system-owned, supervisor-mode library. TKB expects to find a symbol definition file having the same name as that of the library with a file version of .STB under [1,1] of device LB:.

:[-]SV

The code /SV or /-SV to indicate whether TKB includes mode-switching vectors within the user task. If you specify /SV, TKB includes a 4-word mode-switching vector within the address space of the user task for each call to a supervisor-mode library subroutine. If you specify /-SV, you must provide your own mode-switching vector. Providing your own mode-switching vectors is useful if your library contains threaded code. It is best to use the system-supplied vectors whenever possible.

apr

An integer in the range of 0 through 7 that specifies the first Supervisor Active Page Register that TKB is to reserve for the library. You can specify an APR only for position-independent, supervisor-mode libraries. The default is the lowest available APR.

The library at virtual 0 must have the CSM dispatcher present in the system-supplied completion routine described in Chapter 8.

Default

None

TASK

11.1.28 TASK -- Task Name

The TASK option gives your task an installed name different from its task image name.

Syntax

```
TASK=task-name
```

task-name

A 1- to 6-character name identifying your task.

Default

The first six characters of the task image file name identify the task when the task is installed.

TSKV

11.1.29 TSKV -- Task SST Vector

The TSKV option declares that a global symbol is the address of the task Synchronous System Trap (SST) vector. You must define the global symbol in the main root segment of your task.

Syntax

TSKV=symbol-name:vector-length

symbol-name

A 1- to 6-character name of a global symbol.

vector-length

A decimal integer in the range of 1 through 32 specifying the length of the SST vector in words.

Default

None

UIC

11.1.30 UIC -- User Identification Code

The UIC option declares the User Identification Code (UIC) for your task when you run it with a time-based schedule request.

Syntax

```
UIC=[group,member]
```

group

An octal number in the range of 1 through 377, or a decimal number in the range of 1 through 255. Decimal numbers must be followed by a decimal point (.).

member

An octal number in the range of 1 through 377, or a decimal number in the range of 1 through 255. Decimal numbers must be followed by a decimal point (.).

Default

The UIC that the Task Builder is running under (normally the terminal UIC).

UNITS

11.1.31 UNITS -- Logical Unit Usage

The UNITS option declares the number of logical units that your task uses.

Syntax

```
UNITS=max-units
```

max-units

A decimal integer in the range of 0 through 250 specifying the maximum number of logical units. A 2-word block is allocated in the task's header for every logical unit. A task that uses many logical units can use a significant portion of dynamic memory because the header is in dynamic memory when the task is executing. The /XH switch affects pool usage by the task header.

Default

```
UNITS=6
```

VSECT

11.1.32 VSECT -- Virtual Program Section

The VSECT option specifies the virtual base address, virtual length, and, optionally, the physical memory allocated to the named program section. Refer to Chapter 5 for more information on virtual program sections.

Syntax

```
VSECT=p-sect-name:base>window[:physical-length]
```

p-sect-name

A 1- to 6-character program section name.

base

An octal value representing the virtual base address of the program section in the range of 0 through 177777. If you use the mapping directives, the value you specify must be a multiple of 4K.

window

An octal value specifying the amount of virtual address space in bytes allocated to the program section. Base plus window must not exceed 177777 (octal).

physical-length

An octal value specifying the minimum amount of physical memory to be allocated to the section in units of 64-byte blocks. TKB rounds this value up to the next 256-word limit. This value, when added to the task image size and any previous allocation, must not cause the total to exceed 2048K bytes. If you do not specify a length, TKB assumes a value of 0.

Default

Physical-length defaults to 0.

WNDWS**11.1.33 WNDWS -- Number of Address Windows**

The WNDWS option declares the number of address windows required by the task in addition to those needed to map the task image, and any mapped array or shared region. The number specified is equal to the number of simultaneously mapped regions the task will use.

Syntax

WWDWS=n

n

An integer in the range 1 through 7 in an RSX-11M system and 1 through 23 in an RSX-11M-PLUS system.

Default

WWDWS=0

APPENDIX A

TASK BUILDER INPUT DATA FORMATS

An object module is the fundamental unit of input to the Task Builder (TKB). You create an object module by using any of the standard language processors (for example, MACRO-11 or FORTRAN) or by using TKB itself (symbol definition file). The RSX-11M/M-PLUS librarian (LBR) gives you the capability to combine a number of object modules into a single library file.

An object module consists of variable-length records of information that describe the contents of the module. These records guide TKB in translating the object language into a task image. Six record (block) types are included in the object language:

- Declare global symbol directory (GSD) record (type 1)
- End of global symbol directory (GSD) record (type 2)
- Text information (TXT) record (type 3)
- Relocation directory (RLD) record (type 4)
- Internal symbol directory (ISD) record (type 5)
- End-of-module record (type 6)

TKB requires at least five of these record types in each object module. The only record type that it does not require is the internal symbol directory.

The various record types are defined according to a prescribed format, as illustrated in Figure A-1. An object module must begin with a declare-GSD record and end with an end-of-module record. Additional declare-GSD records can occur anywhere in the file, but must occur before an end-of-GSD record. An end-of-GSD record must appear before the end-of-module record, and at least one RLD record must appear before the first TXT record. Additional RLD and TXT records can appear anywhere in the file. The ISD records can appear anywhere in the file between the initial declare-GSD record and the end-of-module record.

Object module records are variable length and are identified by a record type code in the first byte of the record. The format of additional information in the record depends on the record type.

TASK BUILDER INPUT DATA FORMATS

The following sections describe each of the six record types in greater detail. The outline of these sections is as follows:

- A.1 Declare Global Symbol Directory Record
 - A.1.1 Module Name (Type 0)
 - A.1.2 Control Section Name (Type 1)
 - A.1.3 Internal Symbol Name (Type 2)
 - A.1.4 Transfer Address (Type 3)
 - A.1.5 Global Symbol Name (Type 4)
 - A.1.6 Program Section Name (Type 5)
 - A.1.7 Program Version Identification (Type 6)
 - A.1.8 Mapped Array Declaration (Type 7)
 - A.1.9 Completion Routine Name (Type 10)
- A.2 End of Global Symbol Directory Record
- A.3 Text Information Record
- A.4 Relocation Directory Record
 - A.4.1 Internal Relocation (Type 1)
 - A.4.2 Global Relocation (Type 2)
 - A.4.3 Internal Displaced Relocation (Type 3)
 - A.4.4 Global Displaced Relocation (Type 4)
 - A.4.5 Global Additive Relocation (Type 5)
 - A.4.6 Global Additive Displaced Relocation (Type 6)
 - A.4.7 Location Counter Definition (Type 7)
 - A.4.8 Location Counter Modification (Type 10)
 - A.4.9 Program Limits (Type 11)
 - A.4.10 Program Section Relocation (Type 12)
 - A.4.11 Program Section Displaced Relocation (Type 14)
 - A.4.12 Program Section Additive Relocation (Type 15)
 - A.4.13 Program Section Additive Displaced Relocation (Type 16)
 - A.4.14 Complex Relocation (Type 17)
 - A.4.15 Resident Library Relocation (Type 20)
- A.5 Internal Symbol Directory Record
- A.6 End of Module Record

A.1 DECLARE GLOBAL SYMBOL DIRECTORY RECORD

The global symbol directory (GSD) record contains all the information required by TKB to assign addresses to global symbols and to allocate the virtual address space required by a task.

GSD records are the only records processed by TKB in its first pass; therefore, you can save substantial time by placing all GSD records at the beginning of a module (because the Task Builder has to read less of the file).

GSD records contain nine types of entries:

- Module name (type 0)
- Control section name (type 1)
- Internal symbol name (type 2)
- Transfer address (type 3)
- Global symbol name (type 4)

TASK BUILDER INPUT DATA FORMATS

- Program section name (type 5)
- Program version identification (type 6)
- Mapped array declaration (type 7)
- Completion routine name (type 10)

TASK BUILDER DATA FORMATS

GSD	Initial Declare GSD
RLD	Initial Relocation Directory
GSD	Additional GSD
TXT	Text Information
TXT	Text Information
RLD	Relocation Directory
.	
.	
.	
GSD	Additional GSD
END GSD	End of GSD
ISD	Internal Symbol Directory
ISD	Internal Symbol Directory
TXT	Text Information
TXT	Text Information
TXT	Text Information
END MODULE	End of Module

ZK-444-81

Figure A-1 General Object Module Format

Each entry type is represented by four words in the GSD record. As shown in Figure A-2, the first two words contain six Radix-50 characters, the third word contains a flag byte and the entry type identification, and the fourth word contains additional information about the entry.

TASK BUILDER INPUT DATA FORMATS

A.1.1 Module Name (Type 0)

The module name entry (two words) declares the name of the object module. The name need not be unique with respect to other object modules (that is, modules are identified by file, not module name), but only one such declaration can occur in any given object module. Figure A-3 illustrates the module entry name format.

0	RECORD TYPE = 1
RAD50 NAME	
ENTRY TYPE	FLAGS
VALUE	
RAD50 NAME	
TYPE	FLAGS
VALUE	

•
•
•

RAD50 NAME	
TYPE	FLAGS
VALUE	
RAD50 NAME	
TYPE	FLAGS
VALUE	

ZK-445-81

Figure A-2 Global Symbol Directory Record Format

MODULE NAME (2 WORDS)	
ENTRY TYPE = 0	0
0	

ZK-446-81

Figure A-3 Module Name Entry Format

TASK BUILDER INPUT DATA FORMATS

A.1.2 Control Section Name (Type 1)

Control sections, which include absolute sections (ASECTs), blank, and named control sections (CSECTs), are replaced in RSX-11M by program sections (PSECTs). For compatibility with other systems, TKB processes ASECTs and both forms of CSECTs. Section A.1.6 details the entry generated for a .PSECT directive.

ASECTs and CSECTs are defined in terms of .PSECT directives, as follows:

For a blank CSECT, a program section is defined with the following attributes:

```
.PSECT ,LCL,REL,CON,RW,I,LOW
```

For a named CSECT, the program section is defined as:

```
.PSECT name, GBL,REL,OVR,RW,I,LOW
```

For an ASECT, the program section is defined as:

```
.PSECT . ABS.,GBL,ABS,I,OVR,RW,LOW
```

TKB processes ASECTs and CSECTs as program sections with the fixed attributes defined above. Figure A-4 illustrates the control section entry name format.

CONTROL SECTION NAME (2 WORDS)	
ENTRY TYPE = 1	IGNORED
MAXIMUM LENGTH	

ZK-447-81

Figure A-4 Control Section Name Entry Format

A.1.3 Internal Symbol Name (Type 2)

The internal symbol name entry (two words) declares the name of an internal symbol (with respect to the module). TKB does not support internal symbol tables; therefore, the detailed format of this entry is undefined. If TKB encounters an internal symbol entry while reading the GSD, it ignores that entry. Figure A-5 illustrates the internal symbol name entry format.

SYMBOL NAME (2 WORDS)	
ENTRY TYPE = 2	0
UNDEFINED	

ZK-448-81

Figure A-5 Internal Symbol Name Entry Format

TASK BUILDER INPUT DATA FORMATS

A.1.4 Transfer Address (Type 3)

The transfer address entry declares the transfer address of a module relative to a program section. The first two words of the entry define the name of the program section, and the fourth word defines the relative offset from the beginning of that program section. If a transfer address is not declared in a module, then a transfer address must not be included in the GSD, or a transfer address of 000001 relative to the default absolute program section (. ABS.) must be specified. Figure A-6 illustrates the transfer address entry format.

NOTE

If the program section is absolute, the offset is the actual transfer address (if not 000001).

SYMBOL NAME (2 WORDS)	
ENTRY TYPE = 3	0
OFFSET	

ZK-449-81

Figure A-6 Transfer Address Entry Format

A.1.5 Global Symbol Name (Type 4)

The global symbol name entry declares either a global reference or a definition. Definition entries must appear after the declaration of the program section in which the global symbols are defined and before the declaration of another program section (see Section A.1.6). Global references can be used anywhere within the GSD.

As shown in Figure A-7, the first two words of the entry define the name of the global symbol. The flag byte of the third word declares the attributes of the symbol, and the fourth word defines the value of the symbol relative to the program section in which the symbol is defined.

SYMBOL NAME (2 WORDS)	
ENTRY TYPE = 4	FLAGS
VALUE	

ZK-450-81

Figure A-7 Global Symbol Name Entry Format

Table A-1 lists the bit assignments of the flag byte of the symbol declaration entry.

TASK BUILDER INPUT DATA FORMATS

Table A-1
Symbol Declaration Flag Byte -- Bit Assignments


Bit Number and Name	Setting	Meaning
0 Weak qualifier	0	The symbol has a strong definition and is resolved in the normal manner.
	1	The symbol has a weak definition or reference. TKB ignores a weak reference (bit 3 = 0). It also ignores a weak definition (bit 3 = 1) unless a previous reference has been made.
1		Not used.
2 Definition or reference type	0	Normal definition or reference.
	1	Library definition. If the symbol is defined in a resident library STB file, the base address of the library is added to the value and the symbol is converted to absolute (bit 5 is reset); otherwise, the bit is ignored.
3 Definition	0	Global symbol reference.
	1	Global symbol definition.
4		Not used.
5 Relocation	0	Absolute symbol value.
	1	Relative symbol value.
6		Not used.
7		Not used.

A.1.6 Program Section Name (Type 5)

The program section name entry declares the name of a program section and its maximum length in the module. It also uses the flag byte to declare the attributes of the program section.

TASK BUILDER INPUT DATA FORMATS

You must construct GSD records such that once a program section name has been declared, all global symbol definitions pertaining to it must appear before another program section name is declared. Global symbols are declared with symbol declaration entries. Thus, the normal format is a series of program section names each followed by optional symbol declarations. Figure A-8 illustrates the program section name entry format.



ZK-451-81

Figure A-8 Program Section Name Entry Format

Table A-2 lists the bit assignments of the flag byte of the program section name entry.

Table A-2
Program Section Name Flag Byte -- Bit Assignments

Bit Number and Name	Setting	Meaning
0 Save	0	Normal program section.
	1	The program section is forced into the root of the task.
1 Library program section	0	Normal program section.
	1	The program section is relocatable and refers to a shared region.
2 Allocation	0	Program section references are to be concatenated with other references to the same program section to form the total memory allocated to the section.
	1	Program section references are to be overlaid. The total memory allocated to the program section is the largest request made by individual references to the same program section.

(continued on next page)

TASK BUILDER INPUT DATA FORMATS

Table A-2 (Cont.)
Program Section Name Flag Byte -- Bit Assignments

Bit Number and Name	Setting	Meaning
3		Not used; reserved for future DIGITAL use.
4 Access	0	The program section has read/write access.
	1	The program section has read-only access.
5 Relocation	0	The program section is absolute and requires no relocation.
	1	The program section is relocatable and references to the control section must have a relocation bias added before they become absolute.
6 Scope	0	The scope of the program section is local. References to the same program section are collected only within the segment in which the program section is defined.
	1	The scope of the program section is global. TKB collects references to the program section across segment boundaries. The Task Builder determines the segment in which storage is allocated for a global program section either by the first module that defines the program section on a path, or by direct placement of a program section in a segment using the ODL .PSECT directive.
7 Type	0	The program section contains instruction (I) references.
	1	The program section contains data (D) references.

NOTE

The length of all absolute sections is 0.

TASK BUILDER INPUT DATA FORMATS

A.1.7 Program Version Identification (Type 6)

The program version identification entry declares the version of the module. TKB saves the version identification of the first module that defines a nonblank version. It then includes this identification on the memory allocation map and writes the identification in the label block of the task image file.

The first two words of the entry contain the version identification. The flag byte and fourth words are not used and contain no meaningful information. Figure A-9 illustrates the program version identification entry format.

SYMBOL NAME	
ENTRY TYPE = 6	0
0	

ZK-452-81

Figure A-9 Program Version Identification Entry Format

A.1.8 Mapped Array Declaration (Type 7)

The mapped array declaration entry allocates space within the mapped array area of task memory. The array name is added to the list of task program section names and may be referred to by subsequent RLD records. The length (in units of 64-byte blocks) is added to the task's mapped array allocation. The total memory allocated to each mapped array is rounded up to the nearest 512-byte boundary. The contents of the flag byte are reserved and assumed to be 0.

One additional window block is allocated whenever a mapped array is declared.

Figure A-10 illustrates the mapped array declaration entry format.

MAPPED ARRAY	
NAME	
ENTRY TYPE = 7	FLAGS
LENGTH (NUMBER OF 64-BYTE BLOCKS)	

ZK-453-81


Figure A-10 Mapped Array Declaration Entry Format

TASK BUILDER INPUT DATA FORMATS

A.1.9 Completion Routine Definition (Type 10)

The completion routine definition declares the entry point for the completion routine of a supervisor-mode library. This data structure is created by the Task Builder and appears only in symbol definition files of supervisor-mode libraries.

As shown in Figure A-11, the first two words of the entry define the name of the entry point. The third word contains the entry type byte and the flag byte. The flag byte contains no meaningful information. The fourth word contains the symbol value.



ZK-454-81

Figure A-11 Completion Routine Entry Format

A.2 END OF GLOBAL SYMBOL DIRECTORY RECORD

The end of global symbol directory (end-of-GSD) record declares that no other GSD records are contained further on in the module. There must be exactly one end-of-GSD record in every object module. As shown in Figure A-12, this record is one word long.



ZK-455-81

Figure A-12 End of Global Symbol Directory Record Format

A.3 TEXT INFORMATION RECORD

The text information (TXT) record contains a byte string of information that is to be written directly into the task image file. The record consists of a load address followed by the byte string.

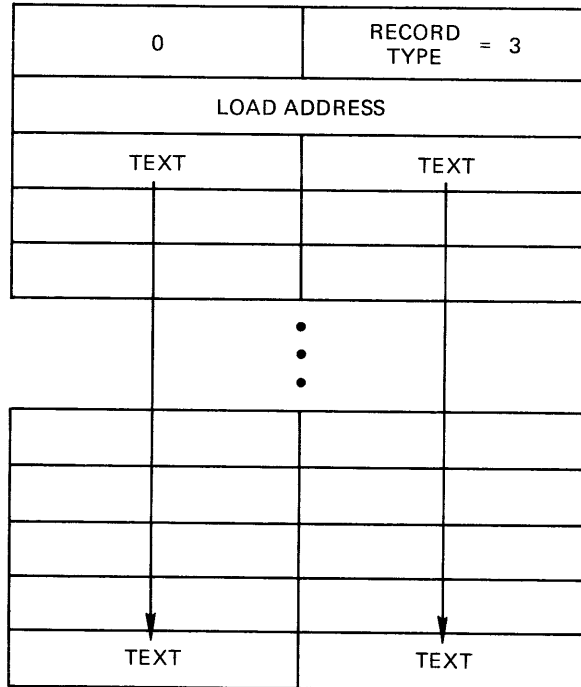
TXT records can contain words and/or bytes of information whose final contents have not yet been determined. This information will be bound by a relocation directory record that immediately follows the text record (see Section A.4). If the TXT record needs no modification, then no relocation directory record is needed. Thus, multiple TXT records can appear in sequence before a relocation directory record.

The load address of the TXT record is specified as an offset from the current program section base. At least one relocation directory record must precede the first TXT record. This directory must declare the current program section.

TASK BUILDER INPUT DATA FORMATS

TKB writes a text record directly into the task image file and computes the value of the load address minus 4. This value is stored in anticipation of a subsequent relocation directory that modifies words and/or bytes contained in the TXT record. When added to a relocation directory displacement byte, this value yields the address of the word and/or byte to be modified in the task image.

Figure A-13 illustrates the TXT record format.



ZK-456-81

Figure A-13 Text Information Record Format

A.4 RELOCATION DIRECTORY RECORD

The relocation directory (RLD) record contains the information necessary to relocate and link the preceding TXT record. Every module must have at least one RLD record that precedes the first TXT record. The first RLD record does not modify a preceding TXT record; rather, it defines the current program section and location. RLD records contain 15 types of entries, classified as relocation or location modification entries:

- Internal relocation (type 1)
- Global relocation (type 2)
- Internal displaced relocation (type 3)
- Global displaced relocation (type 4)

TASK BUILDER INPUT DATA FORMATS

- Global additive relocation (type 5)
- Global additive displaced relocation (type 6)
- Location counter definition (type 7)
- Location counter modification (type 10)
- Program limits (type 11)
- Program section relocation (type 12)
- Program section displaced relocation (type 14)
- Program section additive relocation (type 15)
- Program section additive displaced relocation (type 16)
- Complex relocation (type 17)
- Resident library relocation (type 20)

Each type of entry is represented by a command byte that specifies the type of entry and the word/byte modification, followed by a displacement byte, and then by the information required for the particular type of entry. The displacement byte, when added to the value calculated from the load address of the preceding TXT record (see Section A.3), yields the virtual address in the image that is to be modified.

Table A-3 lists the bit assignments of the command byte of each RLD entry.

Table A-3
Relocation Directory Command Byte --
Bit Assignments

Bit Number and Name	Setting	Meaning	
0-6	Entry type	Potentially, 128 command types can be specified; currently, 15 are implemented.	
7	Modification	0	The command modifies an entire word.
		1	The command modifies only one byte. TKB checks for truncation errors in byte modification commands. If truncation is detected (that is, if the modification value is greater than 255), an error occurs.

Figure A-14 illustrates the RLD record format.


TASK BUILDER INPUT DATA FORMATS

For example:

```

A:      MOV      #A,R0
        or
        .WORD    A
    
```

Figure A-15 illustrates the internal relocation entry format.



ZK-458-81

Figure A-15 Internal Relocation Entry Format

A.4.2 Global Relocation (Type 2)


The global relocation entry relocates a direct pointer to a global symbol. TKB obtains the definition of the global symbol and writes the result into the task image file at the calculated address.

For example:

```

MOV      #GLOBAL,R0
        or
        .WORD    GLOBAL
    
```

Figure A-16 illustrates the global relocation entry format.



ZK-459-81

Figure A-16 Global Relocation Entry Format

A.4.3 Internal Displaced Relocation (Type 3)

The internal displaced relocation entry relocates a relative reference to an absolute address from within a relocatable control section. TKB subtracts the address plus 2 that the relocated value is to be written into from the specified constant, and writes the result into the task image file at the calculated address.


For example:

```

CLR      177550
        or
MOV      177550,R0
    
```

TASK BUILDER INPUT DATA FORMATS

Figure A-17 illustrates the internal displaced relocation entry format.



ZK-460-81

Figure A-17 Internal Displaced Relocation Entry Format

A.4.4 Global Displaced Relocation (Type 4)

The global displaced relocation entry relocates a relative reference to a global symbol. TKB obtains the definition of the global symbol; subtracts the address plus 2 that the relocated value is to be written into from the definition value; and writes the result into the task image file at the calculated address.


For example:

CLR GLOBAL

or

MOV GLOBAL,R0

Figure A-18 illustrates the global displaced relocation entry format.



ZK-461-81

Figure A-18 Global Displaced Relocation Entry Format

A.4.5 Global Additive Relocation (Type 5)

The global additive relocation entry relocates a direct pointer to a global symbol with an additive constant. TKB obtains the definition of the global symbol; adds the specified constant to the definition value; and writes the result into the task image file at the calculated address.

For example:

MOV #GLOBAL+2,R0

or

.WORD GLOBAL-4

Figure A-19 illustrates the global additive relocation entry format.

TASK BUILDER INPUT DATA FORMATS

DISP	B	ENTRY TYPE = 5
SYMBOL NAME (2 WORDS)		
CONSTANT		

ZK-462-81

Figure A-19 Global Additive Relocation Entry Format

A.4.6 Global Additive Displaced Relocation (Type 6)

The global additive displaced relocation entry relocates a relative reference to a global symbol with an additive constant. TKB obtains the definition of the global symbol; adds the specified constant to the definition value; subtracts the address plus 2 that the relocated value is to be written into from the resultant additive value; and writes the result into the task image file at the calculated address.

For example:

```

CLR      GLOBAL+2
      or
MOV      GLOBAL-5,R0
    
```

Figure A-20 illustrates the global additive displaced relocation entry format.

DISP	B	ENTRY TYPE = 6
SYMBOL NAME (2 WORDS)		
CONSTANT		

ZK-463-81

Figure A-20 Global Additive Displaced Relocation Entry Format

A.4.7 Location Counter Definition (Type 7)

The location counter definition entry declares a current program section and location counter value. TKB stores the control base as the current control section; adds the current control section base to the specified constant; and stores the result as the current location counter value.

Figure A-21 illustrates the location counter definition entry format.

TASK BUILDER INPUT DATA FORMATS

0	B	ENTRY TYPE = 7
PROGRAM SECTION NAME (2 WORDS)		
CONSTANT		

ZK-464-81

Figure A-21 Location Counter Definition Entry Format

A.4.8 Location Counter Modification (Type 10)

The location counter modification entry modifies the current location counter. TKB adds the current program section base to the specified constant and stores the result as the current location counter.

For example:

```

.=.+N
    or
.BLKB  N
    
```

Figure A-22 illustrates the location counter modification entry format.

0	B	ENTRY TYPE = 10
CONSTANT		

ZK-465-81

Figure A-22 Location Counter Modification Entry Format

A.4.9 Program Limits (Type 11)


The program limits entry is generated by the .LIMIT assembler directive. TKB obtains the first address above the header (normally the beginning of the stack) and the highest address allocated to the task. It then writes these two addresses into the task image file at the calculated address and at the calculated address plus 2, respectively.

For example:

```
.LIMIT
```

Figure A-23 illustrates the program limits entry format.

TASK BUILDER INPUT DATA FORMATS



ZK-466-81

Figure A-23 Program Limits Entry Format

A.4.10 Program Section Relocation (Type 12)

The program section relocation entry relocates a direct pointer to the beginning address of another program section (other than the program section in which the reference is made) within a module. TKB obtains the current base address of the specified program section and writes it into the task image file at the calculated address.

For example:


```

      .PSECT A
B:    .
      .
      .
      .PSECT C
      MOV #B,R0

      or

      .WORD B
    
```

Figure A-24 illustrates the program section relocation entry format.



ZK-467-81

Figure A-24 Program Section Relocation Entry Format

A.4.11 Program Section Displaced Relocation (Type 14)

The program section displaced relocation entry relocates a relative reference to the beginning address of another program section within a module. TKB obtains the current base address of the specified program section; subtracts the address plus 2 that the relocated value is to be written into from the base value; and writes the result into the task image file at the calculated address.


TASK BUILDER INPUT DATA FORMATS

For example:

```

      .PSECT A
B:    .
      .
      .
      .PSECT C
      MOV B,R0
    
```

Figure A-25 illustrates the program section displaced relocation entry format.



ZK-468-81

Figure A-25 Program Section Displaced Relocation Entry Format

A.4.12 Program Section Additive Relocation (Type 15)

The program section additive relocation entry relocates a direct pointer to an address in another program section within a module. TKB obtains the current base address of the specified program section; adds this address to the specified constant; and writes the result into the task image file at the calculated address.

For example:

```

      .PSECT A
B:    .
      .
      .
C:    .
      .
      .
      .PSECT D
      MOV #B+10,R0
      MOV #C,R0

      or

      .WORD B+10
      .WORD C
    
```

Figure A-26 illustrates the program section additive relocation entry format.

TASK BUILDER INPUT DATA FORMATS

DISP	B	ENTRY TYPE = 15
PROGRAM SECTION NAME (2 WORDS)		
CONSTANT		

ZK-469-81

Figure A-26 Program Section Additive Relocation Entry Format

A.4.13 Program Section Additive Displaced Relocation (Type 16)

The program section additive displaced relocation entry relocates a relative reference to an address in another program section within a module. TKB obtains the current base address of the specified program section; adds this address to the specified constant; subtracts the address plus 2 that the relocated value is to be written into from the resultant additive value; and writes the result into the task image file at the calculated address.

For example:

```

        .PSECT A
B:      .
        .
        .
C:      .
        .
        .
        .PSECT D
        MOV B+10,R0
        MOV C,R0
    
```

Figure A-27 illustrates the program section additive displaced relocation entry format.

DISP	B	ENTRY TYPE = 16
PROGRAM SECTION NAME (2 WORDS)		
CONSTANT		

ZK-470-81

Figure A-27 Program Section Additive Displaced Relocation Entry Format

TASK BUILDER INPUT DATA FORMATS

A.4.14 Complex Relocation (Type 17)

The complex relocation entry resolves a complex relocation expression. Such an expression is one in which any of the MACRO-11 binary or unary operations are permitted with any type of argument, regardless of whether the argument is an unresolved global symbol; is relocatable to any program section base; is absolute; or is a complex relocatable subexpression.

The RLD command word is followed by a string of numerically specified operation codes and arguments. The operation codes each occupy one byte. The entire RLD command must fit in a single record. The following 15 operation codes are defined:

- No operation -- Byte 0
- Addition (+) -- Byte 1
- Subtraction (-) -- Byte 2
- Multiplication (*) -- Byte 3
- Division (/) -- Byte 4
- Logical AND (&) -- Byte 5
- Logical inclusive OR (!) -- Byte 6
- Negation (-) -- Byte 10
- Complement (^C) -- Byte 11
- Store result (command termination) -- Byte 12
- Store result with displaced relocation (command termination) -- Byte 13
- Fetch global symbol -- Byte 16 (It is followed by four bytes containing the symbol name in Radix-50 representation.)
- Fetch relocatable value -- Byte 17 (It is followed by one byte containing the program section number, and two bytes containing the offset within the program section.)
- Fetch constant -- Byte 20 (It is followed by two bytes containing the constant.)
- Fetch resident library base address -- Byte 21 (If the file is a resident library STB file, the library base address is obtained; otherwise, the base address of the task image is fetched.)

