

VMS

digital

VMS Device Support Reference Manual

Order Number: AA-PBPXA-TE

VMS Device Support Reference Manual

Order Number: AA-PBPXA-TE

June 1990

This manual provides the reference material for the *VMS Device Support Manual*, which describes how to write a driver for a device connected to a VAX processor. This manual describes the data structures, macros, and routines used in device driver programming.

Revision/Update Information: This book supersedes the reference material from the *VMS Device Support Manual*, Version 5.0. The general device support information from that manual is now in the *VMS Device Support Manual*, Version 5.4.

Software Version: VMS Version 5.4

**digital equipment corporation
maynard, massachusetts**

June 1990

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.


Restricted Rights: Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

© Digital Equipment Corporation 1990.

All Rights Reserved.
Printed in U.S.A.

The postpaid Reader's Comments forms at the end of this document request your critical evaluation to assist in preparing future documentation.

The following are trademarks of Digital Equipment Corporation:

CDA	DEQNA	MicroVAX	VAX RMS
DDIF	Desktop-VMS	PrintServer 40	VAXserver
DEC	DIGITAL	Q-bus	VAXstation
DECdtm	GIGI	ReGIS	VMS
DECnet	HSC	ULTRIX	VT
DECUS	LiveLink	UNIBUS	XUI
DECwindows	LN03	VAX	
DECwriter	MASSBUS	VAXcluster	

The following is a third-party trademark:

PostScript is a registered trademark of Adobe Systems Incorporated.

ZK5503

Production Note

This book was produced with the VAX DOCUMENT electronic publishing system, a software tool developed and sold by Digital. In this system, writers use an ASCII text editor to create source files containing text and English-like code; this code labels the structural elements of the document, such as chapters, paragraphs, and tables. The VAX DOCUMENT software, which runs on the VMS operating system, interprets the code to format the text, generate a table of contents and index, and paginate the entire document. Writers can print the document on the terminal or line printer, or they can use Digital-supported devices, such as the LN03 laser printer and PostScript printers (PrintServer 40 or LN03R ScriptPrinter), to produce a typeset-quality copy containing integrated graphics.

Contents

PREFACE	xiii
---------	------

CHAPTER 1 DATA STRUCTURES	1-1
----------------------------------	------------

1.1	CONFIGURATION CONTROL BLOCK (ACF)	1-2
-----	-----------------------------------	-----

1.2	ADAPTER CONTROL BLOCK (ADP)	1-4
-----	-----------------------------	-----

1.3	CHANNEL CONTROL BLOCK (CCB)	1-11
-----	-----------------------------	------

1.4	PER-CPU DATABASE (CPU)	1-12
-----	------------------------	------

1.5	CHANNEL REQUEST BLOCK (CRB)	1-19
1.5.1	Interrupt Transfer Vector Block (VEC)	1-23

1.6	DEVICE DATA BLOCK (DDB)	1-27
-----	-------------------------	------

1.7	DRIVER DISPATCH TABLE (DDT)	1-29
-----	-----------------------------	------

1.8	DRIVER PROLOGUE TABLE (DPT)	1-31
-----	-----------------------------	------

1.9	INTERRUPT DISPATCH BLOCK (IDB)	1-35
-----	--------------------------------	------

1.10	I/O REQUEST PACKET (IRP)	1-37
------	--------------------------	------

1.11	I/O REQUEST PACKET EXTENSION (IRPE)	1-42
------	-------------------------------------	------

1.12	OBJECT RIGHTS BLOCK (ORB)	1-44
------	---------------------------	------

1.13	SCSI CLASS DRIVER REQUEST PACKET (SCDRP)	1-46
------	--	------

Contents

1.14	SCSI CONNECTION DESCRIPTOR TABLE (SCDT)	1-54
1.15	SCSI PORT DESCRIPTOR TABLE (SPDT)	1-60
1.16	SPIN LOCK DATA STRUCTURE (SPL)	1-66
1.17	UNIT CONTROL BLOCK (UCB)	1-68
CHAPTER 2 VMS MACROS INVOKED BY DRIVERS		2-1
	ADPDISP	2-2
	BI_NODE_RESET	2-5
	CASE	2-6
	CLASS_CTRL_INIT	2-7
	CLASS_UNIT_INIT	2-8
	CPUDISP	2-9
	DDTAB	2-12
	\$DEF	2-14
	\$DEFEND	2-15
	\$DEFINI	2-16
	DEVICELOCK	2-17
	DEVICEUNLOCK	2-19
	DPTAB	2-21
	DPT_STORE	2-24
	DSBINT	2-27
	ENBINT	2-28
	\$EQLST	2-29
	FIND_CPU_DATA	2-31
	FORK	2-32
	FORKLOCK	2-33
	FORKUNLOCK	2-35
	FUNCTAB	2-37
	IFNORD, IFNOWRT, IFRD, IFWRT	2-39
	INVALIDATE_TB	2-41
	IOFORK	2-43
	LOADALT	2-44
	LOADMBA	2-45
	LOADUBA	2-46
	LOCK	2-47
	LOCK_SYSTEM_PAGES	2-49

PURDPR	2-51
READ_SYSTIME	2-52
RELALT	2-53
RELCHAN	2-54
RELDPR	2-55
RELMPR	2-56
RELSCHAN	2-57
REQALT	2-58
REQCOM	2-59
REQDPR	2-60
REQMPR	2-61
REQPCHAN	2-62
REQSCHAN	2-63
SAVIPL	2-64
SETIPL	2-65
SOFTINT	2-67
SPI\$ABORT_COMMAND	2-68
SPI\$ALLOCATE_COMMAND_BUFFER	2-69
SPI\$CONNECT	2-70
SPI\$DEALLOCATE_COMMAND_BUFFER	2-72
SPI\$DISCONNECT	2-73
SPI\$FINISH_COMMAND	2-74
SPI\$GET_CONNECTION_CHAR	2-75
SPI\$MAP_BUFFER	2-77
SPI\$RECEIVE_BYTES	2-80
SPI\$RELEASE_BUS	2-81
SPI\$RESET	2-82
SPI\$SEND_BYTES	2-83
SPI\$SEND_COMMAND	2-84
SPI\$SENSE_PHASE	2-87
SPI\$SET_CONNECTION_CHAR	2-88
SPI\$SET_PHASE	2-90
SPI\$UNMAP_BUFFER	2-91
TIMEDWAIT	2-92
TIMEWAIT	2-94
UNLOCK	2-96
UNLOCK_SYSTEM_PAGES	2-97
\$VEC	2-98
\$VECEND	2-99
\$VECINI	2-100
\$YIELD, _YIELD	2-102
WFIKPCH, WFIRLCH	2-104