The STORE commands indicate that the value is to be written into the task image file at the calculated address.

All operands are evaluated as 16-bit signed quantities using two's complement arithmetic. The results are equivalent to expressions that the assembler evaluates internally. The following rules should be noted:

1. An attempt to divide by 0 yields a 0 result. The Task Builder issues a nonfatal diagnostic error message.

TASK BUILDER INPUT DATA FORMATS


2. All results are truncated from the left to fit into 16 bits. No diagnostic error message is issued if the number is too large. If the result modifies a byte, TKB checks for truncation errors as described in Section A.4.
3. All operations are performed on relocated (additive) or absolute 16-bit quantities. PC displacement is applied to the result only.

For example:

```

A:      .PSECT  ALPHA
      .
      .
      .PSECT  BETA
B:      .
      .
      .
      MOV    #A+B-G1/G2&^C177120!G3>>,R1
    
```

Figure A-28 illustrates the complex relocation entry format.



ZK-471-81


Figure A-28 Complex Relocation Entry Format

A.4.15 Resident Library Relocation (Type 20)

The library relocation entry relocates a direct pointer to an address within a resident library.

If the current file is a resident library symbol definition file (STB), TKB obtains the base address of the library; adds this address to the specified constant; and writes the result into the task image file at the calculated address. If the file is not associated with a resident library, TKB uses the task base address.

Figure A-29 illustrates the library relocation entry format.



ZK-472-81

Figure A-29 Resident Library Relocation Entry Format

TASK BUILDER INPUT DATA FORMATS

A.5 INTERNAL SYMBOL DIRECTORY RECORD

Internal symbol directory (ISD) records have two purposes:

1. To pass information to symbolic debuggers via the .STB file
2. To create autoloading vectors dynamically for the entry points of the library

TBK looks for global symbol definitions in the input object modules and looks for ISD records if /DA is specified; otherwise TBK ignores the ISD records. Some ISD records require no relocation and TBK can copy them directly into the .STB file. Others will require modification; after being modified, they can be written to the .STB file. In addition, TBK may need to generate some ISD records of its own in the .STB file.

Except for autoloading library entry points, TBK puts ISD records into the .STB file only if the /DA switch is used in the TBK command line. When TBK outputs the .STB file, it writes three major types of ISD records:


- Type 1 records, TBK generated ISDs. The form of these records is language independent.
- Type 3 records, written for any type 2 records in an input object module. TBK does this after adding data and then changing the type to 3. Type 2 relocatable/relocated records are those that contain both language dependent and independent sections. Language processors generate these records and TBK modifies them. They contain information that can be used to find the absolute task image address of source program entities (variables, program statements, etc.)
- Type 4 records, written to the .STB file without modification. Type 4 records are literal records that contain language dependent information. Apart from the first few bytes, TBK ignores the rest of the record.

These record formats are described in the following sections.

A.5.1 Overall Record Format

ISD records have the same basic structure as all object language records. Because of the variety of different types, the skeleton structure must include additional fields that are common to all ISD record types. The general format of all ISD records is shown in Figure A-30.

TASK BUILDER INPUT DATA FORMATS



ZK-1058-82


Figure A-30 General Format of All ISD Records

ISD record types fall into general categories. The categories are:

- 0 -- Illegal.
- 1 -- TKB-generated.
- 2 -- Compiler-generated relocatable.
- 3 -- Relocated (type 2 after TKB processing).
- 4-127 -- not defined and reserved for future use.
- 128-255 -- literal records; the type code identifies the generating language processor and the internal structure.

A.5.2 TKB Generated Records (Type 1)

The content of this record type is a string of individual items, each with its own format. The items are either start-of-segment items, task identification items, or autoloadable entry point items. The TKB generated record is similar to the structure of an RLD or GSD record. The general format is shown in Figure A-31.

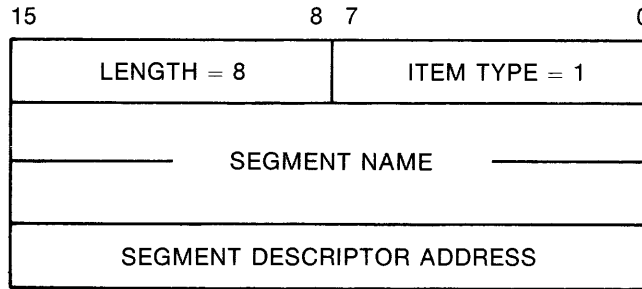


ZK-1059-82

Figure A-31 General Format of a TKB Generated Record

A.5.2.1 Start-of-Segment Item Type (1) - The format of the start-of-segment item type is shown in Figure A-32.

TASK BUILDER INPUT DATA FORMATS



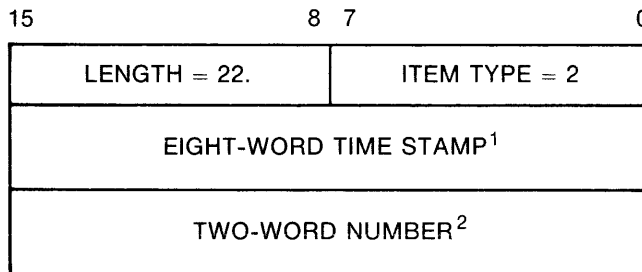
ZK-1060-82

Figure A-32 Format of TKB Generated Start-of-Segment Item (1)

A.5.2.2 Task Identification Item Type (2) - The task identification item type ensures that a .STB file and the task image being debugged were generated at the same time. Otherwise, symbols that are found may not correspond to the actual task.

The task identification item type exists to make the correlation between the .STB file and its related task possible. The contents of this item type correspond exactly to the first ten words of an area in a task image file, which is in the TKB created PSECT called \$\$DBTS.

The format of the task identification item type is shown in Figure A-33.



1. Its form is that which is returned by RSX-11M/M-PLUS directive GTIMS.

2. TKB generates this number as an additional check on correspondence. Currently always zero.


ZK-1061-82

Figure A-33 Format of TKB Generated Task Identification Item (2)

A.5.2.3 Autoloadable Library Entry Point Item Type (3) - TKB outputs the autoloadable library entry point item into a .STB file when building overlaid resident libraries. The ISD record contains the needed information for TKB to dynamically generate autoload vectors for entry points in the library. Autoload vectors appear only for those entry points that are referenced by the task. Unlike the other item types, the autoloadable library entry point item is not for use by debuggers.

TASK BUILDER INPUT DATA FORMATS

The format of the autoloadable entry point item is shown in Figure A-34.



ZK-1062-82

Figure A-34 Format of an Autoloadable Library Entry Point Item (3)

A.5.3 Relocatable/Relocated Records (Type 2)

These records are the central part of TKB's involvement in debugger communication. Every item type in these records must be standardized, and only standard items can appear. The general format of relocatable/relocated records is the same as that shown in Figure A-30.

A language processor outputs these record types as type 2. When TKB processes them, it changes the type to type 3. It also fills in or modifies some fields. In the descriptions of following item types, fields that are filled in by TKB are marked with an asterisk (*). They should be left as zero in language processor output.

A.5.3.1 Module Name Item Type (1) - A module name item should be the first ISD entry of each object module. A debugger can assume that all following ISD information up to the next module name item relates to this module.

The language code is included so that a debugger for a specific language can determine whether to ignore a module if it is written for another language. The language code has the same range of values as that of a language-dependent ISD record (128-255) and has the same meaning.

The format of the module name item type is shown in Figure A-35.

TASK BUILDER INPUT DATA FORMATS

15	8 7	0
LENGTH	ITEM TYPE = 1	
MUST BE 0	LANGUAGE CODE	
MODULE NAME ¹		

1. A counted ASCII string of the required length. A counted ASCII string is a byte string in which the first byte indicates the number of bytes to follow.

ZK-1063-82

Figure A-35 Format of a Module Name Item Type (1)

A.5.3.2 Global Symbol Item Type (2) - One type 2 item must appear for each global symbol definition that the language processor wants the debugger to understand. It need not, however, include definitions generated for the language processor run-time system.

The format of the global symbol item type is shown in Figure A-36.

15	8 7	0
LENGTH	ITEM TYPE = 2	
SYMBOL NAME (RADIX-50)		
VALUE*		
DESCRIPTOR ADDRESS FOR CONTAINING OVERLAY SEGMENT*		
MUST BE ZERO	FLAGS	
FULL SYMBOL NAME ¹		

1. Counted ASCII string of the required length. A counted ASCII string is a byte string in which the first byte indicates the number of bytes to follow.

ZK-1053-82

Figure A-36 Format of a Global Symbol Item Type (2)

TASK BUILDER INPUT DATA FORMATS

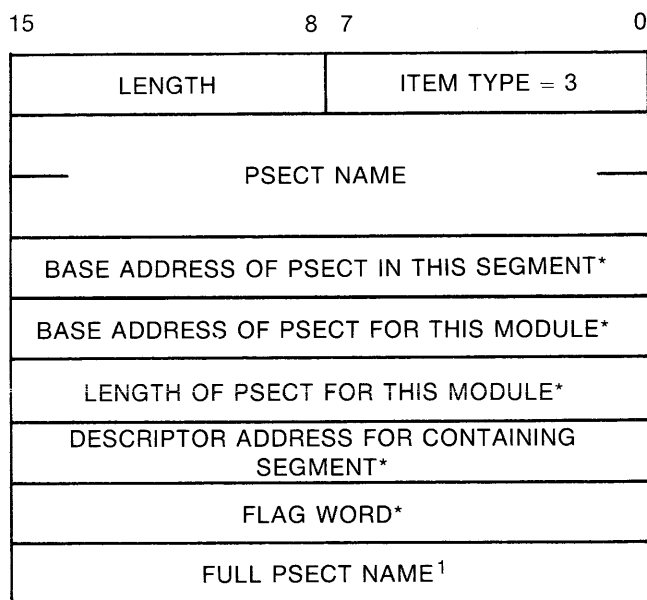
A.5.3.3 PSECT Item Type (3) - A concatenated PSECT has two base addresses: one for the whole PSECT, and another for the part of it that belongs to this module. It is the base address for the part that belongs to this module that may be used by a debugger to convert local symbol values to absolute addresses.

The segment descriptor address is necessary because PSECTs may move to segments other than the one in which it was placed. This address is relevant to languages that provide semi-automatic overlay generation, like COBOL-11. This word may be zero if the PSECT has not moved to another segment.

The flag word is a copy of the flag word built by TKB. It allows for identification of VSECTs.

The full PSECT name may be needed for some languages.

The format of a PSECT item type is shown in Figure A-37.



1. A counted ASCII string of the required length. A counted ASCII string is a byte string in which the first byte indicates the number of bytes to follow.

ZK-1054-82

Figure A-37 Format of a PSECT Item Type (3)

A.5.3.4 Line-Number or PC Correlation Item Type (4) - This item provides the information needed to translate a source line number into a task image address, or a task image address into a source line number.

The format of a line-number or PC correlation item type is shown in Figure A-38.

TASK BUILDER INPUT DATA FORMATS

15	8 7	0
LENGTH	ITEM TYPE = 4	
PSECT		
NAME		
START PC ¹		
DESCRIPTOR ADDRESS OF CONTAINING OVERLAY SEGMENT*		
START PAGE NUMBER		
START LINE NUMBER		
STRING OF ONE-BYTE ITEMS		

1. Offset into PSECT in type 2 records; absolute address in type 3 records.

ZK-1055-82

Figure A-38 Format of a Line-Number or PC Correlation Item Type (4)

A.5.3.5 Internal Symbol Name Item Type (5) - The internal symbol name item allows for the fact that a name may have more than one associated address. For example, a COBOL variable may have three associated addresses: the address of the area that contains the data, the address of a CIS descriptor, and the address of a picture string.

The internal symbol name item is shown in Figure A-39.

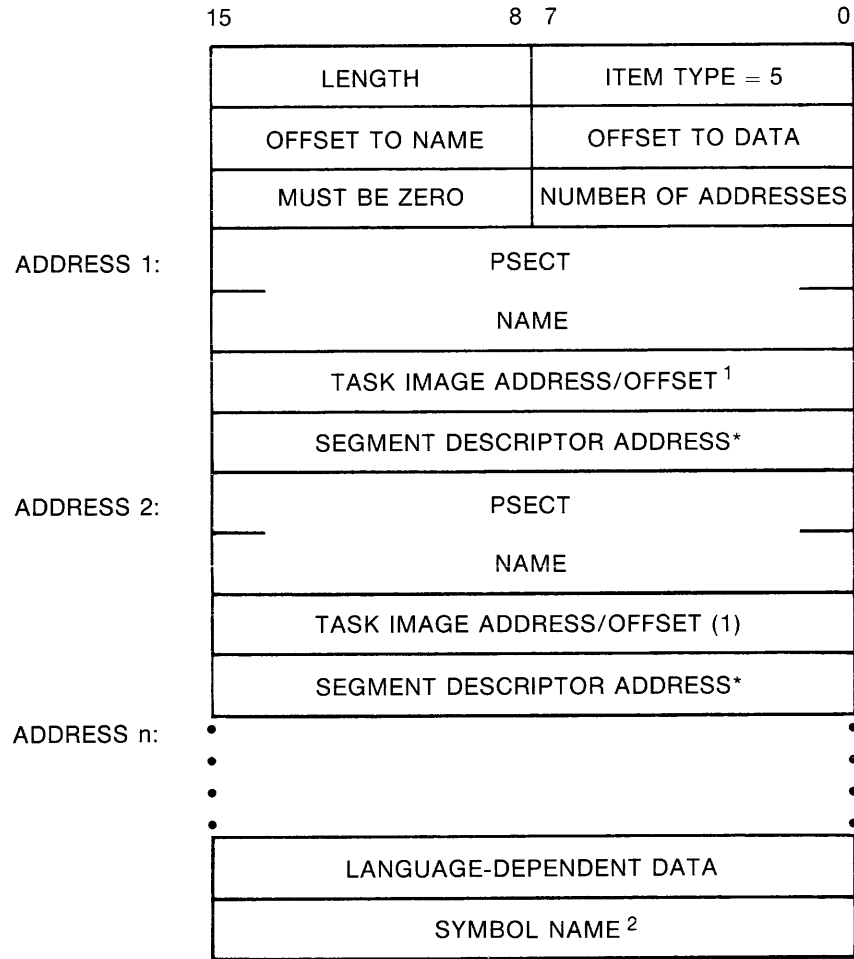
A.5.4 Literal Records (Type 4)

Literal records may take any form, but the two-byte header shown in Figure A-40 must be present.

A.6 END OF MODULE RECORD

The end-of-module record declares the end of an object module. There must be exactly one end-of-module record in every object module. As shown in Figure A-41, this record is one word long.

TASK BUILDER INPUT DATA FORMATS




1. Modified by TKB

2. A counted ASCII string of the required length. A counted ASCII string is a byte string in which the first byte indicates the number of bytes to follow.

ZK-1056-82

Figure A-39 Format of an Internal Symbol Name Item Type (5)



ZK-1057-82

Figure A-40 Format of a Literal Record Type

TASK BUILDER INPUT DATA FORMATS

0	RECORD TYPE = 6
---	--------------------

ZK-473-81

Figure A-41 End-of-Module Record Format

APPENDIX B

DETAILED TASK IMAGE FILE STRUCTURE

Figures B-1 through B-4 illustrate how the Task Builder (TKB) records a task image on disk. As noted in the following sections, parts of the task disk image shown in these figures are optional and may not be recorded for every task image.

The following sections, which provide detailed information on the task image file structure, are organized as follows:

- B.1 Label Block Group
- B.2 Checkpoint Area
- B.3 Header
 - B.3.1 Low-memory Context
 - B.3.2 Logical Unit Table Entry
- B.4 Task Image
 - B.4.1 Autoload Vectors for Conventional Tasks
 - B.4.2 Autoload Vectors for I- and D-Space Tasks
 - B.4.3 Task-Resident Segment Descriptor
 - B.4.4 Window Descriptor
 - B.4.5 Region Descriptor
 - B.4.6 Supervisor-Mode Vector

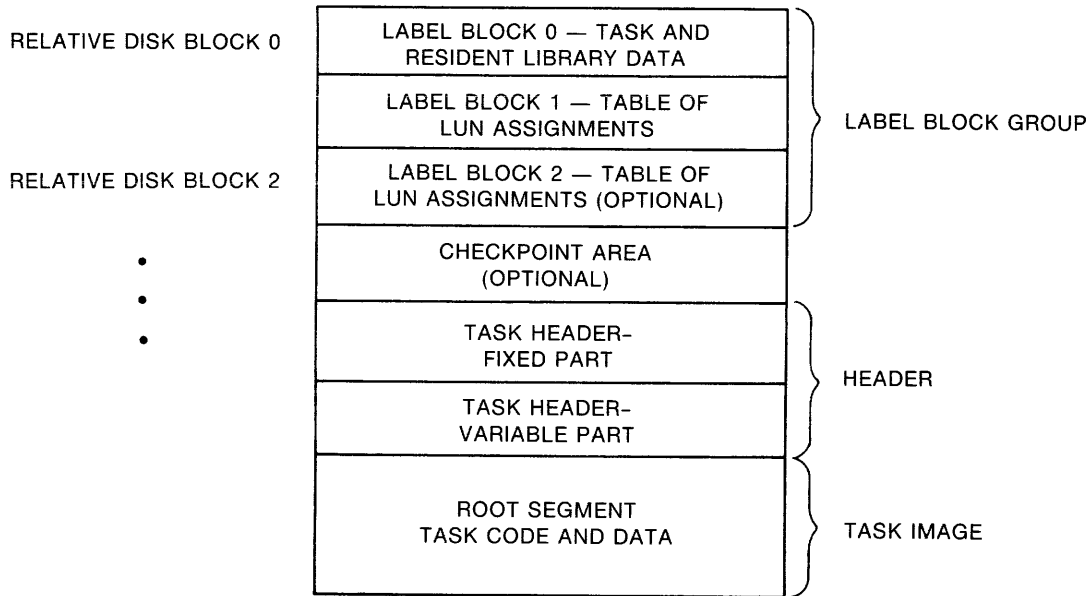
B.1 LABEL BLOCK GROUP

The label block group precedes the task on the disk and contains data that is needed by the system to install and load a task but need not reside in memory during task execution. This group consists of three parts:

- Task and resident library data (label block 0)
- Table of LUN assignments (label blocks 1 and 2)
- The segment load list (label block 3)

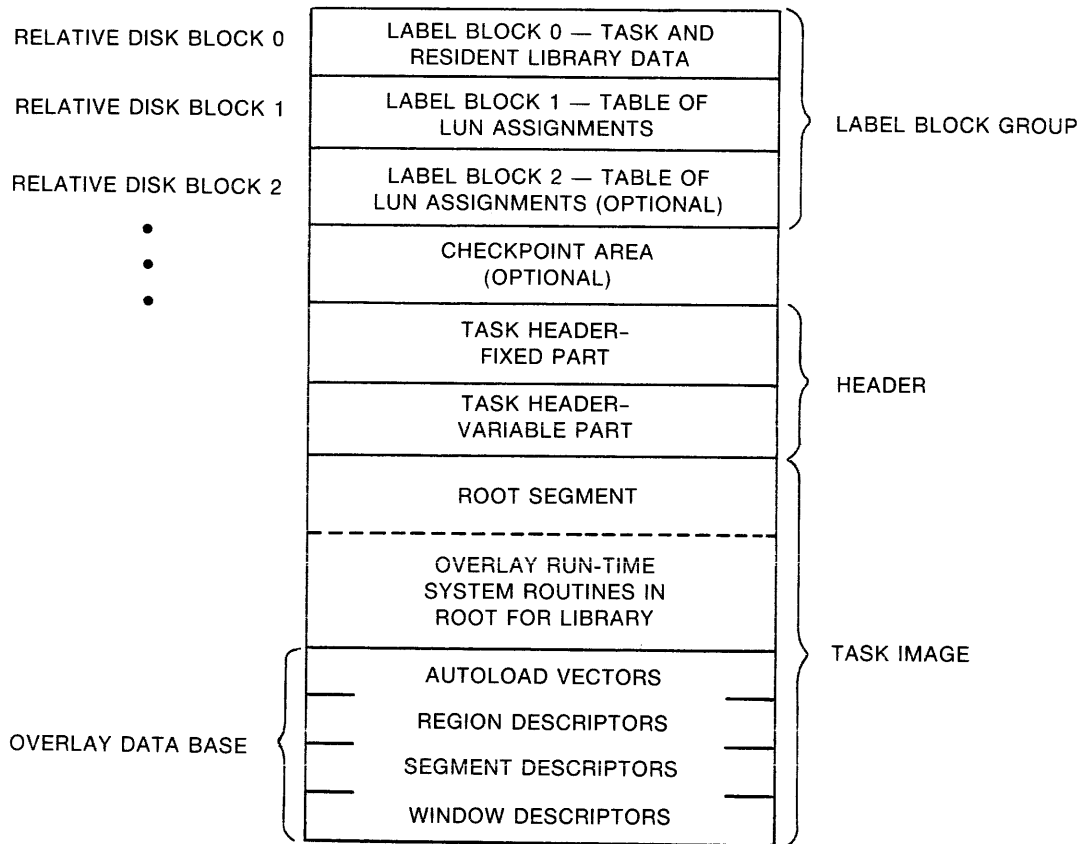
Table B-1 describes the task and resident library data. Figure B-5 illustrates how TKB organizes this data in label block 0. The INSTALL processor verifies the task and resident library data when entering the tasks into the System Task Directory (STD) file. You can obtain the offsets shown in Figure B-5 by calling the LBLDF\$ macro that resides in macro library LB:[1,1] EXEMC.MLB.

DETAILED TASK IMAGE FILE STRUCTURE



ZK-1064-82

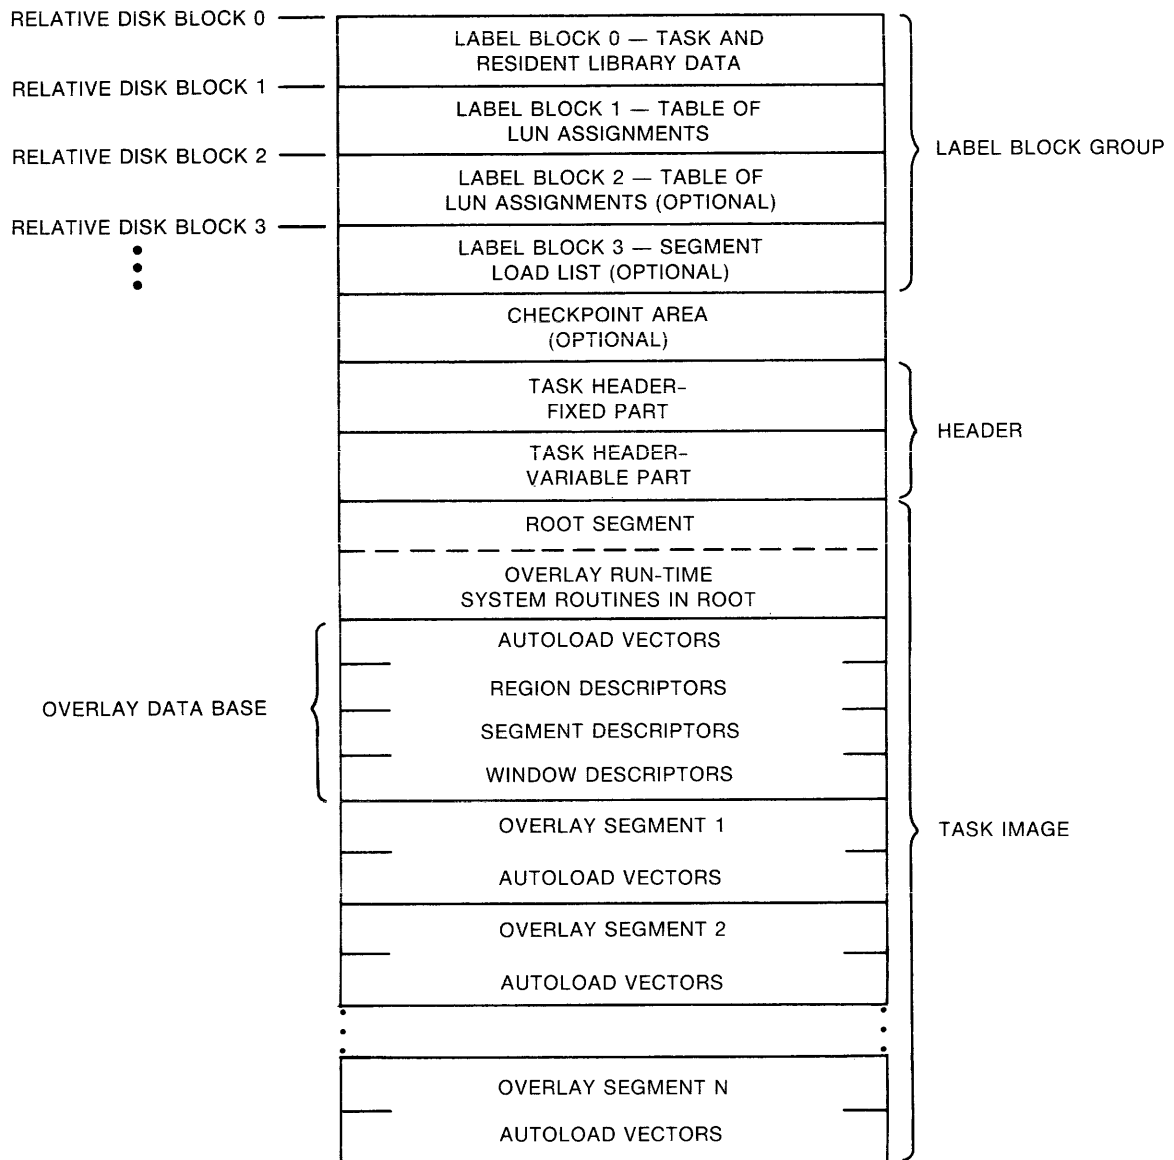
Figure B-1 Image on Disk of Non-Overlaid Conventional Task



ZK-1065-82

Figure B-2 Image on Disk of Conventional Non-Overlaid Task Linked to Overlaid Library

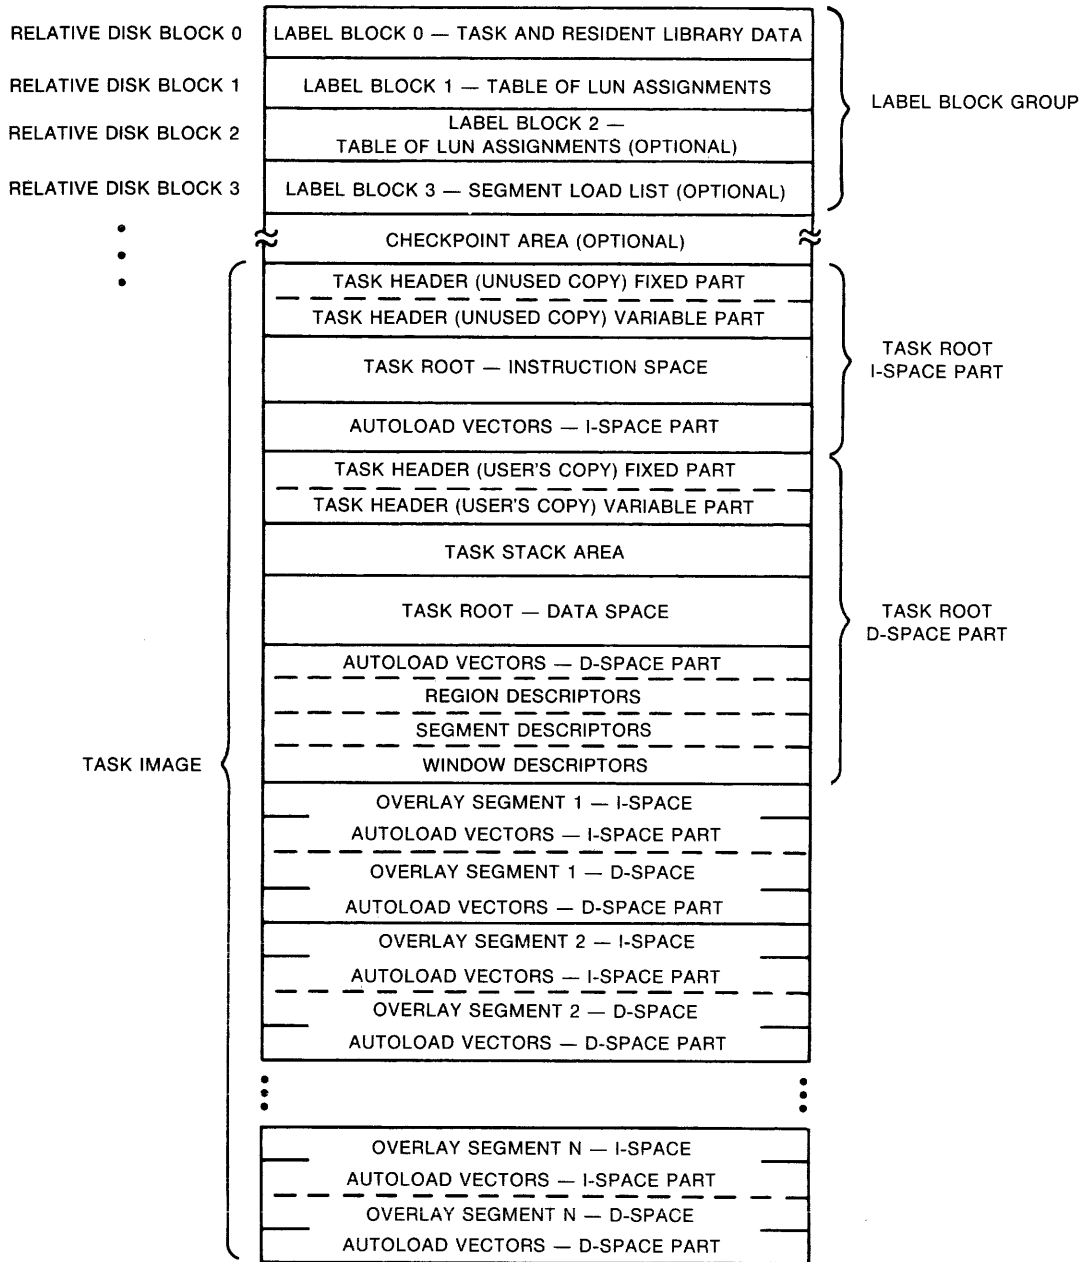
DETAILED TASK IMAGE FILE STRUCTURE



ZK-1066-82

Figure B-3 Image on Disk of Conventional Overlaid Task

DETAILED TASK IMAGE FILE STRUCTURE



ZK-1067-82

Figure B-4 Image on Disk of Overlaid I- and D-Space Task

Table B-1
Task and Resident Library Data

Parameter	Definition
L\$BTSK	Task name consisting of two words in Radix-50 format. This parameter is set by the TASK keyword.
L\$BPAR	Partition name consisting of two words in Radix-50 format. This parameter is set by the PAR keyword.