Contents

CHAPTER 3 OPERATING SYSTEM ROUTINES	3-1
COM\$DELATTNAST	3-2
COM\$DRVDEALMEM	3-3
COM\$FLUSHATTNS	3-4
COM\$POST, COM\$POST_NOCNT	3-5
COM\$SETATTNAST	3-6
ERL\$DEVICERR, ERL\$DEVICTMO, ERL\$DEVICEATTN	3-8
EXE\$ABORTIO	3-10
EXE\$ALLOCBUF, EXE\$ALLOCIRP	3-12
EXE\$ALONONPAGED	3-14
EXE\$ALONPAGVAR	3-15
EXE\$ALOPHYCNTG	3-16
EXE\$ALTQUEPKT	3-17
EXE\$CREDIT_BYTCNT, EXE\$CREDIT_BYTCNT_BYTLM	3-18
EXE\$DEANONPAGED, EXE\$DEANONPGDSIZ	3-19
EXE\$DEBIT_BYTCNT(_NW), EXE\$DEBIT_BYTCNT_BYTLM(_NW)	3-20
EXE\$DEBIT_BYTCNT_ALO, EXE\$DEBIT_BYTCNT_BYTLM_ALO	3-22
EXE\$FINISHIO, EXE\$FINISHIOC	3-24
EXE\$FORK	3-26
EXE\$INSERTIRP	3-27
EXE\$INSIOQ, EXE\$INSIOQC	3-28
EXE\$INSTIMQ	3-29
EXE\$IOFORK	3-30
EXE\$MODIFY	3-31
EXE\$MODIFYLOCK, EXE\$MODIFYLOCKR	3-34
EXE\$ONEPARM	3-37
EXE\$QIODRVPKT	3-38
EXE\$QIORETURN	3-39
EXE\$READ	3-40
EXE\$READCHK, EXE\$READCHKR	3-43
EXE\$READLOCK, EXE\$READLOCKR	3-45
EXE\$RMVTIMQ	3-48
EXE\$SENSEMODE	3-49
EXE\$SETCHAR, EXE\$SETMODE	3-50
EXE\$SNDEVMSG	3-52
EXE\$WRITE	3-54
EXE\$WRITECHK, EXE\$WRITECHKR	3-56
EXE\$WRITELOCK, EXE\$WRITELOCKR	3-58

EXE\$WRTMAILBOX	3-61
EXE\$ZEROPARM	3-62
IOC\$ALOALTMAP, IOC\$ALOALTMAPN, IOC\$ALOALTMAPSP	3-63
IOC\$ALOUBAMAP, IOC\$ALOUBAMAPN	3-65
IOC\$APPLYECC	3-67
IOC\$CANCELIO	3-68
IOC\$DIAGBUFILL	3-69
IOC\$INITIATE	3-70
IOC\$IOPOST	3-72
IOC\$LOADALTMAP	3-74
IOC\$LOADMBAMAP	3-76
IOC\$LOADUBAMAP, IOC\$LOADUBAMAPA	3-77
IOC\$MOVFRUSER, IOC\$MOVFRUSER2	3-79
IOC\$MOVTOUSER, IOC\$MOVTOUSER2	3-80
IOC\$PURGDATAP	3-82
IOC\$RELALTMAP	3-84
IOC\$RELCHAN	3-86
IOC\$RELDATAP	3-87
IOC\$RELMAPREG	3-89
IOC\$RELSCHAN	3-91
IOC\$REQALTMAP	3-92
IOC\$REQCOM	3-94
IOC\$REQDATAP, IOC\$REQDATAPNW	3-96
IOC\$REQMAPREG	3-98
IOC\$REQPCHANH, IOC\$REQPCHANL, IOC\$REQSCHANH, IOC\$REQSCHANL	3-100
IOC\$RETURN	3-102
IOC\$VERIFYCHAN	3-103
IOC\$WFIKPCH, IOC\$WFIRLCH	3-104
LDR\$ALLOC_PT	3-107
LDR\$DEALLOC_PT	3-108
MMG\$UNLOCK	3-109
SMP\$ACQNOIPL	3-110
SMP\$ACQUIRE	3-111
SMP\$ACQUIREL	3-113
SMP\$RELEASE	3-114
SMP\$RELEASEL	3-115
SMP\$RESTORE	3-116
SMP\$RESTOREL	3-117

Contents

CHAPTER 4	DEVICE DRIVER ENTRY POINTS	4-1
	ALTERNATE START-I/O ROUTINE	4-2
	CANCEL-I/O ROUTINE	4-4
	CLONED UCB ROUTINE	4-6
	CONTROLLER INITIALIZATION ROUTINE	4-8
	DRIVER UNLOADING ROUTINE	4-10
	FDT ROUTINES	4-11
	INTERRUPT SERVICE ROUTINE	4-13
	REGISTER DUMPING ROUTINE	4-15
	START-I/O ROUTINE	4-17
	TIMEOUT HANDLING ROUTINE	4-19
	UNIT DELIVERY ROUTINE	4-21
	UNIT INITIALIZATION ROUTINE	4-22
	UNSOLICITED INTERRUPT SERVICE ROUTINE	4-24

INDEX

FIGURES

1-1	The I/O Database _____	1-2
1-2	Configuration Control Block (ACF) _____	1-3
1-3	Adapter Control Block (ADP) _____	1-4
1-4	Channel Control Block (CCB) _____	1-11
1-5	Per-CPU Database (CPU) _____	1-13
1-6	Channel Request Block (CRB) _____	1-19
1-7	Interrupt Transfer Vector Block (VEC) _____	1-24
1-8	Device Data Block (DDB) _____	1-27
1-9	Driver Dispatch Table (DDT) _____	1-29
1-10	Driver Prologue Table (DPT) _____	1-31
1-11	Interrupt Dispatch Block (IDB) _____	1-35
1-12	I/O Request Packet (IRP) _____	1-37
1-13	I/O Request Packet Extension (IRPE) _____	1-43
1-14	Object Rights Block (ORB) _____	1-44
1-15	SCSI Class Driver Request Packet (SCDRP) _____	1-46
1-16	SCSI Connection Descriptor Table (SCDT) _____	1-55
1-17	SCSI Port Descriptor Table (SPDT) _____	1-60
1-18	Spin Lock Data Structure (SPL) _____	1-67
1-19	Composition of Extended Unit Control Blocks _____	1-70

1-20	Unit Control Block (UCB) _____	1-71
1-21	UCB Error-Log Extension _____	1-80
1-22	UCB Local Tape Extension _____	1-81
1-23	UCB Local Disk Extension _____	1-82
1-24	UCB Terminal Extension _____	1-84
2-1	SCSI Bus Phase Longword Returned to SPI\$SENSE_PHASE _____	2-87
2-2	SCSI Bus Phase Longword Supplied to SPI\$SET_PHASE _____	2-90

TABLES

1-1	Contents of Configuration Control Block _____	1-3
1-2	Contents of Adapter Control Block _____	1-6
1-3	Contents of Channel Control Block _____	1-12
1-4	Contents of Per-CPU Database _____	1-15
1-5	Contents of Channel Request Block _____	1-21
1-6	Contents of Interrupt Transfer Vector Block (VEC) _____	1-24
1-7	Contents of Device Data Block _____	1-28
1-8	Contents of Driver Dispatch Table _____	1-30
1-9	Contents of Driver Prologue Table _____	1-33
1-10	Contents of Interrupt Dispatch Block _____	1-36
1-11	Contents of an I/O Request Packet _____	1-38
1-12	Contents of the I/O Request Packet Extension _____	1-44
1-13	Contents of Object Rights Block _____	1-45
1-14	Contents of SCSI Class Driver Request Packet _____	1-49
1-15	Contents of SCSI Connection Descriptor Table _____	1-57
1-16	Contents of SCSI Port Descriptor Table _____	1-63
1-17	Contents of the Spin Lock Data Structure _____	1-67
1-18	UCB Extensions and Sizes Defined in \$UCBDEF _____	1-69
1-19	Contents of Unit Control Block _____	1-72
1-20	UCB Error-Log Extension _____	1-80
1-21	UCB Local Tape Extension _____	1-82
1-22	UCB Local Disk Extension _____	1-83
1-23	UCB Terminal Extension _____	1-86

Preface

The *VMS Device Support Reference Manual* provides the reference material for the *VMS Device Support Manual*, which describes how to write a driver for a device connected to a VAX processor. This manual describes the data structures, macros, and routines used in driver programming.

This manual provides information you need to write a device driver that runs under VMS Version 5.4 and to load the driver into the operating system. Digital makes no guarantee that drivers written for earlier versions of VMS will execute without modification on this version of the operating system. Although the intent is to maintain the existing interface, some unavoidable changes might occur as new features are added.

The use of internal executive interfaces other than those described in this manual is discouraged.

Intended Audience

This manual is intended for system programmers who are already familiar with VAX processors and the VMS operating system.

Document Structure

This manual contains the following four parts:

Chapter 1 contains a set of figures and tables that describe the contents of each data structure in the I/O database.

Chapter 2 lists the VMS macros usually invoked by drivers.

Chapter 3 describes the context, synchronization, and I/O requirements of the operating system routines used by drivers or called as the result of a driver macro invocation.

Chapter 4 supplies a condensed description of the function and environment of each driver entry point routine.

Associated Documents

Before reading the *VMS Device Support Reference Manual*, you should have an understanding of the material discussed in the following documents:

- The *VMS Device Support Manual* is the driver programming companion document
- *VAX Hardware Handbook*
- I/O-related portions of the *VMS System Services Reference Manual*

Preface

- The section on VMS naming conventions in the *Guide to Creating VMS Modular Procedures*
- *VMS I/O User's Reference Manual: Part I* and *VMS I/O User's Reference Manual: Part II*

You may also find useful some of the material in your processor's hardware documentation, as well as in the following books:

- *VMS System Dump Analyzer Utility Manual*
- *Guide to Maintaining a VMS System*
- *VAX/VMS Internals and Data Structures*
- *VMS Delta/XDelta Utility Manual*

Conventions

This manual describes code transfer operations in three ways:

- 1 The phrase “issues a system service call” implies the use of a CALL instruction.
- 2 The phrase “calls a routine” implies the use of a JSB or BSB instruction.
- 3 The phrase “transfers control to” implies the use of a BRB, BRW, or JMP instruction.

Typographical conventions used in this book include the following:

- Generally, when first introduced in the text, a new term appears in bold print. For example:

Under the VMS operating system, a **device driver** is a set of routines and tables that the system uses to process an I/O request for a particular device type.

- Terms that serve as arguments to macros appear in boldface in the text of the manual. For example:

If an at sign (@) character precedes the **oper** argument, then the **exp** argument describes the address of the data with which to initialize the field.

- Brackets indicate that the enclosed item is optional. For example:

DSBINT [ip] [dst]

Brackets are not optional, however, in the syntax of a directory name within a file specification or in the syntax of a substring specification within an assignment statement.

- A vertical ellipsis means either that not all data that the system would display in response to the command is shown or that not all data a user would enter is shown. For example:

```
JSB      @UCB$L_FPC(R5)          ; Restore the driver process.  
.  
.  
.  
;Between these instructions, the interrupt service routine  
;can make no assumptions about the contents of R0 through R4.  
.  
.  
.  
POPR    #^M<R0,R1,R2,R3,R4,R5> ; Restore interrupt registers.
```


1

Data Structures

This chapter provides a condensed description of those data structures referenced by driver code. It lists their fields in the order in which they appear in the structures. All data structures discussed in this chapter—with the exception of the channel control block (CCB)—exist in nonpaged system memory.

Many of these structures—including the adapter control block (ADP), channel control block (CCB), channel request block (CRB), configuration control block (ACF), device data block (DDB), driver dispatch table (DDT), driver prologue table (DPT), object rights block (ORB), I/O request packet (IRP), I/O request packet extension (IRPE), and unit control block (UCB)—are collectively known as the I/O database (see Figure 1-1). The structures in the **I/O database** help the VMS operating system and device drivers monitor the status and control the functions of the I/O subsystem. They provide the following types of information:

- Descriptions of each pending and in-progress I/O request
- Characteristics of each device type
- Number and type of each device unit
- Status of current activity on each device unit
- External entry points to all device drivers
- Entry points for controller and device unit initialization routines
- Code that dispatches interrupts to the appropriate servicing routines
- Addresses of device registers
- Bit maps describing the allocation of data paths and map registers

Aside from the I/O database structures, this chapter includes descriptions of those data structures VMS uses to maintain multiprocessing synchronization and record processor-specific information: the spin lock data structure (SPL) and the per-CPU database structure (CPU), respectively. Additionally, it describes the structures that implement the SCSI port interface that supports the creation of SCSI class driver.

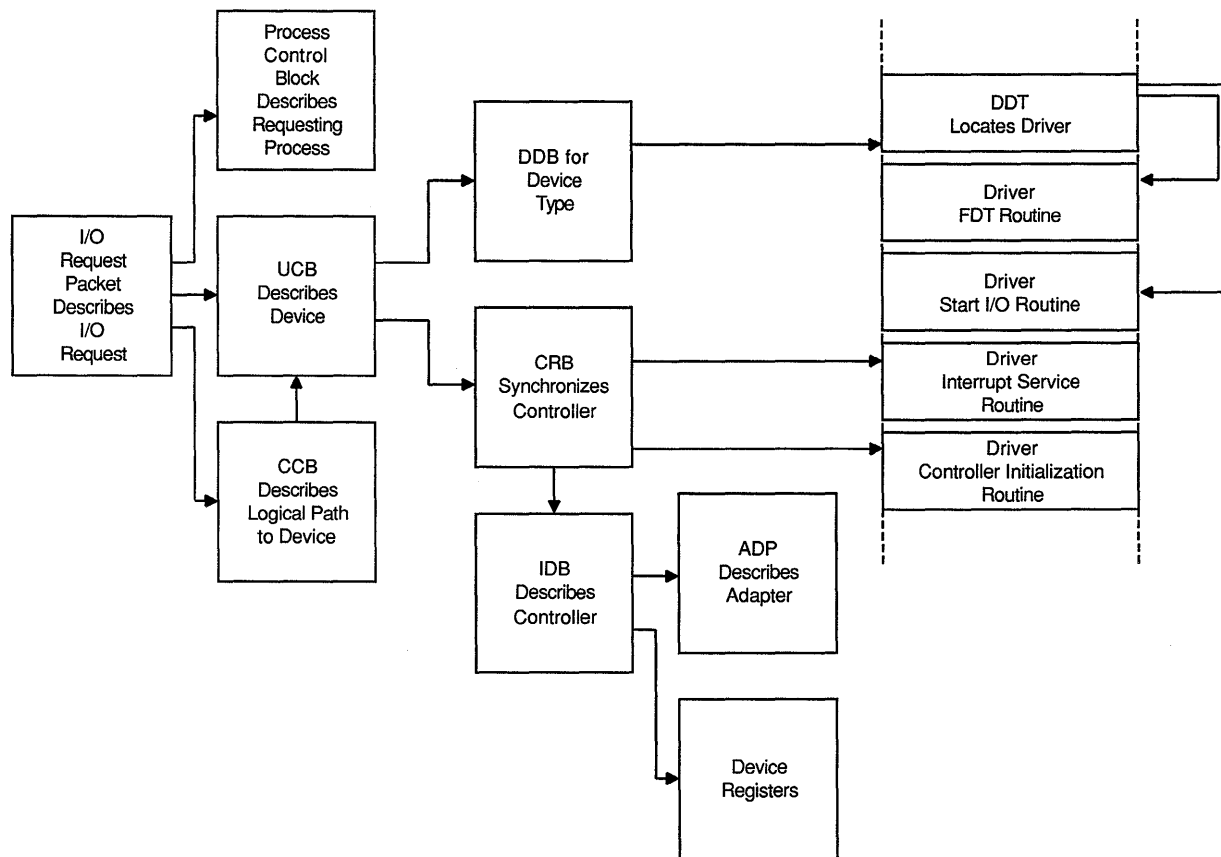
Notes: Driver code must consider fields marked by asterisks (*) to be read-only fields.