(continued on next page)

DETAILED TASK IMAGE FILE STRUCTURE

Table B-1 (Cont.)
Task and Resident Library Data

Parameter	Definition																																				
L\$BSA	Starting address of task. Marks the lowest task virtual address. This parameter is set by the PAR keyword.																																				
L\$BHG	Highest virtual address mapped by address window 0.																																				
L\$BMXV	Highest task virtual address. When the task does not have memory-resident overlays, the value is set to L\$BHG.																																				
L\$BLDZ	Task load size in units of 64-byte blocks. This value represents the size of the root segment.																																				
L\$BMXZ	Task maximum size in units of 64-byte blocks. This value represents the size of the root segment plus any additional physical memory needed to contain task overlays.																																				
L\$BOFF	Task offset into partition in units of 64-byte blocks. This value represents the size of the mapped array area, which precedes the task's code and data in the partition.																																				
L\$BWND	Number of task window blocks less library window blocks -- Low byte																																				
L\$BSYS	System ID -- High byte																																				
L\$BWND	Number of task windows (excluding resident libraries).																																				
L\$BSEG	Size of overlay segment descriptors (in bytes).																																				
L\$BFLG	Task flags word. The following flags are defined:																																				
	<table border="1"> <thead> <tr> <th>Bit</th> <th>Flag</th> <th>Meaning When Bit = 1</th> </tr> </thead> <tbody> <tr> <td>15</td> <td>TS\$PIC</td> <td>Task contains position-independent code (PIC).</td> </tr> <tr> <td>14</td> <td>TS\$NHD</td> <td>Task has no header.</td> </tr> <tr> <td>13</td> <td>TS\$ACP</td> <td>Task is ancillary control processor.</td> </tr> <tr> <td>12</td> <td>TS\$PMD</td> <td>Task generates Postmortem Dump.</td> </tr> <tr> <td>11</td> <td>TS\$SLV</td> <td>Task can be slaved.</td> </tr> <tr> <td>10</td> <td>TS\$NSD</td> <td>No SEND can be directed to task.</td> </tr> <tr> <td>9</td> <td></td> <td>(Not used)</td> </tr> <tr> <td>8</td> <td>TS\$PRV</td> <td>Task is privileged.</td> </tr> <tr> <td>7</td> <td>TS\$CMP</td> <td>Task is built in compatibility mode.</td> </tr> <tr> <td>6</td> <td>TS\$CHK</td> <td>Task is not checkpointable.</td> </tr> <tr> <td>5</td> <td>TS\$RES</td> <td>Task has memory-resident overlays.</td> </tr> </tbody> </table>	Bit	Flag	Meaning When Bit = 1	15	TS\$PIC	Task contains position-independent code (PIC).	14	TS\$NHD	Task has no header.	13	TS\$ACP	Task is ancillary control processor.	12	TS\$PMD	Task generates Postmortem Dump.	11	TS\$SLV	Task can be slaved.	10	TS\$NSD	No SEND can be directed to task.	9		(Not used)	8	TS\$PRV	Task is privileged.	7	TS\$CMP	Task is built in compatibility mode.	6	TS\$CHK	Task is not checkpointable.	5	TS\$RES	Task has memory-resident overlays.
Bit	Flag	Meaning When Bit = 1																																			
15	TS\$PIC	Task contains position-independent code (PIC).																																			
14	TS\$NHD	Task has no header.																																			
13	TS\$ACP	Task is ancillary control processor.																																			
12	TS\$PMD	Task generates Postmortem Dump.																																			
11	TS\$SLV	Task can be slaved.																																			
10	TS\$NSD	No SEND can be directed to task.																																			
9		(Not used)																																			
8	TS\$PRV	Task is privileged.																																			
7	TS\$CMP	Task is built in compatibility mode.																																			
6	TS\$CHK	Task is not checkpointable.																																			
5	TS\$RES	Task has memory-resident overlays.																																			

(continued on next page)

DETAILED TASK IMAGE FILE STRUCTURE

Table B-1 (Cont.)
Task and Resident Library Data

Parameter	Definition		
	Bit	Flag	Meaning When Bit = 1
L\$BFLG (Cont.)	4	TS\$IOP	Privileged task does not map I/O page
	3	TS\$SUP	Image linked as supervisor-mode library (RSX-11M-PLUS systems only).
	2	TS\$XHR	Task was built with external header
	1	TS\$NXH	Task was built with pool resident header (non external)
L\$BDAT	Three words containing the task creation date as 2-digit integer values as follows:		
	<ul style="list-style-type: none"> ● Year since 1900 ● Month of year ● Day of month 		
L\$BLIB	Resident library entries.		
L\$BPRI	Task priority set by the PRI keyword.		
L\$BXFR	Task transfer address. Used to initiate a bootable core image, for example, the resident executive.		
L\$BEXT	Task extension size in units of 32-word blocks. This parameter is set by the EXTTSK keyword.		
L\$BSGL	Relative block number of segment load list. Set to 0 if no list is allocated.		
L\$BHRB	Relative block number of header.		
L\$BBLK	Number of blocks in label block group.		
L\$BLUN	Number of logical units.		
L\$BROB	Relative block number of R/O image.		
L\$BROL	R/O load size in 32-word blocks.		
L\$BRDL	Size of R/O data in 32-word blocks.		
L\$BHDB	Relative block number of data header.		
L\$BDHV	High virtual address of data window 1.		
L\$BDMV	High virtual address of data.		
L\$BDLZ	Load size of data		
L\$BDMZ	Maximum size of data		

DETAILED TASK IMAGE FILE STRUCTURE

Label	Offset		
L\$BTSK	0	Task	
	2	Name	
L\$BPAR	4	Task	
	6	Partition	
L\$BSA	10	Base address of task	
L\$BHG	12	Highest window 0 virtual address	
L\$BMXV	14	Highest virtual address in task	
L\$BLDZ	16	Load size in 64-byte blocks	
L\$BMXZ	20	Maximum size in 64-byte blocks	
L\$BOFF	22	Task offset into partition	
L\$BWND/L\$BSYS	24	System I.D. Number of window blocks*	
L\$BSEG	26	Size of overlay segment descriptors	
L\$BFLG	30	Task flag word	
L\$BDAT	32	Task creation date – Year	
	34	– Month	
	36	– Day	
L\$BLIB	40	Library/common	
	42	Name	R\$LNAM
	44	Base address of library	R\$LSA
	46	Highest address in first library window	R\$LHG
	50	Highest address in library	R\$LMXV
	52	Library load size (64-byte blocks)	R\$LLDZ
	54	Library maximum size (64-byte blocks)	R\$LMXZ
	56	Library offset into region	R\$LOFF
	60	Number of library window blocks	R\$LWND
	62	Size of library segment descriptors	R\$LSEG
	64	Library flag word	R\$LFLG
	66	Library creation date – Year	R\$LDAT
	70	– Month	
	72	– Day	
⋮	⋮		
344	0		
L\$BPRI	346	Task priority	
L\$BXFR	350	Task transfer address	
L\$BEXT	352	Task extension (64-byte blocks)	
L\$BSGL	354	Block number of segment load list	
L\$BHRB	356	Block number of header	
L\$BBLK	360	Number of blocks in label	
L\$BLUN	362	Number of logical units	
L\$BROB	364	Relative block of R-O image	
L\$BROL	366	R/O load size	
L\$BRDL	370	R/O data size in 32-word blocks	
L\$BHDB	372	Relative block number of data header	
L\$BDHV	374	High virtual address of data window 1	
L\$BDMV	376	High virtual address of data	
L\$BDLZ	400	Load size of data	
L\$BDMZ	402	Maximum size of data	
⋮	⋮		
	0		

Library Request (maximum of 7 14-word entries in RSX-11M systems and maximum of 15 14-word entries in RSX-11M+ systems)

*Less library window blocks.

Figure B-5 Label Block 0 -- Task and Resident Library Data

Table B-2 describes the contents of the resident shared region name block. TKB constructs this block by referring to the disk image of the resident shared region. The format is identical to words 3 through 16 of the label group block.

DETAILED TASK IMAGE FILE STRUCTURE

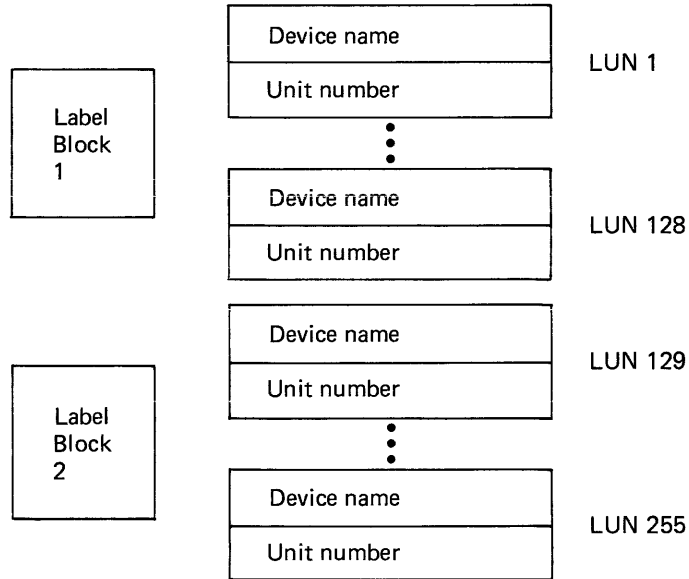
Table B-2
Resident Library/Common Name Block Data

Parameter	Definition														
R\$LNAM	Shared region name consisting of 2 words in Radix-50 format.														
R\$LSA	Base virtual address of library or common.														
R\$LHGV	Highest address mapped by first library window.														
R\$LMXV	Highest virtual address in library or common.														
R\$LLDZ	Shared region load size in 64-byte blocks.														
R\$LMXZ	Library maximum size in 64-byte blocks. This value represents the size of the root segment plus the sum of all memory-resident overlays.														
R\$LOFF	Size of mapped array space allocated by resident library. This value is added to the mapped array area of the task.														
R\$LWND	Number of window blocks required by library.														
R\$LSEG	Size of library overlay segment descriptors in bytes.														
R\$LFLG	Library flags word. The following flags are defined: <table border="1" data-bbox="474 1024 1230 1459"> <thead> <tr> <th>Bit</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>15</td> <td>LD\$ACC -- Access intent (1=read/write, 0=read-only)</td> </tr> <tr> <td>14</td> <td>LD\$RSV -- APR was reserved</td> </tr> <tr> <td>13</td> <td>LD\$CLS -- Library is part of a cluster</td> </tr> <tr> <td>3</td> <td>LD\$SUP -- Supervisor-mode library (1=yes)</td> </tr> <tr> <td>2</td> <td>LD\$REL -- Position-independent code (PIC) flag (1=PIC)</td> </tr> <tr> <td>1</td> <td>LD\$TYP -- Shared region type (1 = common, 0 = library)</td> </tr> </tbody> </table>	Bit	Meaning	15	LD\$ACC -- Access intent (1=read/write, 0=read-only)	14	LD\$RSV -- APR was reserved	13	LD\$CLS -- Library is part of a cluster	3	LD\$SUP -- Supervisor-mode library (1=yes)	2	LD\$REL -- Position-independent code (PIC) flag (1=PIC)	1	LD\$TYP -- Shared region type (1 = common, 0 = library)
Bit	Meaning														
15	LD\$ACC -- Access intent (1=read/write, 0=read-only)														
14	LD\$RSV -- APR was reserved														
13	LD\$CLS -- Library is part of a cluster														
3	LD\$SUP -- Supervisor-mode library (1=yes)														
2	LD\$REL -- Position-independent code (PIC) flag (1=PIC)														
1	LD\$TYP -- Shared region type (1 = common, 0 = library)														
R\$LDAT	Three words containing the shared region creation date in 2-digit integer values as follows: <ul style="list-style-type: none"> • Year since 1900 • Month of year • Day of month 														

The table of LUN assignments, illustrated in Figure B-6, contains the name and logical unit number of each device assigned. Label block 2 (the second block of LUN assignments) is allocated only if the number of LUNs exceeds 128.


DETAILED TASK IMAGE FILE STRUCTURE

TKB creates the segment load list if the image contains only memory-resident overlays. The segment load list is used only in RSX-11S systems for loading tasks that have resident overlays. Figure B-7 illustrates the segment load list. Each entry in the list gives the length, in bytes, of a memory-resident overlay segment.



ZK-476-81

Figure B-6 Label Blocks 1 and 2 -- Table of LUN Assignments



ZK-477-81

Figure B-7 Label Block 3 -- Segment Load List

B.2 CHECKPOINT AREA

The checkpoint area is created by the /AL switch (refer to Chapter 10). The checkpoint area is as large as the task image plus any areas created by the EXTTSK, PAR, or VSECT options. The checkpoint area does not include space for the external header if the /XH switch was specified.

DETAILED TASK IMAGE FILE STRUCTURE

B.3 HEADER

As shown in Figures B-1 through B-4, the task header starts on a block boundary and is immediately followed by the task image. The header is read into memory with the task image.

The header is divided into two parts: a fixed part as shown in Figure B-8; and a variable part as shown in Figure B-9. The offsets for the fixed part are defined by macro HDRDF\$ residing in LB:[1,1]EXEMC.MLB.

The variable part of the header contains window blocks that describe the following:

- The task's virtual-to-physical mapping
- Logical unit data
- Task context

Although the header is fully accessible to the task, you should consider only the information in the low-memory context (H.DSW through H.VEXT) in the fixed part of the header to be accurate. In a mapped system, the Executive copies the header of an active task to protected memory. Subsequent Executive updates to the header are made to this copy, not to the header copy within the running task.

The following sections provide more detail on the low-memory context and on Logical Unit Table entries (the Logical Unit Table is part of the variable part of the header; see Figure B-9).

NOTE

To save the identification, you should move the initial value set by the Task Builder to local storage. When the program is fixed in memory and being restarted without being reloaded, you must test the reserved program words for their initial values to determine whether the contents of R3 and R4 should be saved.

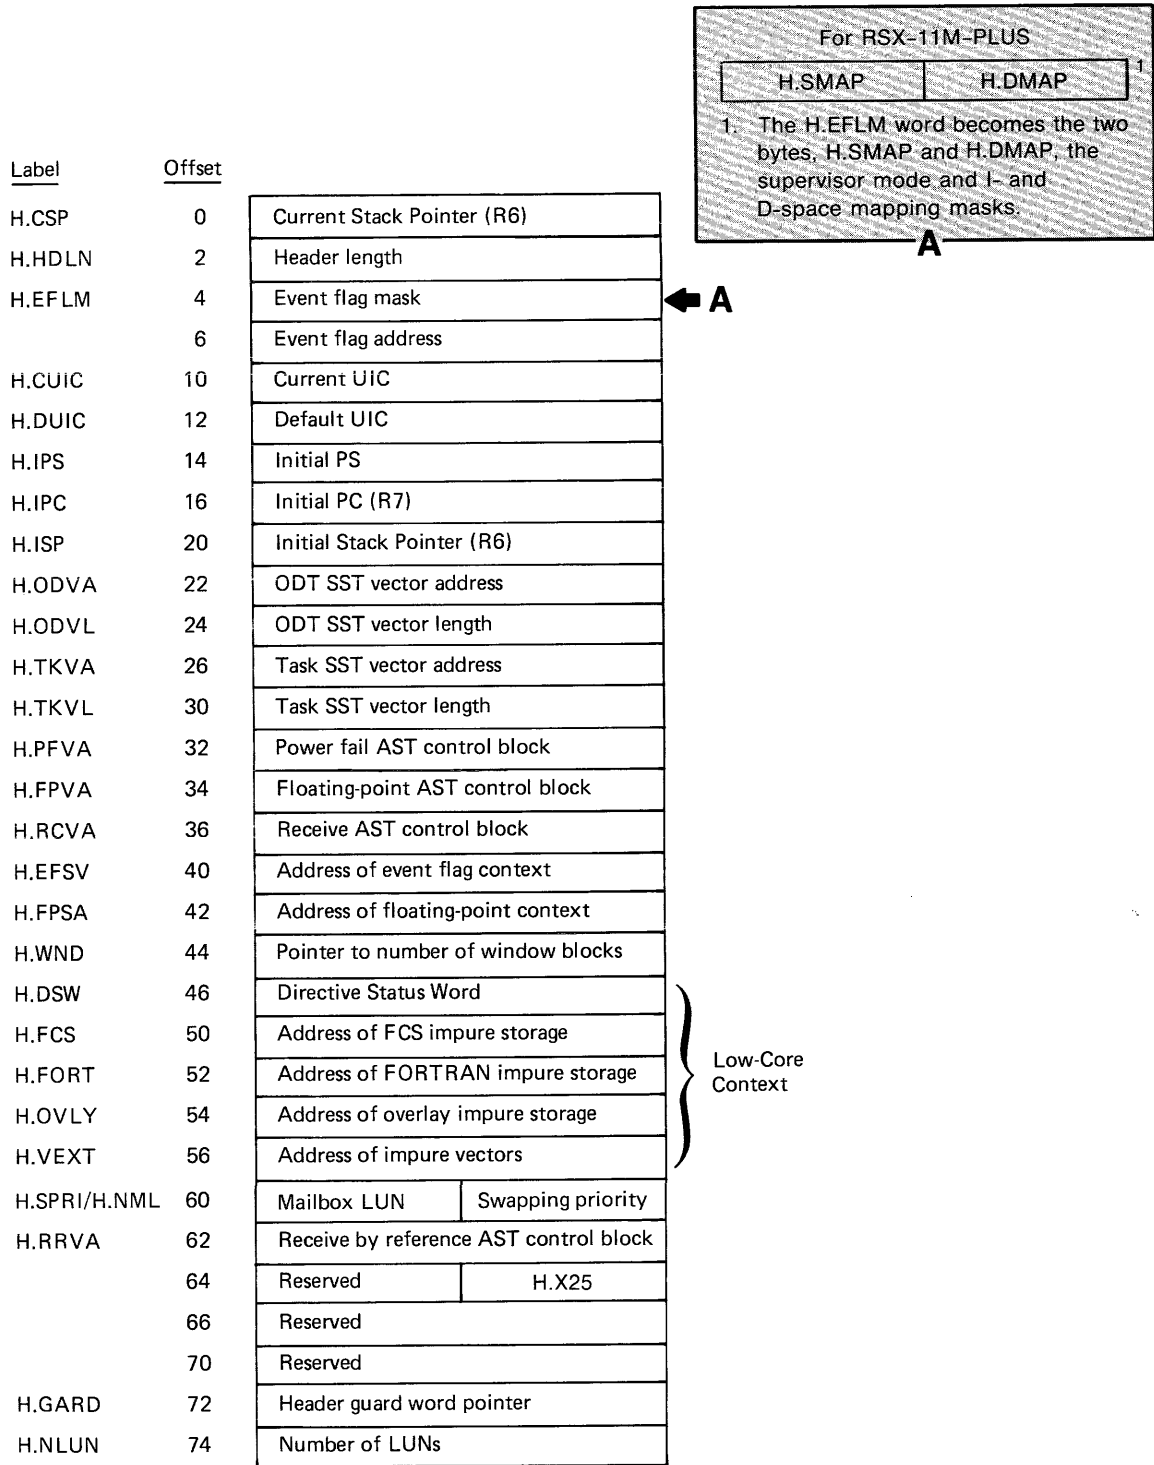
The contents of R0, R1, and R2 are only set when a debugging aid is included in the task image.

B.3.1 Low-Memory Context

The low-memory context for a task consists of the Directive Status Word and the impure area vectors. TKB recognizes the following global names:

Name	Meaning
.FSRPT	File Control Services work area and buffer pool vector
\$OTSV	FORTTRAN OTS work area vector
N.OVPT	Overlay run-time system work area vector
\$VEXT	Vector extension area pointer

DETAILED TASK IMAGE FILE STRUCTURE



ZK-478-81

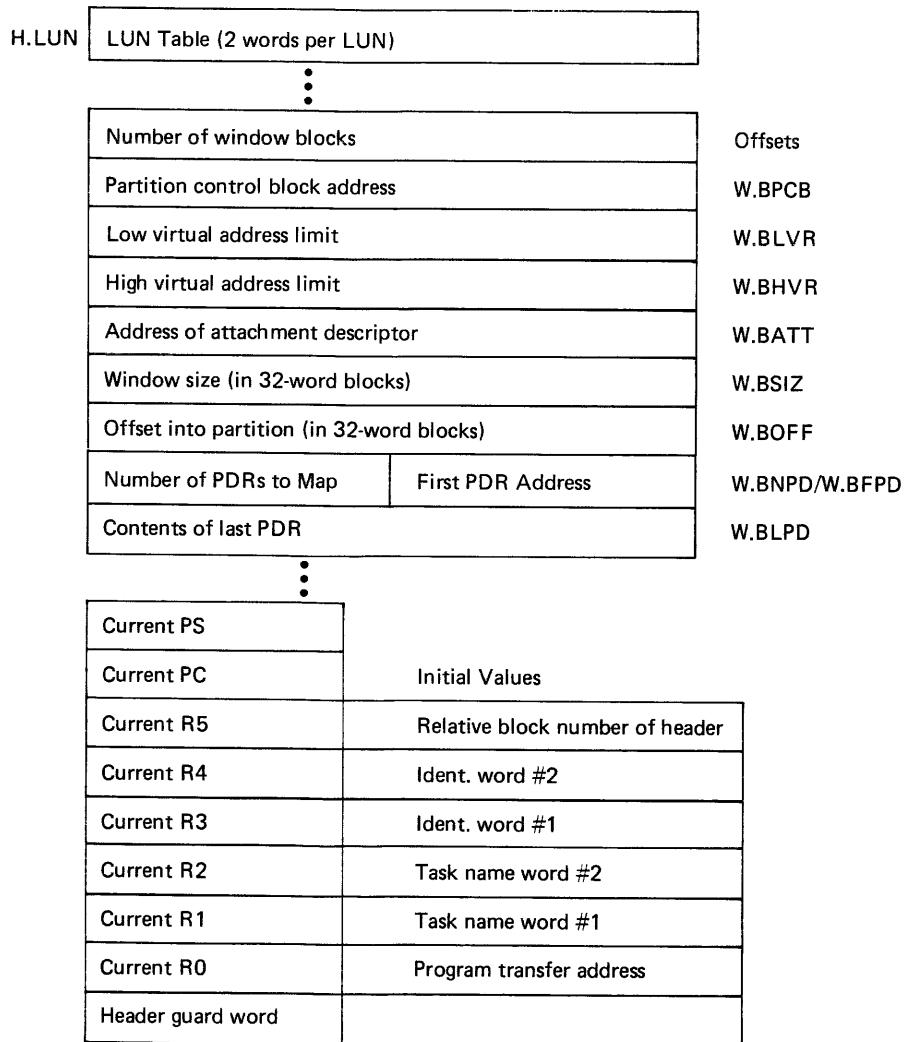
Figure B-8 Task Header, Fixed Part

The only proper reference to these pointers is by symbolic name. The pointers are read-only. If you write into them, the result will be lost on the next context switch.

The impure area pointers contain the addresses of the storage used by the reentrant library routines listed above.

DETAILED TASK IMAGE FILE STRUCTURE

The address contained in the vector extension pointer locates an area of memory that can contain additional impure area pointers.




ZK-479-81

Figure B-9 Task Header, Variable Part

Figure B-10 illustrates the format of the vector extension area. Each location within this area contains the address of an impure storage area that can be referred to by subroutines within a resident library; these subroutines must be reentrant. The address of this area (location \$VEXTA) is contained at absolute address \$VEXT in the task header. Addresses below \$VEXTA, referred to by negative offsets, are reserved for DIGITAL applications. Addresses above \$VEXTA, referred to by positive offsets, are allocated for user applications.

DETAILED TASK IMAGE FILE STRUCTURE



ZK-480-81

Figure B-10 Vector Extension Area Format

The program sections \$\$VEX0 and \$\$VEX1 have the attributes D, GBL, RW, REL, and OVR.

The program section attribute OVR facilitates defining the offset to the vector and initializing the vector location at link time. For example:

```

        .GLOBL $VEXTA      ; MAKE SURE VECTOR AREA IS LINKED
        .PSECT  $$VEX1,D,GBL,REL,OVR

$$$=.          ; POINT TO BASE OF POINTER TABLE

        .BLKW  N          ; OFFSET TO CORRECT LOCATION
                          ; IN VECTOR AREA

LABEL:        .WORD  IMPURE ; SET IMPURE AREA ADDRESS
OFFSET==LABEL-BEG ; DEFINE OFFSET

        .PSECT

IMPURE:

        .
        .
        .
    
```

You should centralize all offset definitions within a single module from which the actual vector space allocation is made. Also, you should write the source code with conditional statements to create two object modules: one that reserves the vector storage; and one that defines the global offsets that will be referred to by your resident library's subroutines.

Note that the sequence of instructions above intentionally redefines the global symbol. The Task Builder reports an error if this value differs from the centralized definition.


DETAILED TASK IMAGE FILE STRUCTURE

You can locate your vector through a sequence of instructions similar to the following:

```
MOV @#VEXT,R0          ; GET ADDRESS OF VECTOR EXTENSIONS
MOV OFFSET(R0),R0      ; POINT TO IMPURE AREA
.END
```

B.3.2 Logical Unit Table Entry

Figure B-11 illustrates the format of each entry in the Logical Unit Table.



ZK-481-81

Figure B-11 Logical Unit Table Entry

The first word contains the address of the device unit control block in the Executive system tables. That block contains device-dependent information.

The second word is a pointer to the window block if the device is file structured.

The UCB address is set during task installation if a corresponding ASG parameter is specified at task-build time. You can also set this word at run time with the Assign LUN Directive to the Executive.

The window block pointer is set when a file is opened on the device whose UCB address is specified by word 1. The window block pointer is cleared when the file is closed.

B.4 TASK IMAGE

The system reads the task image into memory beginning with the task header (see Figures B-1 through B-4). The root segment of a conventional task image is a set of contiguous disk blocks; it is therefore loaded with a single disk access. However, an I- and D-space task root contains 2 sets of contiguous blocks. Therefore, an I- and D-space task requires two disk accesses, one for the D-space part and one for the I-space part. The D-space part is loaded first. Additionally, each segment of an overlaid I- and D-space task requires two disk accesses if it contains both I- and D-space.

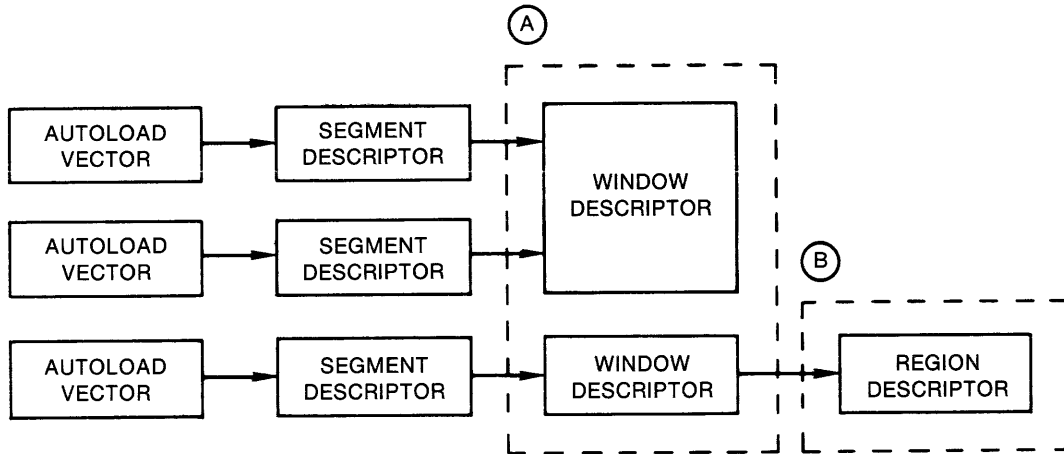
Each overlay segment of the task image begins on a block boundary (see Figure B-3). Note that a given overlay segment occupies as many contiguous disk blocks as it needs to supply its space request. The maximum size for any segment, including the root, is 32K minus 32 words.