Fields marked “Reserved” or “Unused” are reserved for future use by Digital unless otherwise specified.

When referring to locations within a data structure, a driver should use symbolic offsets, *not* numeric offsets, from the beginning of the structure. Numeric offsets are likely to change with each new release of the VMS operating system. The figures in this chapter list VMS Version 5.4 numeric offsets to aid in driver debugging.

Data Structures

Figure 1-1 The I/O Database



ZK-1766-GE

1.1 Configuration Control Block (ACF)

The configuration control block (ACF) is used by the SYSGEN autoconfiguration facility to describe the device it is adding to the system. Device drivers can gain access to this data structure only if they have specified a unit delivery routine in the DPT and only when that routine is executing. Under certain conditions, the information stored in the ACF might be useful to a unit delivery routine.

The fields described in the configuration control block are illustrated in Figure 1-2 and described in Table 1-1. An asterisk (*) indicates a read-only field in tables and figures.

Data Structures

1.1 Configuration Control Block (ACF)

Figure 1-2 Configuration Control Block (ACF)

ACF\$L_ADAPTER*				0
ACF\$L_CONFIGREG*				4
ACF\$B_AFLAG*	ACF\$B_AUNIT*	ACF\$W_AVECTOR*		8
ACF\$L_CONTRLREG*				12
ACF\$W_CUNIT*		ACF\$W_CVECTOR*		16
ACF\$L_DEVNAME*				20
ACF\$L_DRVNAME*				24
ACF\$B_COMBO_VEC*	ACF\$B_CNUMVEC*	ACF\$W_MAXUNITS*		28
Unused		ACF\$B_NUMUNIT*	ACF\$B_COMBO_CSR*	32
ACF\$L_DLVR_SCRH				36

*A read-only field

Table 1-1 Contents of Configuration Control Block

Field Name	Contents
ACF\$L_ADAPTER*	Address of ADP for adapter currently being configured.
ACF\$L_CONFIGREG*	Address of configuration register for adapter currently being configured.
ACF\$W_AVECTOR*	Offset from base of SCB to interrupt vector of adapter currently being configured.
ACF\$B_AUNIT*	Adapter unit number of device or controller currently being configured.
ACF\$B_AFLAG*	Flags associated with autoconfiguration operation. Flags defined in this field include the following: <ul style="list-style-type: none"> ACF\$V_RELOAD Reloading driver code. ACF\$V_CRBBLT CRB and IDB already built for device. ACF\$V_SCBVEC CVECTOR is offset into SCB. ACF\$V_NOLOAD_DB Do not load I/O database, only load driver. ACF\$V_SUPPORT VMS-supported device. ACF\$V_GETDONE Addresses of data structures in I/O database have been obtained. ACF\$V_BVP Multiport BVP adapter.
ACF\$L_CONTRLREG*	Address of CSR for controller currently being configured.

(continued on next page)

Data Structures

1.1 Configuration Control Block (ACF)

Table 1-1 (Cont.) Contents of Configuration Control Block

Field Name	Contents
ACF\$W_CVECTOR*	Offset into ADP vector table to longword that contains transfer address of interrupt vector used by controller currently being configured (if ACF\$V_SCBVEC is not set). If ACF\$V_SCBVEC is set, this field is the offset from the SCB base to the interrupt vector of the controller currently being configured.
ACF\$B_CUNIT*	Unit number of device currently being configured.
ACF\$L_DEVNAME*	Address of counted ASCII string that gives name of controller currently being configured.
ACF\$L_DRVNAME*	Address of counted ASCII string that gives driver name for controller currently being configured.
ACF\$W_MAXUNITS*	Maximum number of units that can be connected to controller currently being configured.
ACF\$B_CNUMVEC*	Number of interrupt vectors to configure for controller currently being configured.
ACF\$B_COMBO_VEC*	Offset to vectors for combo device. (The name of this field is ACF\$B_COMBO_VECTOR_OFFSET.)
ACF\$B_COMBO_CSR*	Offset to start of control registers of combo device. (The name of this field is ACF\$B_COMBO_CSR_OFFSET.)
ACF\$B_NUMUNIT*	Number of units to be configured for controller currently being configured.
ACF\$L_DLVR_SCRH	Field available for use by unit delivery routine. SYSGEN never alters this field.

1.2 Adapter Control Block (ADP)

Each MASSBUS adapter, UNIBUS adapter, Q22 bus, and VAXBI node configured in a VAX system is represented to VMS and driver routines by an adapter control block (ADP). The ADP stores adapter-specific static and dynamic data such as the adapter CSR address and map-register wait queues.

Depending upon the type of I/O adapter being described, the ADP size is variable and subject to the length of the bus-specific ADP extension. Table 1-2 defines the fields that appear in a UNIBUS ADP; these fields are pictured in Figure 1-3. Bus-specific extensions start at offset ADP\$L_HOSTNODE in the ADP.

Figure 1-3 Adapter Control Block (ADP)

ADP\$L_CSR*			0
ADP\$L_LINK*			4
ADP\$B_NUMBER*	ADP\$B_TYPE*	ADP\$W_SIZE*	8
ADP\$W_ADPTYPE*		ADP\$W_TR*	12

(continued on next page)

Data Structures

1.2 Adapter Control Block (ADP)

Figure 1–3 (Cont.) Adapter Control Block (ADP)

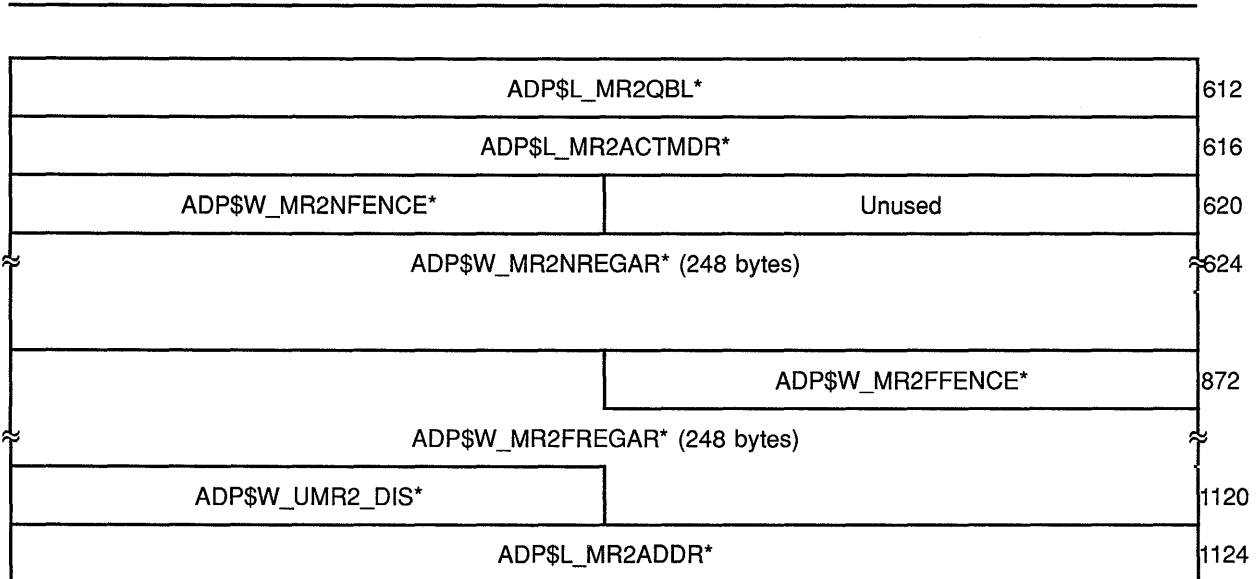
ADP\$L_VECTOR*			16
ADP\$L_DPQFL*			20
ADP\$L_DPQBL*			24
ADP\$L_AVECTOR*			28
ADP\$L_BI_IDR*			32
ADP\$W_BI_VECTOR*	ADP\$W_BI_FLAGS*		36
ADP\$L_SCB_PAGE*			40
ADP\$L_BIMASTER*			44
ADP\$B_ADDR_BITS*	Unused	ADP\$W_ADPDISP_FLAGS*	48
Reserved			52
ADP\$L_MRQFL*/ADP\$L_HOSTNODE*			56
ADP\$L_MRQBL*			60
ADP\$L_INTD_UBA* (12 bytes)			64
ADP\$L_UBASCB* (16 bytes)			76
ADP\$L_UBASPTE*			92
ADP\$L_MRACTMDRS*			100
ADP\$W_MRNFFENCE*	ADP\$W_DPBITMAP*		104
ADP\$W_MRNREGARY* (248 bytes)			108
ADP\$W_MRFREGARY* (248 bytes)		ADP\$W_MRFFENCE*	356
ADP\$W_UMR_DIS*			604
ADP\$L_MR2QFL*			608

(continued on next page)

Data Structures

1.2 Adapter Control Block (ADP)

Figure 1-3 (Cont.) Adapter Control Block (ADP)



*A read-only field

Table 1-2 Contents of Adapter Control Block

Field Name	Contents
ADP\$L_CSR*	Virtual address of adapter configuration register. For a generic VAXBI adapter, this field contains the address of the base of the adapter's node space. The VMS adapter initialization routine writes this field. The configuration register marks the base of adapter register space, an area that contains data path registers, map registers, or any other registers appropriate to the implementation of the adapter.
ADP\$L_LINK*	Address of next ADP. The VMS adapter initialization routine writes this field. A value of 0 indicates that this is the last ADP.
ADP\$W_SIZE*	Size of ADP. The VMS adapter initialization routine writes this field when the routine creates the ADP. For nondirect-vector UNIBUS adapters, ADP\$W_SIZE includes the space allocated for the four UNIBUS interrupt service routines (for BR4 to BR7) and the vector jump table.
ADP\$B_TYPE*	Type of data structure. The VMS adapter initialization routine writes the symbolic constant DYN\$C_ADP into this field when the routine creates the ADP.
ADP\$B_NUMBER*	Number of this type of adapter (for example, the number for a third MASSBUS adapter is 2). The VMS adapter initialization routine writes this field when the routine creates the ADP.

(continued on next page)

Data Structures

1.2 Adapter Control Block (ADP)

Table 1–2 (Cont.) Contents of Adapter Control Block

Field Name	Contents
ADP\$W_TR*	Nexus number of adapter. The VMS adapter initialization routine writes this field when the routine creates the ADP. The driver-loading procedure compares the nexus number specified in a CONNECT command with this field of each ADP in the system to determine to which adapter a device is attached. For a generic VAXBI adapter, this field contains its VAXBI node ID.
ADP\$W_ADPTYPE*	Type of adapter. The VMS adapter initialization routine writes the symbolic constant AT\$_UBA into this field when the routine creates an ADP for a UNIBUS adapter or Q22 bus; AT\$_MBA for a MASSBUS adapter; and AT\$_GENBI for a generic VAXBI adapter.
ADP\$L_VECTOR*	<p>Address of adapter dispatch table. The table is 512 bytes of longword vectors that correspond to device interrupt vectors (0₈–777₈).</p> <p>On VAX processors that handle direct-vector interrupts, ADP\$L_VECTOR points to the second (or subsequent) page of the SCB. The CPU uses this page when it dispatches the device interrupt to the driver interrupt service routine. Each vector entry that corresponds to a vector in use contains the address of the controller's interrupt dispatcher (CRB\$L_INTD). (The actual stored value is CRB\$L_INTD+1, the set low bit of the address indicating that the interrupt stack is to be used in servicing interrupts.)</p> <p>On VAX processors that handle non-direct-vector interrupts, ADP\$L_VECTOR points to a page allocated from nonpaged pool called the adapter dispatch table (or vector jump table). Each longword in the page that corresponds to a vector in use contains the address of the controller's interrupt dispatcher (CRB\$L_INTD+2). When the UNIBUS adapter interrupts on behalf of a UNIBUS device, the UNIBUS adapter interrupt service routine saves R0 through R5, determines the vector address of the interrupting device, indexes into the vector-jump table, and jumps to the instruction at CRB\$L_INTD+2.</p> <p>For both types of VAX processor, adapter dispatch table entries that correspond to unused vectors contain the address of the adapter's unexpected-interrupt service routine.</p>
ADP\$L_DPQFL*	<p>Data path wait queue forward link. IOC\$REQDATAP and IOC\$RELDATAP read and write this field. When a driver fork process requests a buffered data path and none is currently available, IOC\$REQDATAP saves driver context in the device's UCB fork block, inserts the fork block address in the data path wait queue, and suspends the driver fork process.</p> <p>When another driver calls IOC\$RELDATAP to release a buffered data path, the routine dequeues a UCB fork block address from the data path wait queue, allocates a data path to the driver, and reactivates that driver fork process.</p> <p>This field is also known as ADP\$L_MBASCB. For MASSBUS adapters and generic VAXBI adapters, the VMS adapter initialization routine stores the address of the adapter's interrupt vector in this field. Certain power failure recovery operations use the contents of ADP\$L_MBASCB to refresh the SCB vectors. The actual stored value is CRB\$L_INTD+1, the set low bit of the address indicating that the interrupt stack is to be used in servicing interrupts.</p>

(continued on next page)