DETAILED TASK IMAGE FILE STRUCTURE

NOTE

One exception to the block boundary alignment of segments occurs when shared regions contain resident overlays. When this occurs, the image is compressed and, instead of being aligned on block boundaries, segments are aligned on 32-word boundaries. This facilitates the loading of regions.

Figures B-12 and B-13 illustrate the structure and principal components of the task-resident overlay data base.



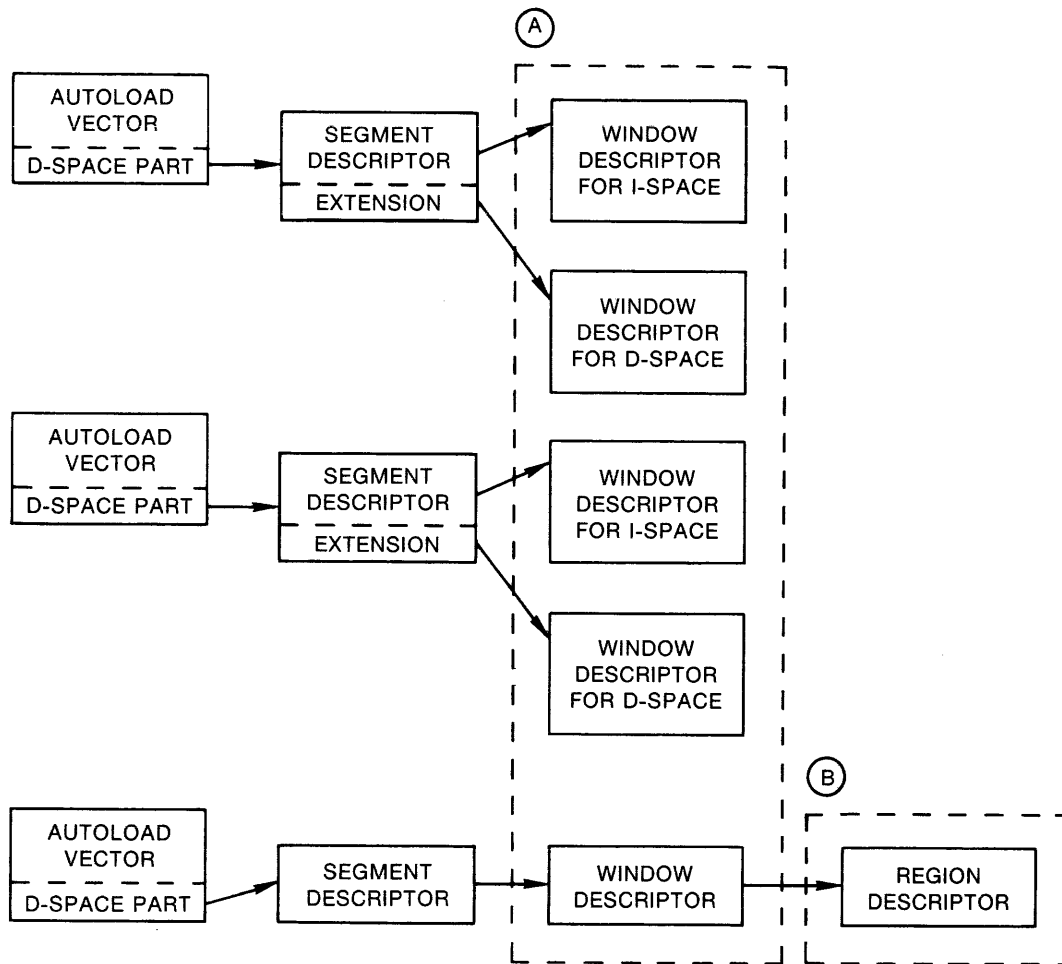
(A) Window descriptors are necessary for the windows that the Overlay Run-Time System uses to map memory resident overlays. The Overlay Run-Time System also needs window descriptors to map disk-resident overlays that are up-tree from memory-resident overlay segments.

(B) The Overlay Run-Time System uses region descriptors to map overlaid libraries.

ZK-1068-82

Figure B-12 Task-Resident Overlay Data Base
for a Conventional Overlaid Task

DETAILED TASK IMAGE FILE STRUCTURE



(A) Window descriptors are necessary for the windows that the Overlay Run-Time System uses to map memory resident overlays. The Overlay Run-Time System also needs window descriptors to map disk-resident overlays that are up-tree from memory-resident overlay segments.

(B) The Overlay Run-Time System uses region descriptors to map overlaid libraries.

ZK-1069-82

Figure B-13 Task-Resident Overlay Data Base for an I- and D-Space Overlaid Task

Autoload vectors are generated whenever a reference is made to an autoloadable entry point in a segment located farther away from the root than the segment making the reference.

One segment descriptor is generated for each overlay segment in the task or shared region. The segment descriptor contains information on the size, virtual address, and location of the segment within the task image file. In addition, it contains a set of link words that point to other segments. The overlay structure determines the link word contents.

DETAILED TASK IMAGE FILE STRUCTURE

Segment descriptors for I- and D-space tasks have an extension for the D-space part that contains the disk block address, virtual load address, segment length in bytes, and window pointer.


The window descriptor contains information required to issue the mapping directives. TKB allocates one window descriptor for each memory-resident overlay in the structure.

The region descriptor contains information required to attach a resident library or common block. There is one region descriptor for each shared region containing memory-resident overlays.

The following sections describe each data base component in greater detail.

B.4.1 Autoload Vectors for Conventional Tasks

The autoload vector table consists of one entry (put into the task image for each autoload entry point) in the form shown in Figure B-14.



ZK-1070-82

Figure B-14 Autoload Vector Entry for Conventional Tasks

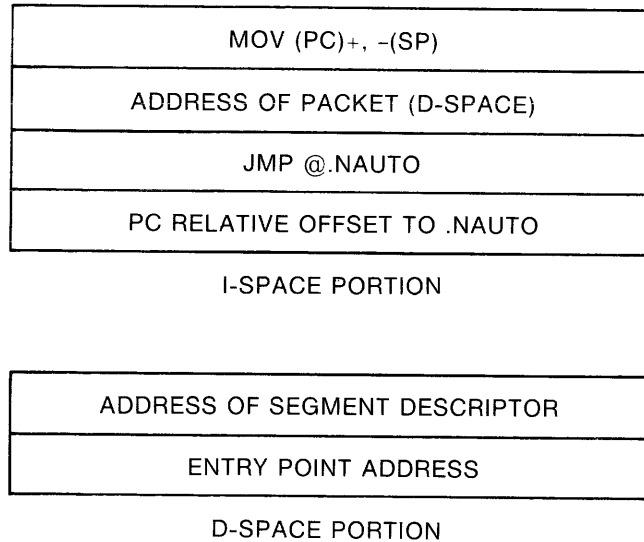
The autoload vector executes an indirect JSR instruction to \$AUTO through .NAUTO. Following the JSR instruction is a pointer to the descriptor for the segment to be loaded. Following the descriptor is the real address of the required entry point.

B.4.2 Autoload Vectors for I- and D-Space Tasks

The autoload vector table consists of two entries (put into the task image for each autoload entry point) in the form shown in Figure B-15.

The I-space part of the autoload vector contains a MOV instruction that places the address of the D-space part of the vector on the stack. The vector then executes an indirect JMP to \$AUTO through .NAUTO. The D-space part of the vector contains the segment descriptor address and the entry point address of the required routine.

DETAILED TASK IMAGE FILE STRUCTURE



ZK-1071-82

Figure B-15 Autoload Vector Entry for I- and D-Space Tasks

B.4.3 Segment Descriptor

The segment descriptor for a conventional task consists of a fixed part and two optional parts. The fixed part is six words long. If the manual-load feature is used (\$LOAD), two words are added containing the segment name. When a memory-resident overlay structure is included, a ninth word is appended that points to the window descriptor.

The segment descriptor for an I- and D-space task consists of a fixed part that is 9 words long and an optional part that is 4 words long. This optional part is always present for task segments and not present for library segments.

Figure B-16 illustrates the contents of the segment descriptor.

DETAILED TASK IMAGE FILE STRUCTURE

TASK-RESIDENT SEGMENT DESCRIPTOR OFFSETS

15	12 11	0	BYTE
FLAGS	RELATIVE DISK BLOCK ADDRESS		0
VIRTUAL LOAD ADDRESS OF SEGMENT		F	2
LENGTH OF SEGMENT IN BYTES			4
LINK UP			6
LINK DOWN			10
LINK NEXT			12
SEGMENT NAME (2-WORD RADIX 50)			14
WINDOW DESCRIPTOR ADDRESS			20

FLAGS: 15-TASK RESIDENT FLAG (ALWAYS 1)
 14-SEGMENT HAS DISK ALLOCATION (1=NO)
 13-SEGMENT IS LOADED FROM DISK (1=YES)
 12-SEGMENT IS LOADED AND MAPPED (0=YES)

F: 0-SEGMENT FOR I- AND D-SPACE TASK (1=YES)

TASK-RESIDENT SEGMENT DESCRIPTOR EXTENSION OFFSETS FOR I- AND D-SPACE TASKS ONLY

15	12 11	0	BYTE
UNUSED	D-SPACE DISK BLOCK ADDRESS		0
D-SPACE VIRTUAL LOAD ADDRESS			2
D-SPACE SEGMENT LENGTH IN BYTES*			4
D-SPACE WINDOW DESCRIPTOR ADDRESS			6

*ZERO IF ONLY I-SPACE SEGMENT

ZK-1072-82

Figure B-16 Segment Descriptor

Word 0 contains the relative disk address in bits 0 through 11 and the segment status in bits 12 through 15. Each segment in the task image file begins on a disk block boundary. The relative disk address is the block number of the segment relative to the start of the root segment. The segment status flags are defined as follows:

Bit	Setting
15	Always set to 1
14	0 = Segment has disk allocation. 1 = Segment does not have disk allocation.
13	0 = Segment is not loaded from disk. 1 = Segment is loaded from disk.
12	0 = Segment is loaded and mapped. 1 = Segment is either not loaded or not mapped.

DETAILED TASK IMAGE FILE STRUCTURE


Word 1 contains the load address of the segment. This address is the first virtual address of the area where the segment will be loaded.

Word 2 specifies the length of the segment in bytes.


Words 3, 4, and 5 point to the following segment descriptors:

- Link up -- The next segment away from the root (0=none).
- Link down -- The next segment toward the root (0=none).
- Link next -- The adjoining segment; the link-next pointers are linked in circular fashion.

When the system loads a segment, the overlay run-time system follows the links to determine which segments are being overlaid and should therefore be marked out of memory. For example:



The segment descriptors are linked as follows:



If there is a co-tree, the link next for the root segment descriptor points to the co-tree root segment descriptor.

Words 6 and 7 contain the segment name in Radix-50 format.

Word 8 points to the window descriptor used to map the segment (0=none).

B.4.4 Window Descriptor

TKB allocates window descriptors only if you define a structure containing memory-resident overlays. Figure B-17 illustrates the format of a window descriptor.

Words 0 through 7 constitute a window descriptor in the format required by the mapping directives. If the memory-resident overlay is part of the task, the region ID is 0. If the memory-resident overlay is part of a shared region, the overlay loading routine fills in the ID at run time.

Words 8 and 9 contain additional data that is referred to by the overlay routines. Bit 15 of the flags word, if set, indicates that the window is currently mapped into the task's address space.

DETAILED TASK IMAGE FILE STRUCTURE

Word 9 contains the address of the associated region descriptor. If the memory-resident overlay is part of the task, and no region descriptor is allocated, this value is 0.

<u>Word</u>	15	8 7	0
0	Base Active Page Register		Window ID
1	Virtual base address		
2	Window size in 64-byte blocks		
3	Region ID		
4	Offset in partition		
5	Length to map		
6	Status word		
7	Send/receive buffer address (always 0)		
8	Flags word		
9	Address of region descriptor		

ZK-485-81

Figure B-17 Window Descriptor

B.4.5 Region Descriptor

The region descriptor is allocated only when the memory-resident overlay structure is part of a shared region. Figure B-18 illustrates the format of a region descriptor.

<u>Word</u>	
0	Region ID
1	Size of region
2	Region
3	name
4	Region
5	partition
6	Region status
7	Protection codes (always 0)
8	Flags

ZK-486-81

Figure B-18 Region Descriptor

DETAILED TASK IMAGE FILE STRUCTURE

Words 0 through 7 constitute a region descriptor in the format required by the mapping directives. The flags word is referred to by the overlay load routine. Bit 15 of the flags word, if set, indicates that a valid region identification is in word 0. If this bit is clear, the overlay load routine issues an Attach Region directive (with protection code set to 0) to obtain the identification.

APPENDIX C
HOST AND TARGET SYSTEMS

C.1 INTRODUCTION

You can build a task on one system (the host), and run it on another (the target). For example, your installation might consist of one large computer system with mapping hardware, and several smaller unmapped systems. On the large system you could create and debug tasks, and then transfer them to the smaller systems to run.

For example, if you are developing a task named TK3, using the default partition of your host system, the TKB command could be:

```
>TKB TK3,TK3=SQ1,SQ2
```

When you are ready to move TK3 to a target system, you build it again, indicating the mapping status of the target system and naming the partition in which the task is to reside:

```
>TKB  
TKB>TK3/-MM,TK3=SQ1,SQ2  
TKB>/  
Enter Options:  
TKB>PAR=PART1:100000:40000  
TKB>//
```

The resulting task image is ready to run on the unmapped target system.

You can transfer a task from the host system to the target system by following these steps:

1. Build the task image specifying the partition in which the task will run. If the target system is an unmapped system, specify the partition's base address and size.
2. Ensure that any shared regions accessed by the task are present in both systems.
3. If the target system and the host system do not have the same mapping status, set the Memory Management switch (/MM or /-MM) to reflect the mapping status of the target system.

The task code must not use any hardware options (FPP, EIS, EAE, and so forth) that are not present on the target system. This is particularly important if you are a FORTRAN user because FORTRAN tasks often use mathematics routines that are hardware dependent. (Refer to the IAS/RSX-11 FORTRAN IV Installation Guide and the IAS/RSX-11 FORTRAN IV User's Guide for more information on FORTRAN requirements).

HOST AND TARGET SYSTEMS

C.2 EXAMPLE C-1: TRANSFERRING A TASK FROM A HOST TO A TARGET SYSTEM

In this section, the resident library LIB and the task that refers to it, MAIN (from Example 4, Chapter 5), are rebuilt to run on an unmapped system. To save space, only the Task Builder command sequences are shown.

Assuming that the target system has a partition within it named LIB, you need to make only two changes to the original command sequence that builds the library:

1. You must attach the negated memory management switch (/MM) to the image file specification.
2. You must specify the partition base and length.

The modified command sequence is as follows:

```
TKB>LIB/-HD/PI/-MM,LIB/-WI,LIB=LIB
TKB>/
Enter Options:
TKB>STACK=0
TKB>PAR=LIB:136000:20000
TKB>//
```

If the target system does not contain a partition of the same name as that of the shared region, you must change the name of the shared region to match the name of an existing partition in the target system. This is a requirement of RSX-11M; on RSX-11M-PLUS systems it is not.

Assuming that the target system has a partition named GEN and that the task MAIN is to run in that partition in the target system, you must make three changes to the command sequence that builds the task MAIN:

1. You must attach the negated memory management switch (/MM) to the task image file specification.
2. You must eliminate the APR parameter of the RESLIB keyword.
3. You must explicitly specify the base address and length of the partition in which the task is to reside.

The modified command sequence is as follows:

```
TKB>MAIN/-MM,MAIN/MA/-WI=MAIN
TKB>/
Enter Options:
TKB>RESLIB=LIB/RO
TKB>PAR=GEN:30100:40000
TKB>//
```

Example C-1, Part 1 shows the map file of the resident library LIB for an unmapped system. LIB is bound to the partition base specified by the PAR keyword in the task-build command sequence. Note that the shared region is declared position independent even though it is bound to the partition base 136000. The position-independent declaration is not necessary in this example because the referencing task MAIN does not require the program section names within the library in order to refer to it. However, in applications involving tasks that require the program section names from the library, you must declare the library position independent so that TKB will place the program section names in the library's symbol definition file.

HOST AND TARGET SYSTEMS

Example C-1, Part 1 Task Builder Map for LIB.TSK

LIB.TSK;1 Memory allocation map TKB M40.10 Page 1
 10-DEC-82 11:50

Partition name : LIB
 Identification : 01
 Task UIC : [303,3]
 Task attributes: -HD,PI
 Total address windows: 1.
 Task image size : 64. words
 Task address limits: 000000 000163
 R-W disk blk limits: 000002 000002 000001 00001.

*** Root segment: LIB

R/W mem limits: 000000 000163 000164 00116.
 Disk blk limits: 000002 000002 000001 00001.

Memory allocation synopsis:

<u>Section</u>	<u>Title</u>	<u>Ident</u>	<u>File</u>
. BLK.:(RW,I,LCL,REL,CON)	000000 000000 00000.		
AADD : (RO,I,GBL,REL,CON)	000000 000024 00020.		
	000000 000024 00020.	LIB 01	LIB.OBJ;1
DIVV : (RO,I,GBL,REL,CON)	000024 000026 00022.		
	000024 000026 00022.	LIB 01	LIB.OBJ;1
MULL : (RO,I,GBL,REL,CON)	000052 000024 00020.		
	000052 000024 00020.	LIB 01	LIB.OBJ;1
SAVAL : (RO,I,GBL,REL,CON)	000076 000042 00034.		
	000076 000042 00034.	LIB 01	LIB.OBJ;1
SUBB : (RO,I,GBL,REL,CON)	000140 000024 00020.		
	000140 000024 00020.	LIB 01	LIB.OBJ;1

Global symbols:

AADD 000000-R MULL 000052-R SUBB 000140-R
 DIVV 000024-R SAVAL 000076-R

*** Task builder statistics:

Total work file references: 368.
 Work file reads: 0.
 Work file writes: 0.
 Size of core pool: 7086. words (27. PAGES)
 SIZE OF WORD FILE: 768. words (3. PAGES)

Elapsed time:00:00:03

HOST AND TARGET SYSTEMS

Example C-1, Part 2 shows the map file of the task MAIN for an unmapped system. The task is bound to the partition base 30100 and linked to the shared region LIB, which begins at 136000.

Example C-1, Part 2 Task Builder Map for MAIN.TSK

MAIN.TSK;1 Memory allocation map TKB M40.10 Page 1
11-DEC-82 13:41

Partition name : GEN
Identification : 01
Task UIC : [303,3]
Stack limits: 000274 001273 001000 00512.
PRG xfr address: 001634
Total address windows: 2.
Task image size : 1152. WORDS
Task address limits: 00000 004327
R-W disk blk limits: 000002 000006 000005 00005.

*** Root segment: MAIN

R/W mem limits: 000000 004327 004330 02264
Disk blk limits: 000002 000006 000005 00005.

Memory allocation synopsis:

<u>Section</u>	<u>Title</u>	<u>Ident</u>	<u>File</u>
. BLK.:(RW,I,LCL,REL,CON)	001274 002620 01424.		
	001274 000530 00344.	MAIN 01	MAIN.OBJ;1
	.		
	.		
	.		

Global symbols:

AADD	160000-R	SAVAL	060000-R	\$CBDSG	003110-R	\$CBTMG	003132-R
DIVV	160000-R	SUBB	060000-R	\$CBOMG	003116-R	\$CBVER	003116-R
IO.WVB	011000	\$CBDAT	003074-R	\$CBOSG	003124-R	\$CDDMG	003656-R
MULL	060000-R	\$CBDMG	003102-R	\$CBTA	003154-R	\$CDTB	003312-R
	.						
	.						
	.						

*** Task builder statistics:

Total work file references: 1889.
Work file reads: 0.
Work file writes: 0.
Size of core pool: 7086. WORDS (27. PAGES)
Size of work file: 1024. WORDS (4. PAGES)

Elapsed time:00:00:07

APPENDIX D

MEMORY DUMPS

The RSX-11M/M-PLUS Postmortem Dump task (PMD) generates postmortem memory dumps of tasks that are abnormally terminated. In addition, PMD can produce edited dumps, called snapshot dumps, for tasks that are running. Section D.1 describes Postmortem Dumps in general; Section D.2 discusses the specific case of snapshot dumps. Both types of dumps are very useful debugging aids.

D.1 POSTMORTEM DUMPS

You can make a task eligible for a Postmortem Dump in any of three ways:

- At task-build time, by specifying the /PM switch for the task file. /-PM disables dumps; it is the default condition.
- When you install a task by using the /PMD switch to override the taskbuild option. /PMD=YES enables dumping; /PMD=NO disables dumping.
- When you use the MCR command ABORT (described in the RSX-11M/M-PLUS MCR Operations Manual), by including the PMD switch in the command line to specify a dump.

You should install PMD in a 4K partition in which all other tasks are checkpointable. This allows the dump to be generated in a timely manner, and prevents the system from being locked up while the dump is being generated. PMD can dump either from memory or from the checkpoint image of your task. The PMD is sensitive to the location of the aborted task; therefore, if the aborted task is checkpointed during the dump, PMD switches to reading the checkpoint image. Once the task is checkpointed, PMD locks it out of memory until it has completed formatting the dump.

Dumps are always generated on the system disk under UFD [1,4]; therefore, to avoid errors from PMD, you must create a UFD for [1,4] before installing the task. When PMD finishes generating the dump, it attempts to queue the dump to the print spooler for subsequent printing. If no spooler is installed, the dump file is left on the disk and can be printed at a later time using the Peripheral Interchange Program (PIP, described in the RSX-11 Utilities Manual).

NOTE

Dump files tend to be somewhat large. The dump of an 8K partition averages about 340 blocks. Therefore, if there is little space on the disk, it is important to print and delete the dump

MEMORY DUMPS

file without delay. The print spooler automatically deletes all files with the type .PMD after printing them.

Example D-1 shows the contents of a Postmortem Dump and snapshot dump; the notes that follow the figure are keyed to the figure and provide a description of the dumps contents. Snapshot Dumps are explained more fully in Section D.2.

Example D-1 Sample Postmortem Dump (Truncated)

POST-MORTEM DUMP 1

TASK: TT6 2 TIME: 5-OCT-76 15:06

PC: 000720 3 IOT EXECUTION 3

REGS: R0 - 000345 R1 - 074400 R2 - 000120 R3 - 140130 } 4
R4 - 000000 R5 - 000000 SP - 000304 PS - 170000

TASK STATUS: MSG AST DST -CHK HLT STP REM MCR 5

EVENT FLAG MASK FOR <1-16> 000001 6

CURRENT UIC: [007,001] DSW: 1. 7

PRIORITY: DEFAULT - 50. RUNNING - 50. I/O COUNT: 0. TI DEVICE - TT6: 8

LOAD DEVICE - DB0: LBN: 1,160034 9

FLOATING POINT UNIT

STATUS - 000000

R0 - 000000 000000 000000 000000 } 10
R1 - 000000 000000 000000 000000
R2 - 000000 000000 000000 000000
R3 - 000000 000000 000000 000000
R4 - 000000 000000 000000 000000
R5 - 000000 000000 000000 000000

LOGICAL UNITS

UNIT	DEVICE	FILE STATUS
1	DB0:	
2	DB0:	
3	DB0:	
4	DB0:	

11

OVERLAY SEGMENTS LOADED AND RESIDENT LIBRARIES MAPPED

STARTING RELATIVE BLOCK: 000002 BASE: 000000 LENGTH: 001454 } 12
STARTING RELATIVE BLOCK: 000004 BASE: 001454 LENGTH: 000264

TASK STACK

ADDRESS	CONTENTS	ASCII	RAD50
000304	000045	%	7

13

TASK IMAGE

MEMORY DUMPS

Example D-1 (Cont.) Sample Postmortem Dump (Truncated)

```

PARTITION: GEN          VIRTUAL LIMITS: 000000 - 001777

000000 000304 000162 000001 067426 ! D6 B4 A Q08!
      304 000 162 000 001 000 026 157 !D r o!
000010 003401 003401 170017 000352 !AD3 AD3 8PQ E4!
      001 007 001 007 017 360 352 000 ! p j !
000020 000304 000000 000000 000000 ! D6 !
      304 000 000 000 000 000 000 000 !D !
000030 000000 000000 000000 000000 ! !
      000 000 000 000 000 000 000 000 ! !
000040 000000 140162 074106 000001 ! 01Z SIO A!
      000 000 162 300 106 170 001 000 ! r@ Fx !
000050 000000 000000 001104 000000 ! NT !
      000 000 000 000 104 002 000 000 ! D !
000060 000373 000000 000000 000000 ! Fk !
      373 000 000 000 000 000 000 000 ! !
000070 000000 074150 000004 051646 ! SJX D MON!
      000 000 150 170 004 000 246 123 ! hx &S!
000100 000000 051646 000000 051646 ! MON MON!
      000 000 246 123 000 000 246 123 ! &S &S!
000110 000000 051646 000000 000001 ! MON A!
      000 000 246 123 000 000 001 000 ! &S !
000120 067020 000000 001777 061404 !QXP YW O3.!
      020 156 000 000 377 003 004 143 ! n c!
000130 000020 000000 000600 007406 ! P IX BPF!
      020 000 000 000 200 001 006 017 ! !
000140 170000 000720 000000 000000 !8P KX !
      000 360 320 001 000 000 000 000 ! p P !
000150 140130 000120 074400 000345 !01 B SNP E/!
      130 300 120 000 000 171 345 000 !x@ P y e !
000160 000000 000000 000000 000000 ! !
      000 000 000 000 000 000 000 000 ! !

*** DUPLICATE THROUGH 000236 ***

000240 000000 000000 001110 000000 ! NX !
      000 000 000 000 110 002 000 000 ! H !
000250 001454 000264 000000 000000 ! TL DT !
      054 003 264 000 000 000 000 000 !, 4 !
000260 000001 001612 074360 003413 ! A VZ SN AEC!
      001 000 212 003 360 170 013 007 ! px !
000270 063014 131574 000000 000000 !PMD ... !
      014 146 174 263 000 000 000 000 ! f 3 !
000300 001051 000001 000045 050114 ! M3 A 7 L36!
      051 002 001 000 045 000 114 120 ! ) % LP!
000310 000000 000001 000100 000304 ! A AX D6!
      000 000 001 000 100 000 304 000 ! @ D !
000320 000524 000000 000000 000000 ! HT !
      124 001 000 000 000 000 000 000 !T !
000330 000000 000000 000000 063014 ! PMD!
      000 000 000 000 000 000 014 146 ! f!
000340 131574 047123 052120 052123 !... LUK MSX MS$!
      174 263 123 116 120 124 123 124 ! 3 SN PT ST!
000350 000000 016746 177734 012746 ! D1N 7T CTF!
      000 000 346 035 334 377 346 025 ! f \ f !
000360 001037 104377 103456 005046 ! MW U61 UYF AX8!
      037 002 377 210 056 207 046 012 ! . & !

```

14

MEMORY DUMPS

- ① Type of dump: Postmortem or snapshot. If it is a snapshot dump, the dump identification is printed.
- ② The name of the task being dumped, and the date and time the dump was generated.
- ③ The program counter at the time of the dump. If it is a Postmortem Dump, the reason the task was aborted is printed.
- ④ The general registers, stack pointer, and processor status at the time of the dump.
- ⑤ The task status flags at the time of the dump. See the description of ATL or TAL in the RSX-11M/M-PLUS MCR Operations Manual for the meaning of the flags.
- ⑥ The task event flag mask word at the time of the dump. If the dump is a snapshot dump, the efn specified in the SNAP\$ macro will be ON (see Section D.2.2).
- ⑦ The task UIC and the current value of the directive status word.
- ⑧ The task's priority and default priority, number of outstanding I/O requests, and the terminal from which the task was initiated (TI:).
- ⑨ The task load device and the logical block number for the start of the task image on the device.
- ⑩ The floating-point unit (FPU) registers or the extended arithmetic element (EAE) registers if the task is using one of these hardware features. If the task is not using the FPU or EAE, these registers are not printed. If the task uses the FPU and does not specify /FP on the task image file, or if it uses the EAE unit and has not specified the /EA switch, the registers are not printed. If the machine you are using has both an FPU and an EAE, PMD assumes you are using the FPU because it is the unit of choice for arithmetic computations.
- ⑪ The logical unit assignments at the time of the dump. UNIT is the logical unit number, and DEVICE is the device to which the logical unit is assigned. For snapshot dumps, the file names of any open files are displayed under FILE STATUS. Postmortem Dumps do not display this information because all of the files have been closed as a result of the I/O rundown on the aborted task.
- ⑫ The following are displayed: the overlay segments loaded and resident libraries mapped at the time of the dump; the relative block number of the segment; the base address; the length of the segment; and, for tasks using manual load, the segment names. For resident libraries, the library name is also displayed. The block number can be used to determine which segment is loaded, by reference to the memory allocation file generated by the Task Builder. The starting block number for each segment is the relative block number of the segment. By obtaining a match, you can determine the name of the segment in memory. Zero-length segments are usually co-tree roots.

MEMORY DUMPS

- 13 The task stack at the time of the dump. The address is displayed, along with the contents, in octal, ASCII, and Radix-50. Each word on the stack is dumped. If the stack pointer is above the initial value of the stack (H.ISP), only one word is dumped. The rest is dumped as part of the task image.
- 14 The task image itself. The partition being dumped and the limits of interest are displayed. For Postmortem Dumps, all address windows in use are dumped. For snapshot dumps, the virtual task limits that you request are displayed. The dump routine rounds the requested low limit down to the nearest multiple of eight bytes and rounds the requested high limit up to the nearest multiple of eight bytes. The dump image displays the virtual starting address of a 4-word block of memory, the data in both octal and Radix-50 on the first line, and byte octal and ASCII on the second line. A 4-word block that is repeated in a contiguous region of memory is printed once, and then noted by the message

*** DUPLICATE THROUGH xxxxxx ***

where xxxxxx indicates the last word that is duplicated. If the task was aborted, all address windows in use are dumped. If the dump is a snapshot dump, up to four contiguous blocks of memory can be dumped, if requested.

D.2 SNAPSHOT DUMP

Snapshot dumps are edited dumps produced for running tasks. You can request a snapshot dump any number of times during the execution of a task. The information generated is under the control of the programmer.

Snapshot dumps are generated by the following macros:

- SNPDEF\$ -- Defines offsets in the snapshot dump control block and defines control bits, which control the format of the dump
- SNPBK\$ -- Allocates the snapshot dump control block (see Table D-1)
- SNAP\$ -- Causes a snapshot dump to be generated

SNPBK\$ and SNAP\$ issue calls to SNPDEF\$; so, you need not explicitly issue the SNPDEF\$ macro call. Sections D.2.1 and D.2.2 describe the SNPBK\$ macro and the SNAP\$ macro, respectively.

MEMORY DUMPS

Label	Offset	
SB.CTL	0	CONTROL FLAGS
SB.DEV	2	DEVICE MNEMONIC
SB.UNT	4	UNIT NUMBER
SB.EFN	6	EVENT FLAG
SB.ID	10	SNAP IDENTIFICATION
SB.LM1 (L1)	12	MEMORY BLOCK 1
(H1)	14	LIMITS
(L2)	16	MEMORY BLOCK 2
(H2)	20	LIMITS
(L3)	22	MEMORY BLOCK 3
(H3)	24	LIMITS
(L4)	26	MEMORY BLOCK 4
(H4)	30	LIMITS
SB.PMD	32	"PMD..."
	34	IN RADIX-50

ZK-488-81

Figure D-1 Snapshot Dump Control Block Format

D.2.1 Format of the SNPBK\$ Macro

The format of the SNPBK\$ macro call is:

SNPBK\$ dev,unit,ctl,efn,id,L1,H1,L2,H2,L3,H3,L4,H4

dev

The 2-character ASCII name of the device to which the dump is directed. If it is a directory device, the UFD [1,4] must be on the volume. The dump is written to the disk and then spooled to the line printer. If there is no print spooler, the file is left on the disk. If the device is not a directory device, the dump goes directly to the device.

unit

The unit number of the device to which the dump is directed.

ctl

The set of flags that control the format of the dump and the data to be printed. The flags are:

- SC.HDR Print the dump header (items 3 to 10 in Figure D-1). Items 1 and 2 are always printed.
- SC.LUN Print information on all assigned LUNs (item 11).
- SC.OVL Print information about all loaded overlay segments (item 12).

MEMORY DUMPS

SC.STK Print the user stack (item 13).
SC.WRD Print the requested memory in octal words and
 Radix-50 (item 14).
SC.BYT Print the requested memory in octal bytes and ASCII
 (item 14).

efn

The event flag to be used to synchronize your program and PMD.

id

A number that identifies the snapshot dump. Because dumps can be requested at different times and under different conditions, this ID is used to identify the place or reason for the dump.

L1,L2,L3,L4

The starting addresses of the memory blocks to be dumped.

H1,H2,H3,H4

The ending addresses of the memory blocks to be dumped.

NOTE

If no memory is to be dumped, each limit
(L1,L2,L3,L4,H1,H2,H3,H4) should be 0.

Only one snapshot dump control block is allowed. It generates the global label `..SPBK`.

NOTE

Because `SNPBK$` is used to allocate storage for the snapshot dump control block, all arguments except `dev` must be valid arguments for `.WORD` or `.BYTE` directives.

D.2.2 Format of the SNAP\$ Macro

The format of the `SNAP$` macro is:

```
SNAP$ ctl,efn,id,L1,H1,L2,H2,L3,H3,L4,H4
```

ctl

The set of flags that control the format of the dump and the data to be printed. The flags are:

SC.HDR Print the dump header.
SC.LUN Print information on all assigned LUNs.
SC.STK Print the user stack.
SC.OVL Print information about all loaded overlay
 segments.

MEMORY DUMPS

SC.WRD Print the requested memory in octal words and Radix-50.

SC.BYT Print the requested memory in octal bytes and ASCII.

efn

The event flag to be used to synchronize your program and PMD. A Wait-For-Single-Event-Flag directive is always generated to perform synchronization.

id

A number that identifies the snapshot dump. Because dumps can be requested at different times and under different conditions, this ID is used to identify the place or reason for the dump.

L1,L2,L3,L4

The starting addresses of memory blocks to be dumped.

H1,H2,H3,H4

The ending addresses of memory blocks to be dumped.

NOTES

1. If no memory is to be dumped, each limit (L1,L2,L3,L4,H1,H2,H3,H4) should be 0.
2. You can set the control flags in any combination; they are not mutually exclusive. Thus, any number of options can be obtained; for example, SC.HDR!SC.LUN!SC.WRD prints the header, LUNs, and the requested memory in word octal and Radix-50 mode.
3. Arguments should be specified only to override the information already in the snapshot dump control block.
4. Because SNAP\$ generates instructions to move data into the snapshot dump control block, its arguments must be valid source operands for MOV instructions.

D.2.3 Example of a Snapshot Dump

The sample program shown in Example D-2 causes two snapshot dumps to be printed directly on LP0:. The first dump uses the parameters defined in the snapshot dump control block. The header is generated, and the data in relative locations BLK to BLK+220 is displayed, in word octal and Radix-50. The identification on the dump is 1.

The second dump causes the data in the locations BLK to BLK+220 to be displayed in byte octal and ASCII. A header is also generated. The dump identification is 64 (100 octal). Figures D-3 and D-4 show the dumps generated by the sample program.

SNPTST - TEST SNAP DUMP AND PMD MACRO M1010 03-SEP-76 15:57 PAGE 1

```

1          .TITLE  SNPTST - TEST SNAP DUMP AND PMD
2          .IDENT  /01/
3          .MCALL  SNPBK$,SNAP$,CALL
4 000000   BLK:    SNPBK$ LP,0,SC.HDR!SC.OVL!SC.WRD,1,1,BLK,BLK+220
5 000036   BUF:    .ASCIZ  /SNPTST/
   000041      123      116      120
   000044      124      123      124
   000044      000
6
7 000046   .EVEN
8 000216   START: SNAP$
   012700   000036'   MOV      #BUF,R0
9 000222   CALL    $CAT5
10 000226   SNAP$   #SC.HDR!SC.OVL!SC.BYT,,#100
11 000412   IOT
12 000046'   .END    START

```

SNPTST - TEST SNAP DUMP AND PMD MACRO M1010 03-SEP-76 15:57 PAGE 1-1
 SYMBOL TABLE

BLK	000000R	SB.EFN= 000006	SC.BYT= 000040	SC.STK= 000010	\$DSW = ***** GX
BUF	000036R	SB.ID = 000010	SC.HDR= 000001	SC.WRD= 000020	\$\$\$T2 = 000027
IE.ACT=	***** GX	SB.LM1= 000012	SC.LUN= 000002	START 000046R	..SPBK 000000RG
SB.CTL=	000000	SB.PMD= 000032	SC.OVL= 000004	\$CAT5 = ***** GX	...SNP= 000032
SB.DEV=	000002	SB.UNT= 000004			

. ABS. 000000 000
 000414 001

ERRORS DETECTED: 0

VIRTUAL MEMORY USED: 1335 WORDS (6 PAGES)
 DYNAMIC MEMORY AVAILABLE FOR 30 PAGES
 ASSEMBLY TIME (ELAPSED): 00:00:14
 SNPTST,SNPTST=SNPTST

Example D-2 Sample Program That Calls for Snapshot Dumps

D-9

MEMORY DUMPS

MEMORY DUMPS

Example D-3 Sample Snapshot Dump (in Word Octal and Radix-50)

SNAPSHOT DUMP ID: 1

TASK: TT6

TIME: 5-OCT-76 15:06

PC: 000522

REGS: R0 - 000000 R1 - 100104 R2 - 000000 R3 - 140130
R4 - 000000 R5 - 000000 SP - 000304 PS - 170000

TASK STATUS: MSG -CHK STP WFR REM MCR

EVENT FLAG MASK FOR 1-16> 000001

CURRENT UIC: [007,001] DSW: 1.

PRIORITY: DEFAULT - 50. RUNNING - 50. I/O COUNT: 0. TI DEVICE - TT6:

LOAD DEVICE - DB0: LBN: 1,160034

FLOATING POINT UNIT

STATUS - 000000

R0 - 000000 000000 000000 000000
R1 - 000000 000000 000000 000000
R2 - 000000 000000 000000 000000
R3 - 000000 000000 000000 000000
R4 - 000000 000000 000000 000000
R5 - 000000 000000 000000 000000

OVERLAY SEGMENTS LOADED AND RESIDENT LIBRARIES MAPPED

STARTING RELATIVE BLOCK: 000002 BASE: 000000 LENGTH: 001454

TASK IMAGE

PARTITION: GEN

VIRTUAL LIMITS: 000304 - 000524

000300	001051	000001	000025	050114	! M3	A	U	L36!
000310	000000	000001	000001	000304	!	A	A	D6!
000320	000524	000000	000000	000000	!	HT		!
000330	000000	000000	000000	063014	!			PMD!
000340	131574	047123	052120	052123	!...	LUK	MSX	MS\$!
000350	000000	016746	177734	012746	!	D1N	7T	CTF!
000360	001037	104377	103456	005046	!	MW	U61	UYF AX8!
000370	012746	000304	012746	000336	!	CTF	D6	CTF EV!
000400	017646	000000	062766	000002	!	EBV	PLV	B!
000410	000002	017666	000002	000002	!	B	EB8	B B!
000420	012746	0002507	104377	013435	!	CTF	31	U61 UX/!
000430	005046	005046	005046	005046	!	AX8	AX8	AX8 AX8!
000440	012746	000336	017646	000000	!	CTF	EV	EBV !
000450	062766	000002	000002	017666	!	PLV	B	B EB8!
000460	000002	000002	012746	003413	!	B	B	CTF AEC!
000470	104377	103006	022737	177771	!	U61	UQ0	FBO 8I!
000500	000046	001402	000261	000405	!	8	SJ	DQ FU!
000510	016746	177576	012746	001051	!	D1N	5F	CTF M3!
000520	104377	012700	000342	004767	!	U61	CSH	EZ AW1!

MEMORY DUMPS

Example D-4 Sample Snapshot Dump (in Byte Octal and ASCII)

SNAPSHOT DUMP ID: 64

TASK: TT6

TIME: 5-OCT-76 15:06

PC: 000716

REGS: R0 - 000345 R1 - 074400 R2 - 000120 R3 - 140130
R4 - 000000 R5 - 000000 SP - 000304 PS - 170000

TASK STATUS: MSG -CHK STP WFR REM MCR

EVENT FLAG MASK FOR 1-16> 000001

CURRENT UIC: [007001] DSW: 1.

PRIORITY: DEFAULT - 50. RUNNING - 50. I/O COUNT: 0. TI DEVICE - TT6:

LOAD DEVICE - DBO: LBN: 1,160034

FLOATING POINT UNIT

STATUS - 000000

R0 - 000000 000000 000000 000000
R1 - 000000 000000 000000 000000
R2 - 000000 000000 000000 000000
R3 - 000000 000000 000000 000000
R4 - 000000 000000 000000 000000
R5 - 000000 000000 000000 000000

OVERLAY SEGMENTS LOADED AND RESIDENT LIBRARIES MAPPED

STARTING RELATIVE BLOCK: 000002 BASE: 000000 LENGTH: 001454
STARTING RELATIVE BLOCK: 000004 BASE: 001454 LENGTH: 000264

(continued on next page)

MEMORY DUMPS

Example D-4 (Cont.) Sample Snapshot Dump (in Byte Octal and ASCII)

TASK IMAGE

PARTITION: GEN				VIRTUAL LIMITS: 000304 - 000524							
000300	051 002	001 000	045 000	114 120	!)	%	LP!				
000310	000 000	001 000	100 000	304 000	!	@	D !				
000320	124 001	000 000	000 000	000 000	!T		!				
000330	000 000	000 000	000 000	014 146	!		f!				
000340	174 263	123 116	120 124	123 124	! 3	SN	PT ST!				
000350	000 000	346 035	334 377	346 025	!	f	\ f !				
000360	037 002	377 210	056 207	046 012	!	.	& !				
000370	346 025	304 000	346 025	336 000	!f	D	f ^ !				
000400	246 037	000 000	366 145	002 000	!&		ve !				
000410	002 000	266 037	002 000	002 000	!	6	!				
000420	346 025	107 005	377 210	035 207	!f	G	!				
000430	046 012	046 012	046 012	046 012	!&	&	& & !				
000440	346 025	336 000	246 037	000 000	!f	^	& !				
000450	366 145	002 000	002 000	266 037	!ve		6 !				
000460	002 000	002 000	346 025	013 007	!	f	!				
000470	377 210	006 206	337 045	371 377	!		% y !				
000500	046 000	002 003	261 000	005 001	!&	I	!				
000510	346 035	176 377	346 025	051 002	!f	f) !				
000520	377 210	300 025	342 000	367 011	!	@	b w !				

APPENDIX E
RESERVED SYMBOLS

Several global symbols and program section names¹ are reserved for use by TKB.² Special handling occurs when TKB encounters a definition of one of these names in a task image.

The definition of a reserved global symbol in the root segment causes a word in the task image to be modified with a value calculated by TKB. The relocated value of the symbol is taken as the modification address.

The following global symbols are reserved by TKB:

Global Symbol	Modification Value
.FSRPT	Address of file storage region work area (.FSRCB).
.MBLUN	Mailbox logical unit number.
.MOLUN	Error message output device.
.NLUNS	The number of logical units used by the task, not including the message output and overlay units.
.NOVLY	The overlay logical unit number.
N.OVPT	Address of overlay run-time system work area (.NOVLY).
.NSTBL	The address of the segment description tables. This location is modified only when the number of segments is greater than one.
.ODTL1	Logical unit number for the ODT terminal device TI:.
.ODTL2	Logical unit number for the ODT line printer device CL:.

1. In **RSX-11M** and **RSX-11M-PLUS**, absolute sections (ASECTS) and both blank and named control sections (CSECTS) are supplanted by program sections (PSECTS). The .PSECT assembler directive eliminates the need for .ASECT and .CSECT directives, except for compatibility with other systems. This manual refers to all sections as program sections, unless the specific characteristics of ASECTS or CSECTS apply.

2. All symbols and program section names containing a period (.) or a dollar sign (\$) are reserved for DIGITAL-supplied software.

RESERVED SYMBOLS

Global Symbol	Modification Value
.SUML1	P/OS standard utility module LUN.
.PTLUN	Logical unit number for plotter/graphics software.
\$OTSVA	Address of Object Time System work area (\$OTSVA).
.TRLUN	The trace subroutine output logical unit number.
.USLU1	Logical unit number for special purpose user software.
.USLU2	Logical unit number for special purpose user software.
\$VEXT	Address of vector extension area (\$VEXTA).

TKB reserves the following program section names. In some cases, the definition of a reserved program section causes that program section to be extended if you specify the appropriate option.

Source Location	Section Name	Description
TKB	\$\$ALER	Contains code to process or trap Overlay Run-time System segment load errors. Provides named areas in the task for the FORTRAN-IV Object Time System and the RSX-11M Overlay Run-time System.
TKB	\$\$ALVC	Contains the segment autoloader vectors for tasks without I- and D-space.
TKB	\$\$ALVD	Contains the D-space portions of the segment autoloader vectors in an I- and D-space task.
TKB	\$\$ALVI	Contains the I-space portions of the segment autoloader vectors in an I- and D-space task.
TKB	\$\$AUTO	Contains code to determine if a called subroutine in an overlay segment is already in memory or if that overlay segment should be read into memory before control is passed to the subroutine that is called.
Input Module	\$\$DBTS	This symbol should appear in the debugger input module with the symbol \$DBTS as follows: <pre> .PSECT \$\$DBTS \$DBTS:: .PSECT </pre> <p>The task builder extends \$\$DBTS and fills it with time stamp information followed by the filename information of the .STB file.</p>
SYSLIB	\$\$DEVT	The extension length (in bytes) is calculated from the formula: <pre> EXT = <S.FDB+52>*UNITS </pre> <p>The definition of S.FDB is obtained from the root segment symbol table, and UNITS is the number of logical units used by the task, excluding the message output, overlay, and ODT units.</p>

RESERVED SYMBOLS

Source Location	Section Name	Description
SYSLIB	\$\$FSR1	The extension of this section is specified by the ACTFIL option.
SYSLIB	\$\$IOB1	The extension of this option is specified by the MAXBUF option.
TKB	\$\$IOB2	A zero length .PSECT containing a label, IOBFND, that is stored in the work area offset, W.BEND, representing the upper bound of the I/O buffer, \$\$IOB1. TKB uses \$\$IOB2 as a boundary value to determine whether the I/O buffer has overflowed.
TKB	\$\$LOAD	Overlay manual load routine.
TKB	\$\$MRKS	Contains code to properly mark those segments that are not needed any longer or have been overlaid by another segment as being out of memory. This ensures that a fresh copy of the overlay segment will be read in the next time the overlay segment is needed.
SYSLIB	\$\$OBF1	FORTRAN OTS uses this area to parse array type format specifications. This section can be extended by the FMTBUF keyword.
TKB	\$\$OBF2	A zero length .PSECT containing a label, OBFH, that is stored in the work area offset, W.OBFH, which represents the upper bound of the run-time format buffer, \$\$OBF1. TKB uses \$\$OBF2 to determine whether the run-time format buffer has overflowed.
TKB	\$\$OVDT	The Overlay Run-time System impure data area. The symbol .NOVPT in low memory points to this area. This area defines the operational parameters with which the Overlay Run-time System operates on disk-resident and memory-resident overlay structures.
TKB	\$\$OVRS	The .ABS. program section that redefines the Overlay Run-time System impure data area with different symbols, defined as offsets and relative to zero. These offsets are necessary for proper linkages between the subroutines in the Overlay Run-time System. This program section is never included in the memory allocation of the task because of its absolute program section attribute.
TKB	\$\$PDLs	Cluster library service routine.
TKB	\$\$RDSG	Contains the code that reads into memory the overlay segment selected by the code contained in the programs section \$\$AUTO.
TKB	\$\$RGDS	Contains the region descriptors for resident libraries referred to by the task.
TKB	\$\$RTQ	Defines the PSECT used for selective enabling of AST recognition in the Overlay Run-time System. \$\$RTQ is 0 in length if \$AUTOT is not included.

RESERVED SYMBOLS

Source Location	Section Name	Description
TKB	\$\$RTR	Defines the PSECT used for selective disabling of AST recognition in the Overlay Run-time System. \$\$RTR is 0 in length if \$AUTOT is not included.
TKB	\$\$RTS	Contains the return instruction.
TKB	\$\$SLVC	Supervisor-mode library transfer vectors (RSX-11M-PLUS only).
TKB	\$\$SGDO	Contains the program section adjoining the task segment descriptors.
TKB	\$\$SGD1	Contains the task segment descriptors.
TKB	\$\$SGD2	Contains a .WORD 0 following the task segment descriptors.
FORTRAN	\$\$TSKP	TKB fills in the following words in the PSECT: <ul style="list-style-type: none"> ● APR bit map in word \$APRMP ● Task offset into region in word \$LBOFF ● Maximum physical read/write memory needed for task in word \$MXLGH ● Maximum physical read-only memory needed for task in word \$MXLGH+2 ● Task extension in 32-word blocks in word \$LBEXT
TKB	\$\$WNDS	Contains task window descriptors

APPENDIX F

IMPROVING TASK BUILDER PERFORMANCE

This appendix contains procedures to assist you in maximizing Task Builder (TKB) performance. These procedures include:

- Evaluating and improving TKB throughput
- Modifying command switch defaults to provide a more efficient user interface
- Using the Slow Task Builder when large work file space is required

These procedures assume that the program to be linked requires features not found in the Fast Task Builder (FTB) described in Appendix G.

Using the procedures described in this appendix may require relinking TKB. You can do this only in a system that has, as a minimum, a 14K user-controlled or system-controlled partition. In some cases, you can make the modifications without relinking by using the binary patch program ZAP (see the RSX-11 Utilities Manual).

Modifications to the TKB build file imply one or more of the following files located under UFD [1,24] (mapped) or [1,20] (unmapped):

RSX-11M systems:
TKBBLD.CMD
STKBLD.CMD

RSX-11M-PLUS systems:
TKBBLD.CMD
STKBLD.CMD

These files reside on the disk containing the utility object files.

F.1 EVALUATING AND IMPROVING TASK BUILDER THROUGHPUT

Task Builder throughput is determined by three factors:

- The amount of disk latency incurred because of overlays
- The amount of memory available for table storage
- The amount of disk latency due to input file processing

The following sections outline methods for improving throughput in each of these last two cases.

IMPROVING TASK BUILDER PERFORMANCE

F.1.1 Table Storage

The principal factor governing TKB performance is the amount of memory available for table storage. To reduce memory requirements, a work file is used to store symbol definitions and other tables. This work file cannot exceed 65,543 bytes. As long as the size of these tables is within the limits of available memory, the contents of this file are kept in memory and the disk is not accessed. If the tables exceed this limit, some information must be displaced and moved to the disk, degrading performance accordingly.

You can gauge work file performance by consulting the statistics portion of the TKB map. The map displays the following parameters:

- Number of work file references -- Total number of times that work file data was referred to.
- Work file reads -- Number of work file references that resulted in disk accesses to read work file data.
- Work file writes -- Number of work file references that resulted in disk accesses to write work file data.
- Size of core pool -- Amount of in-core table storage in words. This value is also expressed in units of 256-word pages (information is read from and written to disk in blocks of 256 words).
- Size of work file -- Amount of work file storage in words. If this value is less than the pool size, the number of work file reads and writes is 0. That is, no work file pages are removed to the disk. This value is also expressed in pages (256-word blocks).
- Elapsed time -- Amount of time required to build the task image and output the map. This value excludes ODL processing, option processing, and the time required to produce the global cross-reference.

You can reduce the overhead for gaining access to the work file in one or more of the following ways:

- By increasing the amount of memory available for table storage
- By placing the work file on the fastest random access device
- By decreasing system overhead required to gain access to the file
- By reducing the number of work file references

You can increase the amount of table storage by installing TKB in a larger partition or, if TKB is running in a system-controlled partition, by using the INSTALL/INC keyword to allocate more space.

In a system that includes support for the Extend Task directive, TKB automatically increases its size if it is checkpointable and installed in a system-controlled partition. You set the maximum limit. You can increase this maximum by issuing the MCR command SET /MAXEXT.

Increasing the proportion of resident dynamic memory reduces the amount of I/O necessary for access to TKB internal data structures. As stated above, once the resident memory has been filled, the data structures overflow into a temporary work file on the device assigned

IMPROVING TASK BUILDER PERFORMANCE

to the work file logical unit number. This logical unit number (W\$KLUN) is specified in the build command file. Preferably, this unit number should be assigned to a device other than the system device, for example a fixed-head disk.

Displacement of pages to the work file is done on a least recently used basis. The work file extends automatically as necessary to hold all pages displaced. The parameter W\$KEXT is provided in the build command file of TKB and defines the file extension properties. A negative value indicates that the extend is noncontiguous; a positive value indicates that the extend is contiguous. If a contiguous extend fails, a noncontiguous request is attempted; if a noncontiguous extend fails, a fatal work file I/O error is reported. As long as the work file remains contiguous, a higher access rate can be obtained.

It is not possible to state exactly how many symbols TKB can process, because there are many data structures included in virtual memory. The following is a list of the structures that are stored in the virtual memory. All the sizes given are approximate only (sizes vary with characteristics of the task being built and may vary from release to release).

Structure Name	Description	Approximate Size (in Words)
Segment Descriptor	Contains listhead sizes, the pointers defining the overlay tree, the segment name. Part of this structure becomes the segment descriptor in the resultant task image.	80.
P-section Descriptor	Contains the name, address size, and attributes of a p-section.	10.
Symbol Descriptor	Contains symbol name, value, flags, and pointers to defining segment and program section descriptors.	8. (nonoverlaid task) 15. (overlaid task)
Element Descriptor	Contains module name, ident, filename, count of program section and some flags.	8.-18.
Control Section Mapping Table	Table of program section size and program section descriptor addresses.	Two words per program section in each module

The maximum usage of virtual memory occurs during phase three of TKB, when the symbol table is built. However, phase one makes significant use of virtual memory when an overlaid task is being built. It is at this point that all the segment descriptors are allocated, as well as an element descriptor for every file name encountered during the parsing of the tree description. In addition, a p-section descriptor is produced for every .PSECT directive encountered in the overlay description.

IMPROVING TASK BUILDER PERFORMANCE

The parsing of the overlay description also makes use of dynamic memory during the processing of each directive. This memory is released upon completion of the analysis; during the analysis, however, the whole tree description must fit into the resident portion of the storage. If sufficient storage cannot be obtained in the resident dynamic memory, the error message NO DYNAMIC STORAGE AVAILABLE is returned. The method for increasing the ratio of dynamic storage to virtual memory can be applied here, possibly to allow a task with a large overlay description to be built.

The amount of memory required during analysis depends on:

- The number of directives
- The length of .FCTR lines
- The number of operators, that is, commas, dashes, and parentheses)
- The number of file names encountered

TKB links all DEC-supplied tasks in a 14K partition.

There are a number of ways to reduce the amount of virtual memory required during the build of a specific task. Reducing the data structures in virtual memory also increases the speed of searching the tables and reduces the amount of paging to the work file.

1. Extract object modules by name from relocatable object libraries (for example., LIBRY/LB:MOD1:MOD2). This technique requires smaller element descriptors and fewer file name descriptors and is also faster because there are fewer files to open and close.
2. Use concatenated object modules for the same reasons as above.
3. Use shared regions (resident libraries and common areas) for language and overlay run-time systems and file control services. Such use of shared regions allows symbols and p-sections to be defined only once, rather than on multiple branches of the tree.
4. Place modules that occur on parallel branches of the tree in a common segment (for example, closer to root) for the same reasons as in 3 above.
5. Use the /SS switch on symbol table files (.STB) that describe absolute symbol definitions so that only those symbols referenced are extracted from the module.
6. Minimize the number of segments and keep the tree balanced. For example, if one segment is very long, there is no value in putting a tree structure in parallel unless creating one segment in parallel would be longer.

In addition to the above, a version of TKB can be built which has less throughput but requires less virtual memory per element than TKB. This version is built using the command file STKBLD.CMD supplied on the RK05 utility disk, or the RK06 and RP system disks under UFD [1,20] (unmapped) or [1,24] (mapped).

IMPROVING TASK BUILDER PERFORMANCE

There are four error messages associated with the virtual memory system:

- NO DYNAMIC STORAGE AVAILABLE. This error occurs when there is insufficient resident storage for creating some data structures. As much as possible of the data already allocated (all unlocked pages) has been paged to the work file, but there is still not enough free memory. Such a situation might arise during the analysis of the overlay description, early in the task-build run, and particularly if it is a complex tree. Reducing the ODL and extending the Task Builder memory allocation (see above) are the recommended recovery procedures.
- UNABLE TO OPEN WORKFILE. The probable causes of this error are:
 - Device assigned to logical unit 8 of the Task Builder is not mounted.
 - The device is not FILES-11.
 - There is no space on the volume.
 - The device is off line, not ready, write locked, or faulty.
 - There is no such device.

The MCR function LUN ...TKB may be used to determine which device the Task Builder is attempting to use.

- WORKFILE I/O ERROR. The probable causes of this error are:
 - Hardware error (for example, parity error on the disk).
 - Device is not ready, or is write-locked.
 - An extend failure has occurred (for example, the disk is full).
- NO VIRTUAL MEMORY STORAGE AVAILABLE. The addressable limit of the virtual memory has been reached. There is no recovery other than to reduce the virtual memory requirements of the task being built along the lines suggested earlier.

The work file normally resides on the device from which TKB was installed. You can change the device by reassigning logical unit 8 through the Monitor Console Routine or by editing the build file and relinking TKB.