Data Structures

1.2 Adapter Control Block (ADP)

Table 1–2 (Cont.) Contents of Adapter Control Block

Field Name	Contents
ADP\$L_DPQBL*	Data path wait queue backward link. IOC\$REQDATAP and IOC\$RELDATAP read and write this field. This field is also known as ADP\$L_MBASPTE. For generic VAXBI adapters, the VMS adapter initialization routine stores here the contents of the first of 16 SPTEs that map the adapter's node space. For the MASSBUS adapter, the routine stores here the SPTE value that maps MBA address space. Certain recovery operations use the contents of ADP\$L_MBASPTE to restore SPTE values and remap node space following a power failure.
ADP\$L_AVECTOR*	Address of first SCB vector for adapter.
ADP\$L_BI_IDR*	Longword mask specifying, by a single set bit, which VAXBI node is the destination of interrupts from this adapter. In VAX 82x0/83x0 systems, the VAXBI node of the primary processor becomes the destination for interrupts; in VAX 85x0/8700/88x0 and VAX 6000-series systems, it is the VAXBI node at which the memory-interconnect-to-VAXBI adapter (NBIB, PBIB, or DWMBAB) resides.
ADP\$W_BI_FLAGS*	VAXBI device flags field.
ADP\$W_BI_VECTOR*	Offset of the first interrupt vector for this VAXBI node from the start of its SCB page. ADP\$L_AVECTOR contains the address of this vector.
ADP\$L_SCB_PAGE*	Offset to SCB page for this VAXBI device.
ADP\$L_BIMASTER*	Address of the ADP of the master device of the VAXBI (for example, the DWMBAB in a VAX 6000-series system).
ADP\$W_ADPDISP_FLAGS*	Flags used by the ADPDISP macro to control branching according to adapter characteristics. The following bit fields are defined within ADP\$W_ADPDISP_FLAGS:
ADP\$V_ADPDISP_INIT	ADPDISP flags have been initialized
ADP\$V_ADAP_MAPPING	Adapter mapping supported
ADP\$V_DIRECT_VECTOR	Direct-vector interrupts
ADP\$V_AUTOPURGE_DP	Autopurging datapath
ADP\$V_BUFFERED_DP	Buffered datapath supported
ADP\$V_ODD_XFER_BDP	Odd transfers supported on buffered data path
ADP\$V_ODD_XFER_DDP	Odd transfers supported on direct data path
ADP\$V_EXTENDED_MAPREG	Alternate map registers (registers 496 to 8191) supported
ADP\$V_QBUS	Q22-bus adapter
<15:9>	Reserved to Digital
ADP\$B_ADDR_BITS*	Number of adapter address bits. This field contains the value 22 (for Q22-bus systems) and 18 (for UNIBUS adapters).
ADP\$L_HOSTNODE*	The offset to the bus-specific ADP extension.

(continued on next page)

Data Structures

1.2 Adapter Control Block (ADP)

Table 1–2 (Cont.) Contents of Adapter Control Block

Field Name	Contents
ADP\$L_MRQFL*	<p>Standard-map-register wait queue's forward link and the first longword in the UNIBUS adapter extension. IOC\$ALOUBAMAP, IOC\$REQMAPREG, and IOC\$RELMAPREG read and write these fields. When a driver fork process requests a set of standard map registers and the set is not currently available, IOC\$REQMAPREG saves driver fork context in the device's UCB fork block, inserts the fork block address in the standard-map-register wait queue, and suspends the driver fork process.</p> <p>When another driver calls IOC\$RELMAPREG to release a set of standard map registers, the routine dequeues a UCB fork block address from the standard-map-register wait queue, allocates the requested set of map registers to the driver, and reactivates that driver fork process.</p>
ADP\$L_MRQBL*	<p>Standard-map-register wait queue's backward link. IOC\$ALOUBAMAP, IOC\$REQMAPREG, and IOC\$RELMAPREG read and write this field.</p>
ADP\$L_INTD_UBA*	<p>Interrupt transfer vector. The VMS adapter initialization routine places executable code in this field to allow certain Digital-supplied adapters or controllers to dispatch to adapter-specific interrupt and error handling routines.</p>
ADP\$L_UBASCB*	<p>Series of four longwords that contain SCB entry values, one for each bus request (BR) level or interrupt vector. The UNIBUS adapter power failure recovery procedure uses these values.</p>
ADP\$L_UBASPTE*	<p>System page-table entry (PTE) values for base of UNIBUS adapter register space and base of UNIBUS I/O register space. These values contained in this quadword field are used during UNIBUS adapter power failure recovery.</p>
ADP\$L_MRACTMDRS*	<p>Number of active standard map register descriptors in arrays to which ADP\$W_MRNREGARY and ADP\$W_MRFREGARY point. IOC\$REQMAPREG and IOC\$RELMAPREG use these fields when allocating and deallocating standard map registers.</p>
ADP\$W_DPBITMAP*	<p>Data path allocation bit map. IOC\$REQDATAP and IOC\$RELDATAP read and write this field. The VMS adapter initialization routine sets the bit map to show as available all the buffered data paths supported by the UNIBUS adapter. (The adapter initialization routine for certain VAX processors whose UNIBUS adapters or Q22-bus interfaces do not supply buffered data paths marks three data paths as available. This facilitates the writing of machine-independent code that can execute regardless of the presence of buffered data paths.)</p> <p>The state of each of the available buffered data paths (whether in use or available) is recorded in the data path allocation bit map. One data path corresponds to each bit in the field. If a bit is clear, the related data path is currently allocated to a driver fork process.</p>
ADP\$W_MRNFBENCE*	<p>Boundary marker for the array specified by ADP\$W_MRNREGARY; contains –1.</p>
ADP\$W_MRNREGARY*	<p>Standard map register "number of registers" array of 124 words. The number of words, or cells, that are active in this array is contained in ADP\$L_MRACTMDRS. Each active cell gives the number of free standard map registers. For each active cell in this array, there is a corresponding first free map register number in the "first register" array (ADP\$W_MRFREGARY). Together, these values give the base map register and number of free map registers for a block of free map registers. This information is used to allocate and deallocate standard map registers.</p>

(continued on next page)

Data Structures

1.2 Adapter Control Block (ADP)

Table 1–2 (Cont.) Contents of Adapter Control Block

Field Name	Contents
UNIBUS Adapter Extension	
ADP\$W_MRFENCE*	Boundary marker for array specified by ADP\$W_MRFREGARY; contains –1.
ADP\$W_MRFREGARY*	Standard map register “first register” array of 124 words. The number of currently active cells in this array is contained in ADP\$L_MRACTMDRS. Each active cell gives a number of the first free map register within a block of free map registers. For each active cell in this array, there is a corresponding cell in the “number of registers” array (ADP\$W_MRNREGARY) that gives a number of free map registers. Together, these values give the base map register and number of free map registers for a block of free map registers. This information is used to allocate and deallocate standard map registers.
ADP\$W_UMR_DIS*	Number of disabled standard map registers. During system initialization, some standard map registers can be disabled so that their corresponding UNIBUS and Q22-bus addresses can be accessed directly through UNIBUS-space or Q22-bus-space physical addresses.
ADP\$L_MR2QFL*	Alternate-map-register wait queue’s forward link. IOC\$ALOALTMAP, IOC\$REQALTMAP, and IOC\$RELALTMAP read and write this field. When a driver fork process requests a set of Q22-bus alternate map registers and the set is not currently available, IOC\$REQALTMAP saves driver context in the device’s UCB fork block, inserts the fork block address in the alternate-map-register wait queue, and suspends the driver fork process. When another driver calls IOC\$RELALTMAP to release a sufficient number of map registers, the routine dequeues a UCB fork block from the alternate-map-register wait queue, allocates the requested set of map registers to the driver, and reactivates that driver fork process.
ADP\$L_MR2QBL*	Alternate-map-register wait queue’s backward link. IOC\$ALOALTMAP, IOC\$REQALTMAP, and IOC\$RELALTMAP read and write this field when allocating and deallocating from the set of Q22-bus alternate map registers.
ADP\$L_MR2ACTMDR*	Number of active map register descriptors in arrays to which ADP\$W_MR2NREGAR and ADP\$W_MR2FREGAR point. IOC\$ALOALTMAP, IOC\$REQALTMAP, and IOC\$RELMAPREG use these fields when allocating and deallocating Q22-bus alternate map registers.
ADP\$W_MR2NFENCE*	Boundary marker for the array specified by ADP\$W_MR2NREGAR; contains –1.
ADP\$W_MR2NREGAR*	Alternate-map-register “number of registers” array of 124 words. The number of words, or cells, that are active in this array is contained in ADP\$L_MR2ACTMDR. Each active cell gives a number of map registers in a block of free alternate map registers. For each active cell in this array, there is a corresponding first free map register number in the array specified by ADP\$W_MR2FREGAR. Together, these values give the base map register and the number of free map registers for a block of free alternate map registers. IOC\$ALOALTMAP, IOC\$REQALTMAP, and IOC\$RELALTMAP use this information when allocating and deallocating from Q22-bus alternate map registers.