System overhead for work file accesses is incurred in translating a relative block number in the file to a physical disk address. To minimize this overhead, TKB requests disk space in contiguous increments. The size of each increment is equal to the value of symbol W\$KEXT defined in TKB build file. A larger positive value causes the file to be extended in larger contiguous increments and reduces the overhead required to gain access to the file. The increment should be set to a reasonable value because TKB resorts to noncontiguous allocation whenever contiguous allocation fails. ;.b You can reduce the size of the work file by:

- Linking your task to a core-resident library containing commonly used routines (for example, FORTRAN Object Time System) whenever possible

IMPROVING TASK BUILDER PERFORMANCE

- Including common modules, such as components of an object time system, in the root segment of an overlaid task
- Using an object library or file of concatenated object modules if many modules are to be linked

When you use either of the last two procedures, system overhead is also significantly reduced because fewer files must be opened to process the same number of modules.

You can reduce the number of work file references by eliminating unneeded output files and cross-reference processing, or by obtaining the short map. In addition, you can usually exclude selected files, such as the default system object module library, from the map. In this case you can obtain, and retain, a full map at less frequent intervals.

F.1.2 Input File Processing

The procedures for minimizing the size of the work file and number of work file accesses also drastically reduce the amount of input file processing.

A given module can be read up to four times when the task is built:

- To build the symbol table
- To produce the task image
- To produce the long map
- To produce the global cross-reference

Files that are excluded from the long map are read only twice. The third and fourth passes are eliminated for all modules when you request a short map without a global cross reference.

F.1.3 Summary

In summary, you can use the following procedures to improve TKB throughput:

- Use the INSTALL/INC or EXTTSK keyword to allocate more table space.
- Increase maximum task size by raising the system limit for dynamic task extension.
- Reduce disk latency by placing the work file on the fastest random access device.
- Reduce system overhead by modifying the command file to allocate work file space in larger contiguous increments.
- Decrease work file size by using resident libraries, concatenated object files, and object libraries.

IMPROVING TASK BUILDER PERFORMANCE

- Decrease work file size by including common modules into the root segment of an overlaid task.
- Decrease the number of work file references by eliminating the map and global cross-reference, obtaining the short map, or excluding files from the map.

F.2 MODIFYING COMMAND SWITCH DEFAULTS

The default switch settings and values provided by the Task Builder as released may not suit the requirements of all installations. For example, the default /-EA (no Kell Extended Arithmetic Element) would be unsatisfactory at an installation that made frequent use of this hardware.

You can thus tailor the switch defaults by altering the contents of the words that contain initial switch states. Modifying TKB in this way is a 3-step process:

1. Consult Tables F-1 through F-4 to determine the switch word and bit to be altered.
2. Edit the appropriate TKB command file to include the switch word modification through a GBLPAT keyword referring to the global switch word name.
3. Relink TKB using the modified command file.

The command files for system tasks, as provided with the released system, require the standard set of TKB defaults; therefore, you must retain and use an unmodified copy of TKB whenever such tasks are relinked.

You use Tables F-1 through F-4 to alter the defaults as follows:

1. You identify the switch and the file to which it applies.
2. You consult the switch category entry in each table to locate the applicable switch words.
3. You look at the switch settings to find the switch and associated bit.
4. You specify the revised value and switch word as arguments in a GBLPAT keyword.
5. You relink TKB to produce a version containing the appropriate defaults.

For example, to change the TKB Extended Arithmetic Element default to /EA, perform the steps described below.

By consulting Table F-1, you determine that two switch words, \$DFSWT and \$DFTSK, contain task file switches. Of these, \$DFTSK contains the default setting for the /EA switch in bit 13. Setting this bit to 1 changes the initial switch setting to /EA. This new value is combined with the initial contents to yield the revised setting 120002. The required keyword input is:

```
TKB>GBLPAT=TASKB:$DFTSK:120002
```

IMPROVING TASK BUILDER PERFORMANCE

NOTE

The setting of bit positions not listed in the tables must not be altered.

The only switches that have associated values are /AC and /PR. In these cases, the value is the number of the initial APR used to map the task. You can alter the default by changing the value of the GBLDEF keyword for the symbol D\$FAPR in TKB build file. Only values 4 or 5 can be used.

Table F-1
Task File Switch Defaults

Switch Category: Task file

Switch Word: \$DFSWT

Initial Contents: 0

Switch Settings:

Bit	Initial State	Initial Condition
15	0	/-XT Not abort after n diagnostics
11	0	/-SQ Not sequential PSECT allocation
4	0	/-FU Not full overlay tree search
3	0	/RO Recognize memory-resident overlay operator.
2	0	/-ID Not user D-space

Switch Category: Task file

Switch Word: \$DFTSK

Initial Contents: 100002

Switch Settings:

Bit	Initial State	Initial Condition
15	1	/-CP, /-AL Not checkpointable ¹
14	0	/-FP Not Floating Point Processor
13	0	/-EA Not Extended Arithmetic Element

1. The combination of not checkpointable with checkpoint allocation (100000) is illogical and should not be used.

(continued on next page)

IMPROVING TASK BUILDER PERFORMANCE

Table F-1 (Cont.)
Task File Switch Defaults

Switch Settings: (Cont.)

Bit	Initial State	Initial Condition
12	0	/HD Header
11	0	/-CM Not compatibility mode
10	0	/-DA No debugging aid
9	0	/-PI Not position independent
8	0	/-PR Not privileged
7	0	/-TR No trace
6	0	/-PM No Postmortem Dump
5	0	/-SL Not slave task
4	0	/SE Send to task allowed
2	0	/-AC Not ancillary control processor
1	1	/-AL No checkpoint allocation
0	0	/XH External header

Switch Category: Task File

Switch Word: \$DFTSO

Initial Contents: 000010

Switch Settings:

Bit	Initial State	Initial Condition
8	0	/-XH No external header
3	1	/-SG RO and RW PSECTs

IMPROVING TASK BUILDER PERFORMANCE

Table F-2
Map File Switch Defaults

Switch Category: Map file
 Switch Word: \$DFLBS
 Initial Contents: 120000
 Switch Settings:

Bit	Initial State	Initial Condition
15	1	/-MA Do not include system library and STB files in map

Switch Category: Map file
 Switch Word: \$DFMAP
 Initial Contents: 2040
 Switch Settings:

Bit	Initial State	Initial Condition
10	1	/SH Short map
8	0	/SP Spool
6	0	/-CR No CREF
5	1	/WI Wide format

Table F-3
Symbol Table File Switch Defaults

Switch Category: Symbol table file
 Switch Word: \$DFSTB
 Initial Contents: 0
 Switch Settings:

Bit	Initial State	Initial Condition
12	0	/HD Build task with header
9	0	/-PI Task is not position independent

IMPROVING TASK BUILDER PERFORMANCE

Table F-4
Input File Switch Defaults

Switch Category: Input file

Switch Word: \$DFINP

Initial Contents: 000100

Switch Settings:

Bit	Initial State	Initial Condition
15	0	/MA Include file contents in map
6	1	/CC File contains two or more concatenated object modules

F.3 THE SLOW TASK BUILDER

TKB.TSK uses a symbol table structure that can be searched quickly, but which requires more work file space than that of previous versions. You may thus receive the following message in some instances:

NO VIRTUAL MEMORY STORAGE AVAILABLE

If this occurs, you should try to reduce the work file size by using the procedures described in Section F.1. If these procedures do not sufficiently reduce the work file size, you can link another version of TKB, the Slow Task Builder. This version requires less storage, but runs considerably slower than the other versions. The build file is STKBLD.CMD, which resides on the same device and UFD as the other Task Builder command files. The default name of STK.TSK, the Slow Task Builder, is ...TKB. It may be convenient to install the Slow Task Builder with a different name if you want to use both Task Builders in your system.

APPENDIX G

THE FAST TASK BUILDER

The Fast Task Builder (FTB) allows you to build simple tasks about four times faster than the Task Builder (TKB). However, FTB has limited functionality. It can only link single-segment, nonprivileged tasks, and supports a limited number of switches and options.

The (FTB) is intended for use as a load-and-go type of linker. It contains very few options and does not support:

- New map format
- Overlaid programs
- FORTRAN virtual arrays
- Production of symbol table files
- Creation of resident libraries
- Privileged tasks
- Cluster libraries

The only supported switches are:

- /SP on map file (default = /SP)
- /CP on task file (default = /CP)¹
- /EA on task file (default = /-EA)
- /MM on task file (default = /MM)
- /FP on task file (default = /FP)
- /DA on input or task image (default = /-DA)
- /LB on an input file in the form:

>TKB TASK=PROG.OBJ,LIBRARY/LB

but not in the form:

>TKB TASK=PROG.OBJ,LIBRARY/LB:MODULE

1. No checkpoint space is allocated in the task image file.

THE FAST TASK BUILDER

The supported option inputs are:

- ASG (same defaults as TKB)
- STACK (same default as TKB)
- UNITS
- TASK (same default as TKB)
- EXTSCCT
- ACTFIL (same default as TKB)
- MAXBUF (same default as TKB)
- LIBR
- COMMON
- RESLIB (same defaults as TKB)
- RESCOM (same defaults as TKB)
- SUPLIB
- RESSUP (same defaults as TKB)

FTB supports linking to shared regions but not building a shared region. FTB cannot link to clustered libraries. Though FTB can link to supervisor-mode libraries, it cannot link to overlaid supervisor-mode libraries.

FTB allocates symbol table space from the end of its image to the end of the partition. It does not have a virtual symbol table. An Extend Task or equivalent of 8K is recommended. FTB does not dynamically extend itself at run time.

FTB runs approximately four times faster than TKB on an 11/70 with RP04s when TKB is running with a totally resident symbol table. In smaller systems with slower disks, the ratio should be much higher.

FTB also supports shared regions.

FTB uses asynchronous system traps (ASTs) and therefore requires AST support in the Executive.

APPENDIX H
ERROR MESSAGES

The Task Builder (TKB) produces diagnostic and fatal error messages. Error messages are printed in the following forms:

TKB -- *DIAG*-error-message

or

TKB -- *FATAL*-error-message

Some errors are correctable when command input is from a terminal. In such a case, a diagnostic error message can be printed, the error corrected, and the task-building sequence continued. However, if the same error is detected in an indirect command file, a correction cannot be made and the Task Builder aborts.

Some diagnostic error messages merely advise you of an unusual condition. If you consider the condition normal for your task, you can install and run the task image.

NOTE

The Task Builder exits with two statuses: it returns an ERROR status when it encounters a diagnostic error, and a SEVERE ERROR when it encounters a fatal error. (For more information about the Exit-With-Status directive, see the RSX-11M/M-PLUS Executive Reference Manual.)

This appendix tabulates the error messages produced by TKB. Most of the messages are self-explanatory. In some cases, the line in which the error occurred is printed.

A Software Performance Report (SPR) should be submitted to DIGITAL in cases where the explanation accompanying a message refers to a system error.

Allocation failure on file file-name

TKB could not acquire sufficient disk space to store the task image file, or did not have write-access to the UFD or volume that was to contain the file.

ERROR MESSAGES

Blank P-section name is illegal overlay-description-line

The overlay-description-line printed contains a .PSECT directive that does not have a p-section name.

Cluster library element library-name does not have null root

This is a fatal error. All libraries, except the first, must be PLAS-overlaid and have a null root. The first library in the group can be nonoverlaid or overlaid with a null or non-null root.

Command I/O error

An I/O error occurs on a command input device. (Device may not be on line, or possible hardware error.)

Command syntax error command-line

The command-line printed has incorrect syntax.

Complex relocation error - divide by zero: module module-name

A divisor having the value 0 was detected in a complex expression. The result of the divide was set to 0. (Probable cause: division by a global symbol whose value is undefined.)

Conflicting base addresses in cluster library

This conflict arises when you specify APRs, for both PIC and non-PIC libraries that are included in the cluster. See the APR parameter as described in the CLSTR option. This is a fatal error.

Disk image core allocation too large invalid-line

The minimum disk allocation specified in the invalid line is greater than 128.

File file-name attempted to store data in virtual section

The file contains a module that has attempted to initialize a virtual section with data.

File file-name has illegal format

The file file-name contains an object module whose format is not valid.

Illegal APR reservation

An APR specified in a COMMON, LIBR, RESCOM, or RESLIB keyword is outside the range 0-7.

Illegal cluster configuration

If the cluster contains a non-overlaid library, that library must be the first library in the cluster. Check the configuration of the libraries in the cluster. This is a fatal error.

ERROR MESSAGES

Illegal default priority specified option-line

The option-line printed contains a priority greater than 250.

Illegal device/volume invalid-line

The invalid line printed contains an illegal device specification.

Illegal directory invalid-line

The invalid line printed contains an illegal directory name.

Illegal error-severity code octal-list

System error (no recovery). An SPR should be submitted with a copy of the message containing the octal-list as printed.

Illegal filename invalid-line

The invalid-line printed contains a wildcard (*) in a file specification. Using wildcards is prohibited.

Illegal get command line error code

System error (no recovery).

Illegal logical unit number invalid-line

The invalid-line printed contains a device assignment to a unit number larger than the number of logical units specified by the UNITS keyword, or assumed by default if the UNITS keyword is not used.

Illegal multiple parameter sets invalid-line

The invalid-line printed contains multiple sets of parameters for a keyword that allows only a single parameter set.

Illegal number of logical units invalid-line

The invalid-line printed contains a logical unit number greater than 250.

Illegal ODT or task vector size

ODT or SST vector size specified is greater than 32 words.

Illegal overlay description operator invalid-line

The invalid-line printed contains an unrecognizable operator in an overlay description. This error occurs if the first character in a p-section or segment name is a dot (.).

ERROR MESSAGES

Illegal overlay directive invalid-line

The invalid-line printed contains an unrecognizable overlay directive.

Illegal partition/common block specified invalid-line

User-defined base or length is not on a 32-word boundary.

Illegal P-section/segment attribute invalid-line

The invalid-line printed contains a program section or segment attribute that is not recognized.

Illegal reference to library P-section p-sect-name

A task has attempted to reference a p-sect-name existing in a shared region but has not named the shared region in a keyword. This error occurs when you explicitly specify an STB file as an input file, but you have not specified the library to which the STB file belongs in an option.

Illegal switch file-specification

The file-specification printed contains an illegal switch or switch value.

Incompatible reference to library P-section p-sect-name

A task has attempted to reference more storage in a shared region than exists in the shared region definition.

Incorrect library module specification invalid-line

The invalid-line contains a module name with a non-Radix-50 character.

Indirect command syntax error invalid-line

The invalid-line printed contains a syntactically incorrect indirect file specification.

Indirect file depth exceeded invalid-line

The invalid-line printed gives the file reference that exceeded the permissible indirect file depth (2).

Indirect file open failure invalid-line

The invalid-line contains a reference to a command input file that could not be located.

Insufficient APRs available to map read-only root

TKB could not find enough free APRs to map the read-only portion of a multiuser task.

ERROR MESSAGES

Insufficient parameters invalid-line

The invalid-line contains a keyword with an insufficient number of parameters to complete its meaning.

Invalid APR reservation invalid-line

APR is specified on a keyword for an absolute library.

Invalid keyword identifier invalid-line

The invalid-line printed contains an unrecognizable keyword.

Invalid partition/common block specified invalid-line

A partition is invalid for one of the following reasons:

- TKB cannot find the partition name in the host system in order to get the base and length.
- The system is mapped, but the base address of the partition is not on a 4K boundary for a nonrunnable task or is not 0 for a runnable task.
- The memory bounds for the partition overlap a shared region.
- The partition name is identical to the name of a previously defined COMMON or LIBR shared region.
- The top address of the partition for a runnable task exceeds 32K minus 32 words for a mapped system, or exceeds 28K minus 1 for an unmapped system.
- A system-controlled partition was specified for an unmapped system.

Invalid reference to mapped array by module module-name

The module has attempted to initialize the mapped array with data. An SPR should be submitted if DIGITAL-supplied software caused this problem.

Invalid window block specification invalid-line

The number of extra address windows specified exceeds the number permitted. On an RSX-11M system, you can specify as many as 7 extra window blocks; on an RSX-11M-PLUS system, you can specify as many as 15 extra window blocks.

If you build a task on an RSX-11M system and specify more window blocks, you get this error message, but the task will build. However, it cannot be installed and run on an RSX-11M system.

I/O error library image file

An I/O error has occurred during an attempt to open or read the Task Image File of a shared region.

ERROR MESSAGES

I/O error on input file file-name

This error occurs when TKB cannot read an input file specification (for example, when the command line is greater than 80 characters).

I/O error on output file file-name

Label or name is multiply defined

invalid-line

The invalid-line printed defines a name that has already appeared as a .FCTR, .NAME, or .PSECT directive.

Library file filename has incorrect format

A module has been requested from a library file that has an empty module name table.

Library not built as a supervisor mode library

The library referred to in a RESSUP or SUPLIB option was built without a completion (CMPRT=X) routine and is not a supervisor-mode library.

Library library-name not found in any cluster

All task image and symbol table files to be included as cluster elements must reside in LB:[1,1].

Library references overlaid library

invalid-line

An attempt was made to link the resident library being built to a shared region that has memory-resident overlays.

Load addr out of range in module module-name

An attempt has been made to store data in the task image outside the address limits of the segment. This problem is usually caused by one of the following:

- An attempt to initialize a p-section contained in a shared region
- An attempt to initialize an absolute location outside the limits of the segment or in the task header
- A patch outside the limits of the segment to which it applies
- An attempt to initialize a segment having the NODSK attribute

Lookup failure on file file name

invalid-line

The invalid-line printed contains a file name that cannot be located in the directory.

Lookup failure on system library file

TKB cannot find the system Library (SY0:[1,1]SYSLIB.OLB) file to resolve undefined symbols.

ERROR MESSAGES

Lookup failure resident library file - filename.ext

No symbol table or task image file can be found for the shared region "filename.ext." If the shared region was linked to another shared region, ensure that the task image of both regions and the symbol table files exist on the same device and in the same UIC as the UIC referenced by the option RESLIB, RESCOM, LIBR, COMMON, RESSUP, or SUPLIB.

Module module-name ambiguously defines P-section p-sect-name

The p-section p-sect-name has been defined in two modules not on a common path, and referenced from a segment that is common to both paths.

Module module-name ambiguously defines symbol sym-name

Module module-name references or defines a symbol sym-name whose definition exists on two different paths, but is referenced from a segment that is common to both paths.

Module module-name contains incompatible autoloader vectors

You are trying to build an I- and D-space task and link it to an older existing library that contains an old-style vector format. Rebuild the library with the RSX-11M-PLUS Version 2.1 or later version Task Builder to create new-style vectors for the library. Then rebuild your task.

Module module-name illegally defines xfr address p-sect-name addr

This message occurs under any one of the following conditions:

- The start address printed is odd.
- The module module-name is in an overlay segment and has a start address. The start address must be in the root segment of the main tree.
- The address is in a p-section that has not yet been defined. An SPR should be submitted DIGITAL-supplied software caused this problem.

Module module-name multiply defines P-section p-sect-name

- The p-section p-sect-name has been defined more than once in the same segment with different attributes.
- A global p-section has been defined more than once with different attributes in more than one segment along a common path.

Module module-name multiply defines symbol sym-name

Two definitions for the relocatable symbol sym-name have occurred on a common path. Or two definitions for an absolute symbol with the same name but different values have occurred.

Module module-name multiply defines xfr addr in seg segment-name

This error occurs when more than one module making up the root has a start address.

ERROR MESSAGES

Module module-name not in library

TKB could not find the module named on the LB switch in the library.

No dynamic storage available

TKB needs additional symbol table storage and cannot obtain it. (If possible, install TKB in a larger partition.)

No memory available for library library-name

TKB could not find enough free virtual memory to map the specified shared region.

No root segment specified

The overlay description did not contain a .ROOT directive.

No virtual memory storage available

Maximum permissible size of the work file is exceeded. The user should consult Appendix F for suggestions on reducing the size of the work file.

Open failure on file file-name

Option syntax error

invalid-line

The invalid-line printed contains unrecognizable syntax.

Overlay directive has no operands

invalid-line

All overlay directives except .END require operands.

Overlay directive syntax error

invalid-line

The invalid-line printed contains a syntax error or references a line that contains an error.

Partition partition-name has illegal memory limits

- The partition-name defined in the host system has a base address alignment that is not compatible with the target system.
- The user has attempted to build a privileged task in a partition whose length exceeds the task's available address space (8K or 12K).

Pass control stack overflow at segment segment-name

System error. An SPR should be submitted with a copy of the ODL file associated with the error.

PIC libraries may not reference other libraries

invalid-line

The user has attempted to build a position-independent shared region that references another shared region.

ERROR MESSAGES

P-section p-sect-name has overflowed

A section greater than 32K has been created.

Required input file missing

At least one input file is required for a task build.

Required partition not specified

The PAR keyword was not used when running TKB on an RSX-11D host system. The keyword must contain explicit base address and length specifications.

Resident library has incorrect address alignment invalid-line

The invalid-line specifies a shared region that has one of the following problems:

- The library references another library with invalid address bounds (that is, not on 4K boundary in a mapped system).
- The library has invalid address bounds.

Resident library mapped array allocation too large invalid-line

The invalid-line printed contains a reference to a shared region that has allocated too much memory in the task's mapped array area. The total allocation exceeds 2.2 million bytes.

Resident library memory allocation conflict keyword-string

One of the following problems has occurred:

- More than seven shared regions have been specified.
- A shared region has been specified more than once.
- Non-position-independent shared regions whose memory allocations overlap have been specified.

Root segment is multiply defined invalid-line

The invalid-line printed contains the second .ROOT directive encountered. Only one .ROOT directive is allowed.

Segment seg-name has addr overflow: allocation deleted

Within a segment, the program has attempted to allocate more than 32,764 words (32K-1 words). A map file is produced, but no task image file is produced.

Segment seg-name not found for patch

The Task Builder could not locate the named segment for a global patch. The option used was GBLPAT=X:Y:0.

ERROR MESSAGES

Supervisor mode completion routine is undefined

TKB could not locate the symbol X, which was specified in the CMPRT=X option.

Symbol sym-name not found for patch

TKB could not locate symbol Y for a global patch. The option used was GBLPAT=X:Y:0.

Task has illegal memory limits

An attempt has been made to build a task whose size exceeds the partition boundary. If a task image file was produced, it should be deleted.

Task has illegal physical memory limits mapped-array task-image task extension

The sum of the parameters displayed -- mapped array size, task image size, and task extension -- exceeds 2.2 million bytes. The quantities are shown as octal numbers in units of 64-byte blocks. Any resulting task image file should be deleted.

Task image file filename is noncontiguous

Insufficient contiguous disk space was available to contain the task image. A noncontiguous file was created. After deleting unnecessary files, the /CO switch in PIP should be used to create a contiguous copy.

Task requires too many window blocks

The number of address windows required by the task and any shared regions exceeds 8 for RSX-11M tasks and 16 for RSX-11M-PLUS tasks.

Task-build aborted via request option-line

The option-line contains a request from the user to abort the task build.

Too many nested .ROOT/.FCTR directives invalid-line

The invalid-line printed contains a .FCTR directive that exceeds the maximum nesting level (16).

Too many parameters invalid-line

The invalid-line printed contains a keyword with more parameters than required.

Too many parentheses levels invalid-line

The invalid-line printed contains a parenthesis that exceeds the maximum nesting level (16).

ERROR MESSAGES

Truncation error in module module-name

An attempt has been made to load a global value greater than +127 or less than -128 into a byte. The low-order eight bits are loaded.

Unable to open work file

The work file device is not mounted. (The work file is usually located on the same device as is the Task Builder.)

Unbalanced parentheses invalid-line

The invalid-line printed contains unbalanced parentheses.

n Undefined symbols segment seg-name

The segment named contains n undefined symbols. If no memory allocation file is requested, the symbols are printed on the terminal.

Virtual section has illegal address limits option-line

The option-line printed contains a VSECT keyword whose base address plus window size exceeds 177777.

Work file I/O error

An I/O error occurs during an attempt to reference data stored by TKB in its work file.

TASK BUILDER GLOSSARY OF TERMS

ABSOLUTE SHARED REGION

A shared region that has the same virtual addresses in all tasks that refer to it.

AUTOLOAD

The method of loading overlay segments, in which the Overlay Run-Time routines automatically load overlay segments when they are needed and handles any unsuccessful load requests.

AUTOLOAD VECTOR

A transfer of control instruction generated by the Task Builder to resolve an up-tree reference to a global symbol.

CO-TREE

One of one or more secondary tree structures within a multiple tree overlay structure. When a co-tree's root segment contains code or data, the root segment of the co-tree is made resident in physical memory through calls to the Overlay Run-Time routines.

COMMON BLOCK

Another name for resident common.

DISK-RESIDENT

That which resides on disk storage until needed.

DISK-RESIDENT OVERLAY SEGMENT

An overlay segment that shares the same physical memory and virtual address space with other segments. The segment is read in from disk each time it is loaded (compare Memory-Resident Overlay Segment).

GLOBAL CROSS-REFERENCE

A list of global symbols, in alphabetical order, accompanied by the name of each referencing module.

GLOBAL SYMBOL

A symbol whose definition is known outside the defining module.

HEADER

That portion of a task image that contains the task's characteristics and status. Shared regions, although built like a task, do not have a header.

TASK BUILDER GLOSSARY OF TERMS

HOST SYSTEM

The system on which a task is built.

LOGICAL ADDRESSES

The actual physical addresses that the task can access.

LOGICAL ADDRESS SPACE

The total amount of physical memory to which the task has access rights.

MAIN TREE

An overlay tree whose root segment is loaded by the Executive when the task is made active.

MANUAL LOAD

The method of loading overlay segments in which the user includes explicit calls in his routines to load overlays and handles unsuccessful load requests.

MAPPED ARRAY AREA

An area of the task's physical memory, preceding the task image, that is used for storage of large arrays. Space in the area is reserved by means of the VSECT keyword or through a Mapped Array Declaration contained in an object module. Access is through the mapping directives issued at run time.

MEMORY ALLOCATION FILE

The output file created by the Task Builder that lists information about the size and location of components within a task.

MEMORY-RESIDENT

In general, that which resides in memory all the time. The entity, as in the case of memory-resident overlays, may initially reside on disk.

MEMORY-RESIDENT OVERLAY SEGMENT

An overlay segment that shares virtual address space with other segments, but which resides in its own physical memory. The segment is loaded from disk only the first time it is referenced; thereafter, mapping directives are issued in place of disk load requests.

OVERLAY DESCRIPTION LANGUAGE

A language that allows you to describe the overlay structure of a task.

TASK BUILDER GLOSSARY OF TERMS

OVERLAY RUNTIME ROUTINES

A set of system library subroutines linked as part of an overlaid task that are called to load segments into memory.

OVERLAY SEGMENT

A segment that shares virtual address space with other segments, and is loaded when needed.

OVERLAY TREE

A tree structure consisting of a root segment and optionally one or more overlay segments.

PATH

A route that is traced from one segment in the overlay tree to another segment in that tree.

PATH-DOWN

A path toward the root of the tree.

PATH-LOADING

The technique used by the autoloading method to load all segments on the path between a calling segment and a called segment.

PATH-UP

A path away from the root of the tree.

PHYSICAL ADDRESS

The assigned byte location in physical memory, which is usually located in the processing unit.

POSITION-INDEPENDENT REGION

A shared region that can be placed anywhere in a referencing task's virtual address space when the system on which the task runs has memory management hardware.

PRIVILEGED TASK

A task that has privileged memory access rights. A privileged task can access the Executive and the I/O page in addition to its own partition and referenced shared regions.

PROGRAM SECTION

A section of memory that is a unit of the total allocation. A source program is translated into object modules that consist of program sections with attributes describing access, allocation, relocatability, and so forth.

REGION

A contiguous block of physical addresses in which a driver, a task, a resident common, or library resides.

TASK BUILDER GLOSSARY OF TERMS

RESIDENT COMMON

A shared region in which data resides that can be shared by two or more tasks.

RESIDENT LIBRARY

A shared region in which single copies of commonly used subroutines reside that can be shared by two or more tasks.

ROOT SEGMENT

The segment of an overlay tree that, once loaded, remains in memory during the execution of the task.

RUNNABLE TASK

A task that has a header and stack and that can be installed and executed.

SHARED REGION

A shared region is a block of data or code that resides in physical memory and can be used by any number of tasks. A shared region is built and installed separately from the task.

SUPERVISOR-MODE LIBRARY

A library of routines that uses the supervisor-mode memory management APRs to map to both the task and its own routines.

SYMBOL DEFINITION FILE

The output object file created by the Task Builder that contains the global symbol definitions and values and sometimes program section names, attributes, and allocations in a format suitable for reprocessing by the Task Builder. Symbol definition files contain linkage information about shared regions.

TARGET SYSTEM

The system on which a task executes.

TASK IMAGE FILE

The output file created by the Task Builder that contains the executable portion of the task.

VIRTUAL ADDRESSES

The addresses within the task. Task addresses can range from zero through 177777(8) depending on the length of the task.

VIRTUAL ADDRESS SPACE

That space encompassed by the range of virtual addresses that the task uses.

VIRTUAL PROGRAM SECTION

A program section that has virtual memory allocated to it, but not physical memory. Virtual address space is mapped into physical memory at run-time by means of the mapping directives.

TASK BUILDER GLOSSARY OF TERMS

WINDOW

A continuous virtual address space that can be moved to allow the task to examine different parts of a region or different regions.

WINDOW BLOCK

A structure defined by the Task Builder that describes a range of continuous virtual addresses.