(continued on next page)

Table 1–2 (Cont.) Contents of Adapter Control Block

Field Name	Contents
UNIBUS Adapter Extension	
ADP\$W_MR2FFENCE*	Boundary marker for the array specified by ADP\$W_MR2NREGAR; contains –1.
ADP\$W_MR2FREGAR*	Alternate map register “first register” array of 124 words. The number of words, or cells, that are active in this array is contained in ADP\$L_MR2ACTMDR. Each active cell gives the number of the first free map register within a block of free map registers. For each active cell in this array, there is a corresponding cell in the “number of registers” array, ADP\$W_MR2NREGAR. Together, these values give the base map register and the number of free map registers for a block of free map registers.
ADP\$W_UMR2_DIS*	Number of disabled Q22-bus alternate map registers. During system initialization, some map registers can be disabled so that their corresponding Q22-bus addresses can be accessed directly through physical addresses.
ADP\$L_MR2ADDR	Address of the first Q22-bus alternate map register mapped in CPU node private space. The value varies for each processor with alternate map registers. IOC\$LOADUBAMAP reads this field when accessing alternate map registers.

1.3 Channel Control Block (CCB)

When a process assigns an I/O channel to a device unit with the \$ASSIGN system service, EXE\$ASSIGN locates a free block among the process’s preallocated channel control blocks (CCBs). EXE\$ASSIGN then writes into the CCB a description of the device attached to the CCB’s channel.

The channel control block is the only data structure described in this chapter that exists in the control (P1) region of a process address space. It is illustrated in Figure 1–4 and described in Table 1–3.

Figure 1–4 Channel Control Block (CCB)

CCB\$L_UCB*			0
CCB\$L_WIND*			4
CCB\$W_IOC*	CCB\$B_AMOD*	CCB\$B_STS*	8
CCB\$L_DIRP*			12

*A read-only field

Data Structures

1.3 Channel Control Block (CCB)

Table 1-3 Contents of Channel Control Block

Field Name	Contents
CCB\$L_UCB*	Address of UCB of assigned device unit. EXE\$ASSIGN writes a value into this field. EXE\$QIO reads this field to determine that the I/O request specifies a process I/O channel assigned to a device and to obtain the device's UCB address.
CCB\$L_WIND*	Address of window control block (WCB) for file-structured device assignment. This field is written by an ACP or XQP and read by EXE\$QIO. A file-structured device's XQP or ACP creates a WCB when a process accesses a file on a device assigned to a process I/O channel. The WCB maps the virtual block numbers of the file to a series of physical locations on the device.
CCB\$B_STS*	Channel status.
CCB\$B_AMOD*	Access mode plus 1 of the channel. EXE\$ASSIGN writes the access mode value into this field.
CCB\$W_IOC*	Number of outstanding I/O requests on channel. EXE\$QIO increases this field when it begins to process an I/O request that specifies the channel. During I/O postprocessing, the special kernel-mode AST routine decrements this field. Some FDT routines and EXE\$DASSGN read this field.
CCB\$L_DIRP*	Address of IRP for requested deaccess. A number of outstanding I/O requests can be pending on the same process I/O channel at one time. If the process that owns the channel issues an I/O request to deaccess the device, EXE\$QIO holds the deaccess request until all other outstanding I/O requests are processed.

1.4 Per-CPU Database (CPU)

A per-CPU database structure exists for each processor in a VMS multiprocessing environment. The per-CPU database records processor-specific information such as the current process control block (PCB), the priority of the current process, and the physical processor identifier. It points to the processor's interrupt stack and contains the list heads for the processor's fork queues and I/O postprocessing queue.

To ensure that the path of a processor's activity at booting and on the interrupt stack remains independent of the paths of other active processors in the system, VMS places a separate boot stack and a separate interrupt stack (formerly pointed to by EXE\$GL_INTSTK) adjacent to the area allocated for the per-CPU database structure. The processor's boot stack, interrupt stack, and per-CPU database fields are virtually contiguous in system address space, although three no-access guard pages prevent the expansion of the stacks beyond the areas reserved for their use. Offset CPU\$L_INTSTK in the per-CPU database points to the interrupt stack.

The fields described in the per-CPU database are illustrated in Figure 1-5 and described in Table 1-4.

Data Structures

1.4 Per-CPU Database (CPU)

Figure 1-5 Per-CPU Database (CPU)

CPU\$L_CURPCB*				0
CPU\$L_REALSTACK*				4
CPU\$B_SUBTYPE*	CPU\$B_TYPE*	CPU\$W_SIZE*		8
CPU\$B_CUR_PRI*	CPU\$B_CPUMTX*	CPU\$B_STATE*	CPU\$B_BUSYWAIT	12
CPU\$L_INTSTK*				16
CPU\$L_WORK_REQ*				20
CPU\$L_PERCPUVA*				24
CPU\$L_SAVED_AP*				28
CPU\$L_HALTPC*				32
CPU\$L_HALTPSL*				36
CPU\$L_SAVED_ISP*				40
CPU\$L_PCBB*				44
CPU\$L_SCBB*				48
CPU\$L_SISR*				52
CPU\$L_P0BR*				56
CPU\$L_P0LR*				60
CPU\$L_P1BR*				64
CPU\$L_P1LR*				68
CPU\$L_BUGCODE*				72
CPU\$B_CPUDATA* (16 bytes)				76
CPU\$L_MCHK_MASK*				92
CPU\$L_MCHK_SP*				96
CPU\$L_POPT_PAGE*				100

(continued on next page)

Data Structures

1.4 Per-CPU Database (CPU)

Figure 1-5 (Cont.) Per-CPU Database (CPU)

Reserved (408 bytes)			104
CPU\$Q_SWIQFL* (48 bytes)			512
CPU\$L_PSFL*			560
CPU\$L_PSBL*			564
CPU\$Q_WORK_FQFL*			568
CPU\$L_QLOST_FQFL*			576
CPU\$L_QLOST_FQBL*			580
CPU\$B_QLOST_FLCK*	CPU\$B_QLOST_TYPE*	CPU\$W_QLOST_SIZE*	584
CPU\$L_QLOST_FPC*			588
CPU\$L_QLOST_FR3*			592
CPU\$L_QLOST_FR4*			596
CPU\$Q_BOOT_TIME*			600
CPU\$Q_CPUID_MASK*			608
CPU\$L_PHY_CPUID*			616
CPU\$L_CAPABILITY*			620
CPU\$L_TENUSEC*			624
CPU\$L_UBDELAY*			628
CPU\$L_KERNEL* (28 bytes)			632
CPU\$L_NULLCPU*			660

(continued on next page)