INDEX

- Abort
 - TKB
 - during input, 11-4
- ABORT option, 11-4
- Absolute region
 - See Region
- ABSPAT option, 11-5
- /AC switch, 10-5
- Access-code, 2-4 to 2-5
 - grouping program section by, 10-36
- ACP
 - specifying APR, 10-5
 - specifying task as, 10-5
- ACTFIL option, 11-6
- Active Page Register
 - See APR
- Address
 - assigning, 2-1
 - concepts, 2-13
 - logical, 2-14
 - mapped system, 2-18
 - physical, 2-13
 - mapped system, 2-17
 - space, 2-14
 - virtual, 2-14
 - virtual and logical
 - coincidence, 2-15
 - space translation, 2-23
 - transfer, A-5
 - virtual, 2-13
 - mapped system, 2-17 to 2-18
 - virtual and logical
 - relationship, 2-17, 2-21
 - virtual space
 - co-tree and main tree, 3-40
 - disk-resident overlay, 3-1
 - to 3-4
 - division, 2-18
 - division by memory
 - management, 2-18
 - in overlay tree, 3-24
 - memory-resident overlay, 3-6
 - overlaid task, 3-10, 3-14
 - overlay, 3-5, 3-33 to 3-35
 - reducing usage of, 3-1
 - virtual space allocation
 - diagram
 - creating ODL file, 3-36
 - virtual space and memory
 - overlaid task, 3-11 to 3-14
- /AL switch, 10-6
- \$\$ALER
 - reserved PSECT name, E-2
- ALERR module, 5-53
- Allocation-code, 2-4 to 2-5
- ALSCT FORTRAN subroutine, 5-56
 - to 5-58
- \$\$ALVC
 - PSECT, 7-11
 - reserved PSECT name, E-2
- \$\$ALVD
 - PSECT, 7-11
 - reserved PSECT name, E-2
- \$\$ALVI
 - PSECT, 7-11
 - reserved PSECT name, E-2
- Ancillary control processor
 - See ACP
- APR, 2-16
 - I- and D-space
 - allocation in multiuser
 - task, 9-3
 - relocatable region
 - specifying for, 5-6
 - resident common
 - system-owned, 11-11
 - resident library
 - system-owned, 11-11
 - specifying for ACP, 10-5
 - supervisor-mode, 8-2 to 8-3
- Arithmetic element
 - extended
 - specifying, 10-16
- Array declaration
 - mapped, A-10
- ASG option, 11-7
- Asterisk (*)
 - See also Autoload indicator
 - cross-reference
 - of overlaid task, 4-13
 - cross-reference listing, 10-12
- At sign (@)
 - cross-reference
 - of overlaid task, 4-13
 - cross-reference listing, 10-12 to 10-13
 - indirect file, 1-5
- Attribute
 - in .NAME directive
 - DSK, 3-28
 - GBL, 3-28
 - NODSK, 3-28
 - NOGBL, 3-28
 - program section
 - restriction, 2-3
 - save, 2-4

INDEX

- \$\$AUTO
 - PSECT, 5-52 to 5-53
 - reserved PSECT name, E-2
- AUTO module, 5-52
- AUTOL module, 5-52
- Autoload, 4-1
 - applying indicator
 - co-tree root, 4-2
 - .FCTR label name, 4-3
 - file name, 4-2
 - portions of ODL tree, 4-2
 - program section name, 4-3
 - segment name, 4-3
 - code sequence
 - conventional task, 4-5
 - error handling, 4-11 to 4-12
 - indicator, 4-2
 - efficiently placed, 4-6
 - overhead in region, 5-13
 - path loading, 4-3 to 4-4
 - specifying, 4-1
 - vector, 4-4
 - eliminating unnecessary, 4-6
 - I- and D-space, 7-9, 7-11
 - overlay, 3-20
 - vector format
 - conventional task, 4-4
 - I- and D-space, 4-4 to 4-5
- Autoloadable
 - data segment, 4-7
 - making file
 - using file name, 4-2
 - making program section, 4-3
- .BLK
 - See Program section
- Block
 - label, 2-8
- Buffer record
 - maximum size, 11-23
- Build file
 - modifying
 - to improve performance, F-1, F-5
- /CC switch, 10-7
- Checkpoint
 - area
 - in task image, B-9
 - space
 - allocating, 10-6, 10-10
- Checkpointable task
 - specifying a, 10-6, 10-10
- Circumflex ()
 - cross-reference listing, 10-12 to 10-13
 - global cross-reference
 - of an overlaid task, 4-13
- CLSTR option, 11-8 to 11-9
- Cluster library, 5-44
 - See also Library
- /CM switch, 10-8
- \$CMPAL
 - completion routine, 8-8 to 8-9
- CMPAL option, 8-9
- \$CMPCS
 - completion routine, 8-8
- CMPCS module, 8-9
- CMPRT option, 8-9, 11-10
 - use in CSM library, 8-3, 8-8
 - use in supervisor-mode library, 8-7
- /CO switch, 10-9
- Code
 - access, 2-4 to 2-5
 - allocation, 2-4 to 2-5
 - relocation, 2-4
 - scope, 2-4, 2-7
 - type, 2-5, 2-7
- Comma (,)
 - See ODL operator
- Command
 - file
 - indirect, 1-5
 - interaction with indirect, 1-5 to 1-6
 - level of indirection, 1-7
 - with ODL, 3-30
 - line
 - comments in, 1-8
 - form, 1-2
 - multi-line input, 1-3
 - option input, 1-3
 - output file interpretation, 1-3
 - terminating character, 1-3, 1-5 to 1-6
 - to build a task, 1-2
 - UFD convention, 1-9
 - sequence
 - comments in, 1-8
 - simple, 1-2
- Comment
 - in command sequence, 1-8
 - in line, 1-8
- Common, 5-1, 5-26 to 5-31
 - See also Region
 - allocation diagram, 5-20
 - assigning references, 5-23
 - building a linking task, 5-21 to 5-22
 - building and linking to, 5-14, 5-17 to 5-18
 - building and linking to a device, 5-26 to 5-31
 - device, 5-26 to 5-31
 - See also Device common
 - establishing offset in, 5-28
 - in MACRO-11
 - building and linking to, 5-17 to 5-18
 - building and linking to a
 - See Region

INDEX

- Common (Cont.)
 - installing in RSX-11M, 5-2 to 5-3, 5-20
 - installing in RSX-11M-PLUS, 5-2 to 5-3, 5-20
 - linking to region, 5-14
 - map, 5-19
 - PSECT
 - building a linking task, 5-23
 - region, 2-19
 - resident, 5-1 to 5-2
 - declaring, 11-11, 11-28
 - name block data, B-8
 - specifying a, 10-9
 - typical, 5-2
- COMMON option, 11-11 to 11-12
- Completion routine
 - \$CMPAL, 8-2, 8-8 to 8-9, 8-18
 - \$CMPCS, 8-2, 8-8
 - content of a, 8-12 to 8-14
 - CSM library, 8-13
 - definition, A-11
 - identification, A-11
 - name, A-11
 - supervisor-mode library, 8-2, 8-7
 - user-written, 8-21
- Complex relocation, A-22
 - entry, A-23
 - operation codes, A-22
- Concatenated object module
 - using to reduce overhead, F-4
- Control section
 - name, A-5
 - name entry, A-5
- Cotree, 3-31
 - and main tree
 - virtual address space, 3-40
 - global symbol resolution, 3-17
 - null root, 3-31
- ODL statement
 - from allocation diagram, 3-39 to 3-40
 - overlay, 3-34 to 3-35
 - segment
 - affecting symbol search on, 10-19
 - segment loading, 4-2
- Counter
 - location
 - definition, A-17
 - modification, A-18
- /CP switch, 10-10
- /CR switch, 10-11 to 10-13
- Cross-reference
 - global
 - of overlaid task, 4-12 to 4-14
 - listing
 - specifying a, 10-11
- CSM library
 - completion routine, 8-12 to 8-14
 - dispatching, 8-19
 - linking task
 - example of, 8-9
 - supervisor-mode, 8-7
 - building, 8-7 to 8-8
 - .STB file, 8-8
- CTRL/Z
 - effect on Task Builder, 11-4
- /DA switch, 10-14
- Data
 - adjacency in memory, 2-28
 - segment
 - autoloadable, 4-7
 - structure
 - building a, 2-1
 - overlay of a, 3-19
- Data base
 - overlay, B-15
 - I- and D-space task, B-16
- Data format
 - input
 - Task Builder, A-1
- \$\$DBTS
 - reserved PSECT name, E-2
- Debugging aid
 - including a, 10-14
- Declaration flag byte
 - symbol, A-7
- Default of switch
 - modifying, F-7 to F-11
- Descriptor
 - region, B-21
 - segment, B-18
 - window, B-20 to B-21
- Development step
 - program, 1-1
- Device
 - assignment, 11-7
 - common, 5-29
 - building and linking to a, 5-26 to 5-31
- \$\$DEVT
 - reserved PSECT name, E-2
- Directive
 - memory management
 - in mapping, 2-24
 - .NAME, 3-27
 - attributes for, 3-28
 - example of use, 3-28
 - ODL, 3-23 to 3-24
 - .END, 3-23
 - .FCTR, 3-25
 - introduction, 3-23
 - .ROOT, 3-23
 - .PSECT, 3-29
 - use of parentheses, 3-24
- Directory record
 - declare global symbol, A-2
 - end of global symbol, A-11

INDEX

- Directory record (Cont.)
 - global symbol
 - end of, A-11
 - internal symbol, A-24
 - relocation, A-12
 - Disk image, 2-8 to 2-9
 - conventional task, 7-6
 - Disk-resident
 - overlay, 3-1 to 3-2
 - loading, 4-1
 - overlay structure, 3-2
 - Displaced relocation
 - internal, A-15
 - /DL switch, 10-15
 - DSPPAT option, 11-13
 - Dump
 - See Postmortem dump
 - See Snapshot dump
 - memory, D-1
 - Dynamic region, 2-20, 5-40 to 5-43

 - /EA switch, 10-16
 - /EL switch, 10-17
 - Element
 - extended arithmetic specifying, 10-16
 - .END directive, 3-23
 - End of global symbol directory, A-11
 - End of module record, A-24
 - Error
 - exit TKB commands on, 11-4
 - handling
 - \$ALERR entry, 4-12
 - for autoload, 4-11 to 4-12
 - for manual load, 4-12
 - overlay, 4-12
 - message, H-1
 - Exclamation point (!)
 - See ODL operator
 - operator
 - See ODL operator
 - Extend Task directive
 - to improve performance, F-2
 - EXTSCT option, 11-14
 - EXTTSK option, 11-15

 - Factor
 - autoloadable
 - making first component of, 4-3
 - Fast Task Builder, G-1
 - speed of, G-2
 - supported
 - features, G-1 to G-2
 - options, G-2
 - switches, G-1
 - unsupported features, G-1 to G-2
 - .FCTR directive, 3-25
 - argument
 - library modules, 3-26
 - .FCTR directive
 - argument (Cont.)
 - library to resolve
 - references, 3-26
 - named input file, 3-25
 - PSECT name, 3-26
 - segment name, 3-26
 - arguments for, 3-25
 - use of label in, 3-25
 - .FCTR statement
 - allocation diagram
 - creating from, 3-38 to 3-39
- File
 - command
 - level of indirection, 1-7
 - declaring number of active, 11-6
 - indirect command, 1-5
 - input
 - designating as debugging aid, 10-14
 - designating as library file, 10-23
 - directing selective symbol search, 10-47 to 10-49
 - including content of in map, 10-26
 - processing to reduce overhead, F-6
 - specifying as default library, 10-15
 - library
 - declaring a, 10-23
 - making file autoloadable
 - using name, 4-2
 - map
 - printing, 1-2
 - omitting a specific output, 1-3
 - open at one time, 11-6
 - specification
 - convention for, 1-8
 - default for, 1-8
- Floating Point Processor
 - specifying, 10-18
- FMTBUF option, 11-16
- FORTTRAN
 - common block
 - in overlays, 3-19
 - manual load calling sequence, 4-9 to 4-10
 - for I- and D-space task, 4-11
 - run-time support
 - virtual program section, 5-56 to 5-57
- /FP switch, 10-18
- \$\$FSR1
 - reserved PSECT name, E-3
- .FSRPT
 - low-memory context, B-10
 - reserved global symbol, E-1

INDEX

- FTB library
 - overlaid region, 5-13
- /FU switch, 10-19
- GBLDEF option, 11-17
- GBLINC option, 11-18
- GBLPAT option, 11-19
- GBLREF option, 11-20
- GBLXCL option, 11-21
 - use in CSM library, 8-3, 8-7 to 8-8, 8-18
- Global
 - additive relocation, A-16
 - relocation, A-15
 - additive displaced, A-17
 - displaced, A-16
 - symbol
 - address of ODT SST routine, 11-24
 - ambiguously defined in overlay, 3-16
 - declaration directory record, A-2
 - directory record format, A-4
 - end of directory record, A-11
 - from the default library, 3-18
 - in auto-loadable segment, 4-4, 4-6
 - in cross-reference listing, 10-12 to 10-13
 - multiplydefined, 3-16
 - multisegment task, 3-16
 - name, A-6
 - name entry, A-6
 - overlay search sequence, 3-17
 - resolution, 2-7
 - search sequence in overlays, 3-16
 - undefined, 2-7
 - symbol resolution
 - co-tree, 3-17
 - default library, 3-18
 - in multisegment task, 3-16
- /HD switch, 10-20
- Header, 2-8
 - excluding task, 10-20
 - I- and D-space task, 7-11
 - in task image, B-10
 - task
 - fixed part, B-11
 - variable part, B-12
 - vector extension area, B-13
- High-level language
 - overlay program in, 3-40 to 3-41
- Host system, C-1
 - building a task for another system, C-1
- Host system (Cont.)
 - transferring from task, C-1
- Hyphen (-)
 - See ODL operator
- I- and D-space
 - specifying, 10-21
- I- and D-space task
 - See Task
- I/O page
 - specifying, 10-22
- /ID switch, 10-21
- Image
 - disk, 2-8 to 2-9
 - memory, 2-8 to 2-9
- Indirect command file, 1-5, 3-30
 - with ODL, 3-30
- Information record
 - text, A-11
- Input
 - data format
 - Task Builder, A-1
 - multi-line
 - to the Task Builder, 1-3
- Internal displaced relocation, A-15
- Internal symbol
 - directory record, A-24
 - name entry, A-5
- \$\$IOB1
 - reserved PSECT name, E-3
- \$\$IOB2
 - reserved PSECT name, E-3
- /IP switch, 10-22
- Label block, 2-8
 - group, B-1
- Language
 - high-level
 - overlay program in, 3-40
- /LB switch, 10-23 to 10-24
- LBLDF\$ macro, B-1
- /LI switch, 10-25
- LIBR
 - linking to region, 5-14
 - option, 11-11 to 11-12
- Library
 - building a, 5-14
 - cluster, 5-44
 - building, 5-44, 5-48 to 5-51
 - building example, 5-48 to 5-51
 - building rule 1: overlays, 5-45
 - building rule 2: references, 5-46
 - building rule 3: .STB file, 5-47
 - building rule 4: stack, 5-47
 - building rule 5: PIC, 5-47

INDEX

Library

- cluster (Cont.)
 - building rule 6: traps, 5-47
 - building rule summary, 5-44
 - examples, 5-47
 - overlay run-time support, 5-51 to 5-53
 - resolving interlibrary references, 5-49
- declaring a, 10-23
- default
 - controlling symbol search, 10-15, 10-19
 - global symbol resolution, 3-18
 - specifying a, 10-15
- extending a, 10-17
- file
 - declaring a, 10-23
- FTB
 - overlaid region, 5-13
- linking resident to supervisor-mode, 8-21
- modules
 - .FCTR directive, 3-26
- object module
 - placing in overlay structure, 10-24
- old
 - overlaid region, 5-13
- region, 2-20
 - specifying, 10-25
- relocation
 - resident, A-23
- resident, 5-1, 5-3
 - building and linking to a, 5-14, 5-31 to 5-38
 - data in task image, B-4 to B-6
 - declaring, 11-11, 11-28
 - label block 0, B-7
 - label block 1, B-9
 - label block 2, B-9
 - label block 3, B-9
 - name block data, B-8
 - relocation, A-23
 - search, 10-23
- resolving references
 - .FCTR directive, 3-26
- restriction in I- and D-space task, 7-9
- supervisor-mode, 2-24, 8-1
 - building, 8-7
 - building linking task, 8-17 to 8-18
 - building the referencing task, 8-7
 - building with relevant options, 8-3, 8-7, 8-17 to 8-18
 - completion routine, 8-2, 8-7, 8-21

Library

- supervisor-mode (Cont.)
 - converting SCAL to CSM, 8-20
 - data in, 8-3
 - definition, 8-1
 - example of, 8-10 to 8-17
 - linking, 8-7
 - linking a resident, 8-20
 - linking to, 8-21
 - linking to SYSLIB, 8-2, 8-7 to 8-9, 8-18
 - linking with relevant options, 8-3, 8-8, 8-18
 - mapping of, 8-3
 - method of mode switching, 8-1, 8-19
 - mode switching, 8-1, 8-7
 - mode switching compared, 8-1
 - mode-switching vector, 8-1
 - \$MSDS directive, 8-2
 - \$MSDS directive restriction, 8-2
 - multiple, 8-20
 - overlaid, 8-21 to 8-22
 - overlay restriction, 8-21 to 8-22
 - parameter passing, 8-2
 - restrictions on contents, 8-2
 - using as resident, 8-20
 - using system supplied vector, 8-19
 - with I- and D-space task, 8-3
 - your own completion routines in, 8-21
 - your own vector in, 8-21
 - supervisor-mode mapping with conventional task, 8-4 to 8-5
 - with I- and D-space task, 8-6
- SYSLIB
 - replacing as default, 10-23
- Linking
 - module, 2-3
- Listing
 - global cross-reference generating as, 10-11
 - wide
 - specifying a, 10-51
- \$\$LOAD
 - PSECT, 5-53
 - reserved PSECT name, E-3
- LOAD module, 5-53
- \$LOAD routine
 - in manual load, 4-7
- Loading
 - asynchronous example of, 4-10
 - mechanism
 - overlay, 3-16

INDEX

- Location counter
 - definition, A-17
 - modification, A-18
- Logical
 - address, 2-14
 - address space, 2-14
 - unit number
 - assigning physical device to a, 11-7
 - unit table entry, B-9, B-14
 - units
 - number of, 11-39
- Low-memory context, B-10
- LUN
 - See Logical unit number
- /MA switch, 10-26
- Macro
 - SNAP\$, D-6
 - SNPBK\$, D-6
 - SNPDF\$, D-6
- MACRO-11 calling sequence
 - manual load, 4-7
 - I- and D-space tasks, 4-8 to 4-9
- Manual load
 - calling sequence, 4-7 to 4-8
 - error handling, 4-12
 - FORTRAN calling sequence, 4-9 to 4-10
 - for I- and D-space task, 4-11
 - MACRO-11 calling sequence, 4-7
 - I- and D-space tasks, 4-8 to 4-9
- Map
 - common, 5-19
 - file
 - adding cross-reference to a, 10-11
 - content, 10-37 to 10-43
 - description, 10-37 to 10-43
 - example, 10-37 to 10-43
 - general, 10-37 to 10-43
 - inhibiting spooling of a, 10-45
 - printing, 1-2
 - specifying, 10-26
 - including SYSLIB contribution, 10-26
 - multiuser task, 9-7
 - overlaid I- and D-space task, 7-12 to 7-16
 - privileged task, 6-10 to 6-11
 - region, 5-19
 - resident region
 - including symbol definition, 10-26
 - short
 - specifying a, 10-37 to 10-43
 - spooling to print, 10-45
- Map (Cont.)
 - task
 - I- and D-space, 7-9
 - linked to a common, 5-24
- Mapped
 - array
 - area, 5-56, 5-58
 - declaration, A-10
 - declaration entry, A-10
 - region
 - declaring address window, 11-41
 - system, 2-14, 2-18
- Mapping
 - concept, 2-22
 - conventional task
 - supervisor-mode library, 2-26 to 2-27
 - conventional task linked to region
 - I- and D-space system, 7-3
 - I- and D-space task, 7-3 to 7-4
 - in I- and D-space system, 7-4
 - in supervisor-mode, 2-24
 - task, 2-15
- MAXBUF option, 11-23
- .MBLUN
 - reserved global symbol, E-1
- Memory
 - allocation
 - I- and D-space task, 7-9
 - allocation file
 - See Map file
 - dump, D-1
 - image, 2-8 to 2-9
 - layout
 - unmapped system, 2-16
 - management
 - specifying for target system, 10-27
 - physical
 - disk-resident overlay, 3-3 to 3-4
 - memory-resident overlay, 3-6
 - overlay, 3-5, 3-33 to 3-35
 - reducing to build a task, F-4
 - reducing usage of, 3-1
 - resident overlay structure, 3-6
 - saving
 - overlaid task, 3-9 to 3-10
 - virtual allocation
 - I- and D-space task, 7-10
 - Memory management
 - use by task, 2-15 to 2-16
 - Memory Management Unit, 2-14
 - Memory-resident
 - overlay
 - loading, 4-1
 - overlay, 3-1

INDEX

- Message
 - diagnostic
 - eliminating, 10-30
 - error, H-1
 - virtual memory system, F-5
 - inhibiting system queuing of, 10-35
- /MM switch, 10-27
- Mode
 - compatibility
 - in a task, 10-8
- Mode-switching
 - to supervisor-mode, 8-8, 8-12 to 8-13
- vector
 - supervisor library, 8-1, 8-19
- Module
 - extracting from library, 10-23
 - linking, 2-3
 - name, A-4
 - name entry format, A-4
 - object
 - extracting by name, F-4
 - linking, 2-1
 - placing in segment
 - reducing overhead, F-4
 - record
 - end of, A-24
- .MOLUN
 - reserved global symbol, E-1
- /MP switch, 10-28
- \$\$MRKS
 - PSECT, 5-52
 - reserved PSECT name, E-3
- /MU switch, 10-29
- Multi-line input, 1-3
- Multiple-tree
 - example, 3-32
 - structure, 3-31 to 3-32
 - defining a, 3-31 to 3-32
- Multisegment task
 - See Overlay
- Multiuser task, 2-28, 9-1
 - as an overlaid task, 9-2
 - building a, 9-5
 - declaring read-only partition
 - for, 11-33
 - defined, 9-1
 - disk image, 9-2
 - example, 9-6 to 9-7
 - example map, 9-7
 - I- and D-space, 9-3
 - APR allocation, 9-3
 - PSECT allocation, 9-4
 - PSECT in, 9-3
 - windows, 9-5
 - program section allocation, 9-1 to 9-2
 - specifying a, 10-29
 - TKB command sequence, 9-7
- Multiuser task (Cont.)
 - window block assignment, 9-1, 9-3
- N.OVPT
 - low-memory context, B-10
 - reserved global symbol, E-1
- Name
 - global symbol, A-6
- .NAME directive, 3-27
 - attribute
 - DSK, 3-28
 - GBL, 3-28
 - NODSK, 3-28
 - NOGBL, 3-28
 - example of use of, 3-28
- .NLUNS
 - reserved global symbol, E-1
- /NM switch, 10-30
- .NOVLY
 - reserved global symbol, E-1
- .NSTBL
 - reserved global symbol, E-1
- Null
 - root, 3-31
 - in ODL, 3-31
 - segment
 - in ODL, 3-31
- Number sign (#)
 - in cross-reference listing, 10-12 to 10-13
- \$\$OBF1
 - reserved PSECT name, E-3
- \$\$OBF2
 - reserved PSECT name, E-3
- Object code
 - patching, 11-5
- Object module
 - concatenating, 10-7
 - content of, A-1
 - format, A-3
 - linking, 2-1
 - overriding definition in, 11-17
 - relocatable, 2-2
 - selective global symbol
 - using /SS to include, 10-47 to 10-49
- Object Time System
 - usage to extend record buffer, 11-16
- ODL
 - autoload indicator, 4-2
 - directive, 3-23 to 3-24
 - .END, 3-23
 - example use of .NAME, 3-28
 - .FCTR, 3-25
 - introduction, 3-23
 - .NAME, 3-27
 - .NAME attributes, 3-28
 - .PSECT, 3-29
 - .ROOT, 3-23

INDEX

- ODL (Cont.)
 - efficiently placing in
 - autoload indicator, 4-6
 - enabling operator
 - memory-resident overlay, 10-34
 - file
 - declaring a, 10-28
 - multiple-tree
 - defining structure, 3-30
 - example, 3-31 to 3-32
 - structure, 3-30
 - operator
 - , (comma), 3-24
 - ! (exclamation point), 3-24, 3-26 to 3-27
 - (hyphen), 3-24
 - introduction, 3-24
 - summary, 3-49 to 3-52
 - using indirect file with, 3-30
- ODL file
 - creating
 - start of procedure, 3-36
 - with allocation diagram, 3-35 to 3-36
 - .FCTR statement
 - creating from allocation diagram, 3-38 to 3-39
 - .ROOT statement
 - creating from allocation diagram, 3-37
 - virtual address space
 - in allocation diagram, 3-36
- ODL statement
 - co-tree
 - from allocation diagram, 3-39 to 3-40
- ODT vector, 11-24
- .ODTL1
 - reserved global symbol, E-1
- .ODTL2
 - reserved global symbol, E-1
- ODTV option, 11-24
- Operator
 - ODL, 3-24
 - , (comma), 3-24
 - ! (exclamation point), 3-24, 3-26 to 3-27
 - (hyphen), 3-24
 - introduction, 3-24
- Option
 - category of, 11-1
 - general form of, 1-4
 - input
 - in command line, 1-3
 - linking to region, 5-14
 - separation of argument list, 1-4 to 1-5
 - summary, 11-2
 - Task Builder, 1-4, 11-1
 - ABORT, 11-4
 - ABSPAT, 11-5
- Option
 - Task Builder (Cont.)
 - ACTFIL, 11-6
 - ASG, 11-7
 - CLSTR, 11-8 to 11-9
 - CMPRT, 11-10
 - COMMON or LIBR, 11-11 to 11-12
 - DSPPAT, 11-13
 - EXTSCT, 11-14
 - EXTTSK, 11-15
 - FMTBUF, 11-16
 - GBLDEF, 11-17
 - GBLINC, 11-18
 - GBLPAT, 11-19
 - GBLREF, 11-20
 - GBLXCL, 11-21
 - MAXBUF, 11-23
 - ODTV, 11-24
 - PAR, 11-25 to 11-26
 - PRI, 11-27
 - RESCOM or RESLIB, 11-28 to 11-29
 - RESSUP, 11-31 to 11-32
 - ROPAR, 11-33
 - STACK, 11-34
 - SUPLIB, 11-35
 - TASK, 11-36
 - TSKV, 11-37
 - UIC, 11-38
 - UNITS, 11-39
 - VSECT, 11-40
 - WNDWS, 11-41
 - Option summary, 11-3
 - OTS
 - usage to extend record buffer, 11-16
 - \$OTSV
 - low-memory context, B-10
 - reserved global symbol, E-2
 - Output files
 - omitting specific, 1-3
 - OVCTC module, 5-52
 - OVCTL module, 5-52
 - OVCTR module, 5-52
 - OVDAT module, 5-53
 - \$\$OVDT
 - PSECT, 5-53
 - reserved PSECT name, E-3
- Overlay, 2-10
 - allocation diagram
 - creating ODL file with, 3-35 to 3-40
 - autoload vector in, 3-20
 - building an, 3-41 to 3-48
 - building memory-resident for region, 5-9
 - capability of an, 3-1
 - choosing a memory-resident, 3-14
 - co-tree, 3-34 to 3-35
 - data base, B-15
 - I- and D-space task, B-16

INDEX

- Overlay (Cont.)
 - data structure, 3-19
 - linked into root, 3-20
 - defining a multiple-tree, 3-31 to 3-32
 - description
 - effect on performance, F-4
 - disk-resident, 3-1 to 3-2
 - defined, 4-1
 - effect of
 - on physical memory, 3-5
 - on virtual address space, 3-5
 - effect of virtual address space
 - on disk-resident, 3-1 to 3-4
 - effect on memory
 - of a disk-resident, 3-2
 - of disk-resident, 3-1, 3-3 to 3-4
 - effect on physical memory
 - of memory-resident, 3-6
 - effect on virtual address space
 - of a memory-resident, 3-6
 - error handling, 4-12
 - example of building a, 3-41 to 3-48
 - I- and D-space task
 - disk image, 7-8
 - in I- and D-space task, 3-21, 7-5
 - PSECT, 7-6
 - loading
 - asynchronously, 4-10
 - disk-resident, 4-1
 - mechanism, 3-16
 - memory-resident, 4-1
 - methods, 4-1
 - synchronously, 4-8 to 4-9
 - memory-resident, 3-1, 3-6
 - conserving physical memory, 3-14
 - defined, 4-1
 - physical memory usage, 3-6
 - region, 5-9
 - virtual address space, 3-14
 - multiuser task, 9-2
 - of program
 - in high-level language, 3-40 to 3-41
 - operator
 - enabling recognition of, 10-34
 - suppression of a
 - memory-resident, 10-34
 - path loading in an, 4-4
 - physical memory, 3-33 to 3-35
 - program section
 - specifying, 3-19
 - region
 - autoload vector, 5-11
- Overlay
 - region (Cont.)
 - building, 5-9
 - building option, 5-10 to 5-11
 - descriptor in, 3-20
 - example of building, 5-10
 - global symbols in .STB file, 5-11
 - resolving symbol, 5-10 to 5-11
 - .STB file, 5-11
 - vectors in I- and D-space task, 5-11
 - region restrictions, 5-12
 - root segment structure, 3-21
 - run-time, 3-19
 - comparison of sizes in routine, 4-16
 - module sizes, 5-52 to 5-53
 - routine, 3-19, 4-16
 - size of routine, 4-16
 - support requirements, 5-12, 5-51 to 5-53
 - use of routine, 4-14 to 4-15
 - segment
 - alignment, 10-8
 - arrangement, 3-15
 - descriptor in, 3-20
 - processing order, 3-17
 - symbol processing, 3-17
 - structure
 - considerations in creating, 3-1
 - most effective, 3-2
 - multiple-tree, 3-18
 - multiply defined global symbol, 3-16
 - specifying library search in a, 10-23
 - task, 3-7, 3-9
 - global cross-reference of, 4-12 to 4-14
 - segment calls, 3-10
 - virtual address space, 3-10
 - tree, 3-15
 - calling segments in an, 4-6
 - virtual
 - address space, 3-33 to 3-35
 - address space and memory, 3-11 to 3-14
 - window
 - block, 3-49
 - descriptor in, 3-20
 - written in high-level language, 3-40 to 3-41
- Overlay Description Language
 - introduction, 3-15
- Overlay Description Operator
 - See ODL
- Overlay structure
 - global symbol

INDEX

- Overlay structure
 - global symbol (Cont.)
 - ambiguously defined, 3-16
- Overlay tree
 - for I- and D-space task, 7-7
- Overlay tree diagram
 - virtual address space, 3-24
- OVIDC module, 5-52
- OVIDL module, 5-52
- OVIDR module, 5-52
- \$\$OVRS
 - reserved PSECT name, E-3
- Page Address Register
 - See PAR
- Page Description Register
 - See PDR
- PAR, 2-16
 - option, 11-25
- PAR option, 11-26
 - building region, 5-2
- Parentheses
 - use of
 - in ODL, 3-24
- Partition
 - declaring, 11-25
 - in region, 5-28
 - naming for target system, C-1
 - option, 11-25
 - requirement
 - shared region, 5-3
 - requirements
 - region, 5-2
 - size
 - for TKB, F-2
 - specifying for region, 5-28
- Patch
 - D-space, 11-13
 - declaring an object level, 11-5
 - global relative, 11-19
- Path loading, 4-3 to 4-4
 - See also Overlays
 - example of, 4-4
 - in autoload, 4-3 to 4-4
- \$\$PDLS
 - PSECT, 5-52
 - reserved PSECT name, E-3
- PDR, 2-16
- Performance
 - improving TKB, F-1
- Physical
 - address, 2-13
 - mapped system, 2-17
- /PI switch, 10-31
- /PM switch, 10-32
- PMD task
 - installation for timely
 - operation, D-1
- Postmortem dump, D-1
 - content, D-3 to D-5
 - example, D-2
 - sample, D-3 to D-5
- Postmortem dump (Cont.)
 - specifying a, 10-32
- /PR switch, 6-2, 10-33
- /PR:0 privileged task, 6-2, 6-4 to 6-5
 - uses of, 6-5
- /PR:4 privileged task, 6-2, 6-4 to 6-5
 - uses of, 6-5
- /PR:5 privileged task, 6-2, 6-4 to 6-6
 - uses of, 6-5
- PRI option, 11-27
- Priority
 - task
 - declaring, 11-27
- Privileged and nonprivileged
 - task
 - distinction between, 6-1
- Privileged task, 2-25, 6-1
 - accessing
 - Executive with a, 6-4
 - I/O page with a, 6-4
 - building a, 6-6 to 6-11
 - comparison of nonprivileged and, 6-1
 - hazards of a, 6-1 to 6-2
 - in a mapped system, 6-1
 - logging off from, 6-1
 - MAC command sequence, 6-9
 - map of, 6-10 to 6-11
 - mapping of, 6-2 to 6-3
 - /PR:0, 6-2, 6-4 to 6-5
 - /PR:4, 6-2, 6-4 to 6-5
 - /PR:5, 6-2, 6-4 to 6-6
 - processor trap in a, 6-2
 - specifying, 2-25
 - specifying a, 6-2
 - TKB command sequence, 6-10
 - to examine unit control block, 6-6 to 6-11
- Program
 - development step, 1-1
 - limit, A-18
 - section, 5-53
 - additive displaced
 - relocation, A-21
 - additive relocation, A-20
 - adjacency requirement for, 10-46
 - allocation, 2-2
 - attribute, 2-4
 - attribute restriction, 2-3
 - blank (.BLK), 2-3
 - creation, 2-3
 - displaced relocation, A-19
 - element, 2-2
 - extension of, 11-14
 - in shared region, 5-7
 - making autoloadable, 4-3

INDEX

- Program
 - section (Cont.)
 - name, 2-3
 - name conflict, 5-7, 5-39 to 5-40
 - naming restriction, 5-7
 - ordering, 5-7, 10-36, 10-46
 - overlay allocation, 3-19
 - relocation, A-19
 - resolution of, 3-19
 - resolving names, 5-7, 5-39 to 5-40
 - resolving names in region and task, 5-16
 - save attribute, 2-7
 - segregating, 10-36, 10-46
 - sequential ordering of, 10-36, 10-46
 - space allocation, 2-5 to 2-6
 - specifying explicitly in overlay, 3-29
 - specifying in overlay, 3-19
 - virtual, 5-53
 - section name, A-7
 - applying autoloading indicator, 4-3
 - section name entry, A-8
 - section name flag byte, A-8 to A-9
 - version
 - identification, A-10
- Program section
 - \$\$ALVC, 7-11
 - \$\$ALVD, 7-11
 - \$\$ALVI, 7-11
 - I- and D-space task
 - overlay, 7-6
 - multiuser task
 - I- and D-space, 9-3
 - named
 - in region, 5-13
 - virtual
 - allocating physical memory to a, 5-56
 - attaching virtual attribute to a, 5-56
 - building a task using, 5-58 to 5-61
 - creating a, 5-56, 5-58 to 5-61
 - FORTAN run-time support
 - for, 5-56 to 5-57
 - option usage, 5-55
 - specifying, 11-40
 - specifying base address for a, 5-56
 - specifying length, 5-54
 - specifying physical size, 5-54
 - specifying window size, 5-54
 - support for a, 5-56 to 5-57
- PSECT
 - See also Program section
 - See Program section
 - .PSECT directive, 3-29
 - PSECT name
 - .FCTR directive
 - argument, 3-26
 - .PTLUN
 - reserved global symbol, E-2
- \$\$RDSG
 - PSECT, 5-52
 - reserved PSECT name, E-3
- Region, 2-18
 - absolute, 5-9
 - building precautions, 5-7
 - mapping for, 5-8
 - mapping of, 5-7
 - specifying an, 5-7
 - symbol definition file, 5-9
 - allocation
 - diagram, 5-20
 - of window block for, 5-25
- APR
 - specifying, 5-6
- assigning references, 5-23
- building
 - a linking task, 5-21 to 5-22
 - and linking to a, 5-14, 5-17 to 5-18
 - interaction of /CO/LI/PI switch, 5-3
 - options
 - in an overlaid, 5-10
 - options in an overlaid, 5-11
 - use of /CO/LI/PI switch, 5-3
 - with PAR option, 5-2
- common, 2-19 to 2-20
- descriptor, B-21
- descriptor in overlay, 3-20
- dynamic, 2-20, 5-40 to 5-43
 - building a task that
 - creates a, 5-40 to 5-43
- installing in RSX-11M, 5-2 to 5-3, 5-20
- installing in RSX-11M-PLUS, 5-2 to 5-3, 5-20
- library, 2-20
- linked to a region, 5-26
- linked with a region, 5-25
- linking to a, 5-13 to 5-14, 5-26
- map, 5-19
- mapping of an absolute, 5-7
- memory-resident overlaid, 5-9
 - building, 5-9
 - example of building, 5-10
- number of, 5-17
- options for building, 5-17 to 5-18

INDEX

- Region (Cont.)
 - options for linking to, 5-14
 - overlaid, 5-9
 - autoload call overhead, 5-13
 - autoload in, 5-13
 - autoload vector, 5-11
 - FTB and old libraries, 5-13
 - global symbols in .STB file, 5-11
 - I- and D-space task vectors, 5-11
 - named program section, 5-13
 - run-time support for, 5-12 to 5-13
 - .STB file, 5-11
 - partition, 5-28
 - requirements, 5-2
 - procedure for building a, 5-17 to 5-18
 - program section in, 5-7
 - PSECT
 - building a linking task, 5-23
 - relocatable, 5-5
 - mapping, 5-5
 - specifying, 5-5
 - specifying APR for, 5-6
 - .STB file for a, 5-9, 5-14, 5-16
 - resident relocatable, 5-5
 - resolving PSECT names, 5-39 to 5-40
 - shared, 2-20, 5-1
 - allocation of window block for, 5-14
 - autoload vectors in, 5-12
 - defined, 5-1
 - initializing window block for, 5-14
 - installing, 5-16
 - partition requirement, 5-3
 - resolving PSECT names, 5-16
 - restrictions for overlaid, 5-12
 - symbol definition file, 5-16
 - use of /CO/LI/PI switches, 5-9
 - windows in, 5-14
 - size of, 5-17
 - specifying
 - as position independent, 10-9, 10-31
 - partition for, 5-28
 - .STB file, 5-4 to 5-5, 5-14, 5-16
 - for an absolute, 5-7, 5-9, 5-12
 - using /CO/LI/PI switches, 5-4
 - symbol definition file, 5-9
 - task, 2-18 to 2-20
- Region (Cont.)
 - task building options, 5-13 to 5-14
 - type of access to, 5-14
 - use of /CO/LI/PI switches, 5-5
 - window, 5-15, 5-25
 - with linked task
 - in I- and D-space system, 7-3
- Relocatable region
 - See Region
- Relocation
 - code, 2-4
 - complex, A-22
 - entry, A-23
 - operation codes, A-22
 - directory command byte, A-13
 - directory record, A-12
 - entries, A-12 to A-13
 - format, A-14
 - displaced
 - global, A-16
 - global, A-15
 - additive, A-16
 - global additive displaced, A-17
 - internal, A-14
 - displaced, A-15
 - library
 - resident, A-23
 - program section, A-19
 - additive, A-20
 - additive displaced, A-21
 - displaced, A-19
 - resident library, A-23
- RESCOM
 - linking to region, 5-14
 - option, 11-28 to 11-29
- Reserved symbol
 - for the Task Builder, E-1
- Resident
 - common
 - name block data, B-8
 - library
 - name block data, B-8
 - relocation, A-23
 - memory
 - for TKB performance, F-3
 - overlay operator
 - enabling recognition of, 10-34
 - region
 - using to reduce overhead, F-4
- Resident region
 - map file
 - including symbol definition in, 10-26
- RESLIB
 - linking to region, 5-14
 - option, 11-28 to 11-29

INDEX

- RESLIB
 - option (Cont.)
 - in supervisor-mode library, 8-20 to 8-21
- RESSUP
 - option, 8-9, 11-31 to 11-32
 - use in CSM library, 8-3, 8-8, 8-19 to 8-21
 - use in supervisor-mode library, 8-7
- Restarting TKB option, 11-4
- \$\$RGDS
 - reserved PSECT name, E-3
- RLSCT FORTRAN subroutine, 5-56, 5-58
- /RO switch, 10-34
- Root
 - in a co-tree, 3-31
 - null, 3-31
 - in ODL, 3-31
 - structure
 - overlay, 3-21
- .ROOT directive, 3-23
- .ROOT statement
 - allocation diagram
 - creating from, 3-37
- ROPAR option, 11-33
- \$\$RTQ
 - PSECT, 5-52 to 5-53
 - reserved PSECT name, E-3
- \$\$RTR
 - PSECT, 5-52 to 5-53
 - reserved PSECT name, E-4
- \$\$RTS
 - PSECT, 5-53
 - reserved PSECT name, E-4
- Run-time support
 - overlaid region, 5-12 to 5-13
 - autoload, 5-13
 - autoload call overhead, 5-13
 - FTB and old libraries, 5-13
 - named program sections, 5-13
- Save attribute, 2-4
- SCAL to CSM library
 - converting, 8-20
- Scope-code, 2-4, 2-7
- /SE switch, 10-35
- Segment
 - autoloadable
 - data, 4-7
 - global symbol in, 4-4
 - call, 3-10
 - to up-tree, 4-6
 - definition of a, 3-1
 - descriptor, B-19
 - I- and D-space task, 7-11
 - in overlay, 3-20
 - limiting number
 - reducing overhead, F-4
 - loading
 - Segment
 - loading (Cont.)
 - as part of co-tree, 4-2
 - when called, 4-4, 4-6
 - making autoloadable, 4-3
 - mapping, 2-10
 - disk-resident, 2-11
 - memory-resident, 2-12
 - multiple, 3-5
 - global symbol in, 3-16
 - global symbol resolution, 3-17
 - symbol resolution, 3-16
 - name
 - applying autoload indicator to, 4-3
 - .FCTR directive argument, 3-26
 - null
 - in ODL, 3-31
 - overlay
 - arrangement, 3-15
 - root structure, 3-21
 - symbol processing, 3-17
 - processing order, 3-17
 - single, 3-4, 3-8
 - up-tree
 - I- and D-space, 3-22
 - Semicolon (;), 1-8
 - SEND directive
 - enabling for your task, 10-35
 - /SG switch, 10-36
 - \$\$SGDO
 - PSECT, 5-53
 - reserved PSECT name, E-4
 - \$\$SGD1
 - reserved PSECT name, E-4
 - \$\$SGD2
 - PSECT, 5-53
 - reserved PSECT name, E-4
 - /SH switch, 10-37
 - Shared region, 5-1
 - See Region
 - Single segment task, 3-8
 - /SL switch, 10-44
 - Slash
 - double (//), 1-6
 - single (/), 1-5 to 1-6
 - Slow TKB
 - to improve overhead, F-11
 - \$\$SLVC
 - reserved PSECT name, E-4
 - SNAP\$ macro, D-6
 - format of, D-8
 - Snapshot dump, D-5 to D-6
 - example of, D-10 to D-12
 - format of macro for creating, D-7
 - SNPBK\$ macro, D-6
 - format of, D-7
 - SNPDF\$ macro, D-6
 - /SP switch, 10-45
 - /SQ switch, 10-46

INDEX

- /SS switch
 - symbol definition file
 - reducing overhead, F-4
- SST vector address
 - declaring, 11-27
- Stack
 - declaring size, 11-34
 - supervisor-mode, 8-20
- STACK option, 11-34
- .STB file, 5-9
 - absolute region, 5-7, 5-9
 - content of a, 5-5
 - excluding symbol from, 5-11
 - for region, 5-4 to 5-5, 5-14, 5-16
 - I- and D-space task, 7-9
 - including symbol in, 5-11
 - interaction of /CO/LI/PI switches, 5-4
 - overlaid region, 5-11 to 5-12
 - global symbols in, 5-11
 - program sections, 5-16
 - program sections in, 5-5
 - relocatable region, 5-9, 5-16
 - use of /CO/LI/PI switches, 5-5, 5-9
- Structure
 - in TKB
 - size of, F-3
- .SUML1
 - reserved global symbol, E-2
- Supervisor-mode, 2-24
 - library, 2-24
 - See Library
 - mapping, 2-24, 2-26 to 2-27
 - mode switching, 8-8, 8-12 to 8-13
 - stack, 8-20
- SUPLIB
 - option, 8-9, 11-35
 - use in CSM library, 8-3, 8-19
 - use in supervisor-mode library, 8-7
- Switch
 - conflict in, 10-1
 - default
 - modifying, F-7 to F-11
 - modifying default of, F-7 to F-11
 - summary, 10-2 to 10-4
 - syntax, 10-1
 - Task Builder, 10-1
 - /AC[:n], 10-5
 - /AL, 10-6
 - /CC, 10-7
 - /CM, 10-8
 - /CO, 10-9
 - /CP, 10-10
 - /CR, 10-11 to 10-13
 - /DA, 10-14
 - /DL, 10-15
 - /EA, 10-16
- Switch
 - Task Builder (Cont.)
 - /EL, 10-17
 - /FP, 10-18
 - /FU, 10-19
 - /HD, 10-20
 - /ID, 10-21
 - /IP, 10-22
 - /LB, 10-23 to 10-24
 - /LI, 10-25
 - /MA, 10-26
 - /MM, 10-27
 - /MP, 10-28
 - /MU, 10-29
 - /NM, 10-30
 - /PI, 10-31
 - /PM, 10-32
 - /PR[:n], 10-33
 - /RO, 10-34
 - /SE, 10-35
 - /SG, 10-36
 - /SH, 10-37
 - /SL, 10-44
 - /SP, 10-45
 - /SQ, 10-46
 - /SS, 10-47 to 10-49
 - /TR, 10-50
 - /WI, 10-51
 - /XH, 10-52
 - /XT[:n], 10-53
- Symbol
 - affecting search for, 2-7
 - declaration flag byte, A-7
 - definition file
 - excluding symbol from a, 11-21
 - for system-owned region, 11-11
 - for user-owned region, 11-29
 - including symbols, 11-18
 - reducing overhead, F-4
 - directory record
 - declare global, A-2
 - end of global, A-11
 - internal, A-24
 - full search in overlays
 - specifying, 10-19
 - global
 - address of ODT SST routine, 11-24
 - ambiguously defined in overlay, 3-16
 - declaring definition in a task, 11-17
 - default library resolution, 3-18
 - directory record, A-2
 - directory record format, A-4
 - end of directory record, A-11
 - excluding in a task, 11-21

INDEX

- Symbol
 - global (Cont.)
 - from the default library, 3-18
 - in autoloading segment, 4-4, 4-6
 - in cross-reference listing, 10-12 to 10-13
 - including in a task, 11-18
 - multiplydefined, 3-16
 - multisegment task, 3-16
 - name, A-6
 - name entry, A-6
 - overlaid region .STB file, 5-11
 - overlay search sequence, 3-17
 - resolution, 2-7, 3-17
 - resolution in co-tree, 3-17
 - resolution in multisegment task, 3-16
 - undefined, 2-7
 - in cross-reference listing, 10-12
 - internal
 - autoloadable library, A-26 to A-27
 - end-of-module, A-30
 - end-of-module format, A-32
 - global, A-28
 - internal symbol name, A-30
 - internal symbol name format, A-31
 - line-number, A-29 to A-30
 - literal record, A-30
 - literal record format, A-31
 - module name, A-27 to A-28
 - overall format, A-24 to A-25
 - PC correlation, A-29 to A-30
 - PSECT item, A-29
 - relocatable/relocated, A-27
 - start-of-segment, A-25 to A-26
 - task identification, A-26
 - TKB generated, A-25
 - name entry
 - internal, A-5
 - number of processed for performance, F-3
 - reserved for the Task Builder, E-1
 - resolving
 - overlay region, 5-11
 - search
 - selective, 10-47 to 10-49
- Symbol definition
 - SYSLIB.OLB, 2-7
- Symbol definition file
 - See also .STB file
 - supervisor-mode library
 - system-owned, 11-35
 - Symbol definition file
 - supervisor-mode library (Cont.)
 - user-owned, 11-31
- Syntax rule
 - summary of, 1-10
- SYSLIB
 - including contribution in map, 10-26
 - linking to
 - by supervisor-mode libraries, 8-2, 8-7 to 8-9, 8-18
 - replacing as default, 10-23
- SYSLIB.OLB
 - for symbol definition, 2-7
- System
 - host and target, C-1
 - mapped, 2-14
 - physical and virtual space, 2-15
 - object module library, 2-2
 - See also Library
 - See also SYSLIB.OLB
 - target
 - memory management, 10-27
 - unmapped, 2-14
- System-controlled partition
 - extending memory for task in, 11-15
- T-bit trace trap, 10-50
- Table storage, F-2
 - memory for
 - to improve performance, F-2
- Target system, C-1
 - transferring task to a, C-1
- Task
 - access
 - system-owned common or library, 11-11
 - system-owned supervisor-mode library, 11-35
 - user-owned common, 11-28
 - user-owned library, 11-28
 - user-owned supervisor-mode library, 11-31
 - active files
 - declaring number of, 11-6
 - additional memory for, 11-15
 - address windows
 - declaring an additional, 11-41
 - ancillary control processor
 - specifying as, 10-5
 - assigning physical device to LUN, 11-7
 - attaching slave attribute to, 10-44
 - building for target system, C-1
 - changing name of, 11-36

INDEX

- Task (Cont.)
- checkpointable
 - specifying, 10-10
 - command line to build a, 1-2
 - comparison
 - conventional and I- and D-space, 7-2
 - completion routine for a, 11-10
 - conventional
 - autoload vector, B-17
 - disk image, 7-6
 - mapping compared to I- and D-space task, 7-2
 - creating a dynamic region, 5-40
 - creating multiuser, 10-29
 - D-space
 - overlay structure, 3-21
 - data
 - in task image, B-4 to B-6
 - needed by system to install, B-1
 - declaring
 - execution priority for, 11-27
 - maximum stack size of, 11-34
 - number of LUNs for, 11-39
 - object-level patch for, 11-5
 - ODT SST vector in, 11-24
 - disk image, 2-8
 - enabling
 - postmortem Dump for, 10-32
 - T-bit trace trapping in, 10-50
 - extending
 - a program section in a, 11-14
 - memory of, 11-15
 - to partition length, 11-15
 - external header
 - specifying, 10-52
 - floating point processor in
 - specifying, 10-18
 - format buffer
 - declaring length of, 11-16
 - global relative patch
 - declaring, 11-19
 - global symbol
 - excluding a, 11-21
 - including in, 11-18
 - global symbol definition
 - declaring a, 11-17
 - global symbol reference
 - declaring a, 11-20
 - header, 2-8
 - allocating additional (checkpoint) space in, 10-6
 - checkpoint area within, B-9
- Task
- header (Cont.)
 - controlling creation of, 10-20
 - fixed part, B-11 to B-12
 - I- and D-space, 7-11
 - space for EAE context, 10-16
 - space for floating-point context, 10-18
 - host to target system
 - example of transferring, C-2
 - I- and D-space, 2-28, 7-1
 - autoload vector, 7-9, B-17 to B-18
 - differing from conventional task, 2-28
 - manual load calling
 - sequence, 4-8 to 4-9, 4-11
 - map, 7-9
 - mapped in I- and D-space system, 7-4
 - mapping, 7-3
 - mapping summary of, 7-2
 - memory allocation, 7-9
 - overlaid, 7-5
 - overlay region vector, 5-11
 - overlay structure, 3-21
 - patching, 11-13
 - PSECT in overlay, 7-6
 - simplified mapping, 2-29
 - specifying, 10-21
 - .STB file, 7-9
 - with up-tree segment, 3-22
 - I-and-D and conventional
 - mapping compared, 7-2
 - identification
 - for I- and D-space, 7-1
 - identifying partition for, 11-25
 - image, B-1, B-14
 - file structure, B-1
 - image on disk
 - non-overlaid, B-2
 - non-overlaid as linked to library, B-2
 - overlaid, B-3
 - overlaid I- and D-space, B-4
 - including debugging aid (ODT)
 - in, 10-14
 - inhibiting queuing message to, 10-35
 - installed name
 - declaring, 11-36
 - label block, 2-8
 - label block 0, B-7
 - label block 1, B-9
 - label block 2, B-9
 - label block 3, B-9
 - linking

INDEX

Task

- linking (Cont.)
 - to a supervisor-mode library, 8-10 to 8-12
 - to region, 5-13
 - to region in I- and D-space system, 7-3
 - to several libraries, 11-8 to 11-9
- list of attributes, 10-37
- logical units
 - number of, 11-39
- making checkpointable, 10-6
- map
 - linked to a common, 5-24
- mapping, 2-15, 2-20
- maximum record buffer size
 - declaring, 11-23
- memory, 2-10
- memory-resident overlay
 - operator
 - enabling, 10-34
- memory-resident overlay segment
 - changing alignment of, 10-8
- multisegment, 3-5
 - global symbol resolution, 3-17
- multiuser, 2-28
 - See Multiuser task
 - declaring read-only partition, 11-33
 - specifying a, 10-29
- ODT vector, 11-24
- overlaid, 3-7, 3-9
 - global cross-reference of, 4-12 to 4-14
 - memory savings, 3-9 to 3-10
 - segment calls, 3-10
 - virtual address space, 3-10
- overlaid I- and D-space
 - disk image, 7-8
 - map, 7-12 to 7-16
 - tree, 7-7
 - virtual address, 7-7
- overlay, 2-10
- partition
 - declaring, 11-25
- patching of
 - with object code, 11-5
- privileged, 2-25
 - See Privileged task
 - specifying, 2-25
 - specifying a, 10-33
- privileged access right for
 - establishing, 10-33
- program section order
 - effect in creating, 10-36, 10-46
- region, 2-18 to 2-20
- relocation of, 2-2
- resident common
 - system-owned, 11-11

Task (Cont.)

- resident library
 - system-owned, 11-11
- single segment, 3-4, 3-8
- slave
 - specifying a, 10-44
- specifications
 - multiple, 1-5
- specifying
 - data space in an I- and D-space, 7-5
 - KEll-A in, 10-16
- SST vector address
 - declaring, 11-37
- stack size
 - declaring, 11-34
- structure, 2-8
 - label block, 2-8
- supervisor-mode library for a, 11-31, 11-35
- system mapping status of
 - indicating, 10-26
- time-based schedule request
 - declaring UIC for, 11-38
- traceable
 - specifying a, 10-50
- UIC
 - declaring, 11-38
- use of memory management, 2-15 to 2-16
- user
 - data space definition, 7-1
- vector address
 - declaring system SST trap, 11-37
- virtual program section
 - specifying, 11-40
- window, 2-20 to 2-22
 - in I- and D-space, 7-5
 - linking to region, 5-15
- Task Builder
 - command line, 1-2
 - fast
 - See Fast Task Builder
 - function, 2-1
 - option, 1-4
 - switch, 10-1
- TASK option, 11-36
- Text information record, A-11
 - to A-12
 - format, A-12
- Throughput
 - improving TKB, F-1
- TKB
 - slow
 - to improve performance, F-11
- /TR switch, 10-50
- Transfer
 - address, A-6
 - address entry, A-6

INDEX

- Tree
 - applying autoloading indicator, 4-2
 - calling segments in, 4-6
 - calling up-tree segments, 4-6
 - multiple
 - defined, 3-30
 - defining, 3-31
 - structure, 3-18, 3-32
- .TRLUN
 - reserved global symbol, E-2
- \$\$TSKP
 - reserved PSECT name, E-4
- TSKV option, 11-37
- Type-code, 2-5, 2-7

- UFD conventions in command line, 1-9
- UIC
 - declaring in task, 11-38
 - option, 11-38
- UNITS option, 11-39
- Unmapped
 - system, 2-14
 - memory layout, 2-16
- Unnamed program section
 - See Program section
- .USLU1
 - reserved global symbol, E-2
- .USLU2
 - reserved global symbol, E-2

- Vector
 - autoloading, 3-20, 4-4
 - conventional task, B-17
 - eliminating unnecessary, 4-6
 - I- and D-space, 7-11
 - I- and D-space format, 4-5
 - I- and D-space task, 7-9, B-17
 - in region, 5-12
 - in task header
 - extension area format, B-13
- Vector (Cont.)
 - mode-switching
 - supervisor library, 8-19
 - supervisor-mode library, 8-1
 - ODT, 11-24
 - overlaid region, 5-11
 - I- and D-space task, 5-11
 - SST address, 11-27
- \$VEXT
 - low-memory context, B-10
 - reserved global symbol, E-2
- VSECT option, 11-40

- /WI switch, 10-51
- Window, 2-20 to 2-22
 - block, 2-20 to 2-22
 - creating a, 5-14
 - for a region, 5-14
 - in overlay, 3-49
 - definition block, 2-22
 - descriptor, B-20 to B-21
 - in overlay, 3-20
 - for region and linking task, 5-15
 - in I- and D-space task, 7-5
 - option, 11-41
 - region, 5-25
 - wrap around
 - in virtual sections, 5-54
- \$\$WNDS
 - reserved PSECT name, E-4
- WNDWS option, 11-41
- Work file
 - accesses
 - system overhead, F-2, F-5 to F-7
 - parameters for, F-2
 - performance
 - changing device to improve, F-5

- /XH switch, 10-52
- /XT switch, 10-53

READER'S COMMENTS

NOTE: This form is for document comments only. DIGITAL will use comments submitted on this form at the company's discretion. If you require a written reply and are eligible to receive one under Software Performance Report (SPR) service, submit your comments on an SPR form.

Did you find this manual understandable, usable, and well organized? Please make suggestions for improvement.

Did you find errors in this manual? If so, specify the error and the page number.

Please indicate the type of user/reader that you most nearly represent.

- Assembly language programmer
- Higher-level language programmer
- Occasional programmer (experienced)
- User with little programming experience
- Student programmer
- Other (please specify) _____

Name _____ Date _____

Organization _____

Street _____

City _____ State _____ Zip Code _____
or Country

Do Not Tear - Fold Here and Tape

digital




No Postage
Necessary
if Mailed in the
United States

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO.33 MAYNARD MASS.

POSTAGE WILL BE PAID BY ADDRESSEE

BSSG PUBLICATIONS ZK1-3/J35
DIGITAL EQUIPMENT CORPORATION
110 SPIT BROOK ROAD
NASHUA, NEW HAMPSHIRE 03061



Do Not Tear - Fold Here

Cut Along Dotted Line