Data Structures

1.4 Per-CPU Database (CPU)

Figure 1-5 (Cont.) Per-CPU Database (CPU)

CPU\$W_UKERNEL* (14 bytes)		664
CPU\$W_UNULLCPU*		676
CPU\$W_HARDAFF*	CPU\$W_CLKUTICS*	680
CPU\$L_RANK_VEC*		684
CPU\$L_IPL_VEC*		688
CPU\$L_IPL_ARRAY* (128 bytes)		692
CPU\$L_TPOINTER*		820
CPU\$W_SANITY_TICKS*	CPU\$W_SANITY_TIMER*	824
CPU\$L_VP_OWNER*		828
CPU\$L_VP_VARIANT_EXIT*		832
CPU\$L_VP_FLAGS*		836
CPU\$L_VP_CPUTIM*		840
Reserved	CPU\$B_FLAGS*	844
CPU\$L_INTFLAGS*		848

*A read-only field

Table 1-4 Contents of Per-CPU Database

Field	Contents
CPU\$L_CURPCB*	Address of current PCB. The scheduler writes this field.
CPU\$L_REALSTACK*	Physical address of boot stack.
CPU\$W_SIZE*	Size of the per-CPU database, including the size of the boot stack but not the interrupt stack or the interrupt stack's guard pages.
CPU\$B_TYPE*	Type of data structure. VMS writes the value DYN\$C_MP into this field when it creates the per-CPU database.
CPU\$B_SUBTYPE*	Structure subtype. VMS writes the value DYN\$C_MP_CPU into this field when it creates the per-CPU database.
CPU\$B_BUSYWAIT*	Concurrent busywait count for this processor.

(continued on next page)

Data Structures

1.4 Per-CPU Database (CPU)

Table 1-4 (Cont.) Contents of Per-CPU Database

Field	Contents
CPU\$B_STATE*	State of this processor. VMS defines the following processor states: CPU\$C_INIT Processor is being initialized. CPU\$C_RUN Processor is running. CPU\$C_STOPPING Processor is stopping. CPU\$C_STOPPED Processor is stopped. CPU\$C_TIMEOUT Logical console has timed out. CPU\$C_BOOT_REJECTED Processor has refused to join multiprocessing system. CPU\$C_BOOTED Processor has booted, but is waiting to join multiprocessing active set.
CPU\$B_CPUMTX*	Count of acquisitions of CPUMTX mutex.
CPU\$B_CUR_PRI*	Current process priority. The scheduler writes this field.
CPU\$L_INTSTK*	Address of initial interrupt stack.
CPU\$L_WORK_REQ*	Work request bits. A processor sets one or more of these bits in another processor's per-CPU database when directing an interprocessor interrupt to that processor. The following fields are defined within CPU\$L_WORK_REQ: CPU\$V_INV_TBS Request to invalidate single address (SMP\$GL_INVALID) in translation buffer CPU\$V_INV_TBA Request to invalidate all addresses in translation buffer CPU\$V_TBACK Acknowledgment that a processor requested to invalidate its translation buffer has done so CPU\$V_BUGCHK Request to bugcheck CPU\$V_BUGCHKACK Acknowledgment that the processor has saved process context and per-CPU data so that the crash CPU can continue to perform a bugcheck CPU\$V_RECALSCHD Recalculate per-CPU mask and reschedule CPU\$V_UPDASTLVL Request to update processor AST level register (PR\$ASTLVL) CPU\$V_UPDTODR Request to update processor time-of-day register (PR\$_TODR) CPU\$V_WORK_FQP Request to process internal fork queue (CPU\$Q_WORK_IFQ) CPU\$V_QLOST Request to stall until quorum regained CPU\$V_RESCHED Request to initiate software interrupt at IPL 3 CPU\$V_VIRTCONS Request to enter virtual console mode CPU\$V_IOPOST Request to request IPL 4 software interrupt <28:31> Processor-specific work request bits
CPU\$L_PERCPUVA*	Virtual address of this per-CPU database structure.

(continued on next page)

Data Structures

1.4 Per-CPU Database (CPU)

Table 1–4 (Cont.) Contents of Per-CPU Database

Field	Contents
CPU\$_SAVED_AP*	Halt restart code.
CPU\$_HALTPC*	Halt PC for restart.
CPU\$_HALTPSL*	Halt PSL for restart.
CPU\$_SAVED_ISP*	Saved ISP for restart.
CPU\$_PCBB*	PCBB from power down.
CPU\$_SCBB*	SCBB from power down.
CPU\$_SISR*	SISR from power down.
CPU\$_P0BR*	P0 base register (used by system power failure and bugcheck routines).
CPU\$_P0LR*	P0 length register (used by system power failure and bugcheck routines).
CPU\$_P1BR*	P1 base register (used by system power failure and bugcheck routines).
CPU\$_P1LR*	P1 length register (used by system power failure and bugcheck routines).
CPU\$_BUGCODE*	Bugcheck code.
CPU\$_B_CPUDATA*	Processor-specific hardware revision information. The first longword of this 16-byte field always contains the processor's system ID (SID) register, and is also defined as CPU\$_SID.
CPU\$_MCHK_MASK*	Function mask for current machine check recovery block.
CPU\$_MCHK_SP*	Saved SP for return at end of machine check recovery block. This field is zero if there is no current recovery block.
CPU\$_P0PT_PAGE*	System virtual address of a page reserved to this processor that is used as a P0 page table when memory management is being enabled.
CPU\$Q_SWIQFL*	Twelve longwords representing the forward and backward links for the software interrupt queues (fork IPLs 6 through 11).
CPU\$_PSFL*	CPU-specific I/O postprocessing queue forward link.
CPU\$_PSBL*	CPU-specific I/O postprocessing queue backward link.
CPU\$Q_WORK_FQFL*	Work packet queue. This field is also called CPU\$Q_WORK_IFQ.
CPU\$_QLOST_FQFL*	Quorum loss fork queue forward link.
CPU\$_QLOST_FQBL*	Quorum loss fork queue blink link.
CPU\$W_QLOST_SIZE*	Quorum loss fork block size.
CPU\$_B_QLOST_TYPE*	Quorum loss fork block type.
CPU\$_B_QLOST_FLCK*	Quorum loss fork lock.
CPU\$_QLOST_FPC*	Quorum loss fork PC.
CPU\$_QLOST_FR3*	Quorum loss fork R3.
CPU\$_QLOST_FR4*	Quorum loss fork R4.
CPU\$Q_BOOT_TIME*	System time at which this processor was bootstrapped.
CPU\$Q_CPUID_MASK*	Bit mask representing this processor's CPU ID.

(continued on next page)

Data Structures

1.4 Per-CPU Database (CPU)

Table 1-4 (Cont.) Contents of Per-CPU Database

Field	Contents
CPU\$L_PHY_CPUID*	Integer that uniquely identifies the local processor in a multiprocessor configuration. This value is system specific. (For example, in a VAX 8300/8350 configuration, it is the VAXBI node ID. For a VAX 8800, it is the left or right bit from the processor's system ID register (PR\$_SID); for a VAX 8810/8820/8830 it is the CPU number (0 to 3) from PR\$_SID. In a VAX 6000-series configuration, it is the XMI node ID. VMS uses the physical ID principally to locate the per-CPU database and interrupt stack of a processor that it is restarting.)
CPU\$L_CAPABILITY*	Bit mask of this processor's capabilities. VMS defines the following capabilities in \$CPBDEF: CPB\$C_PRIMARY Primary CPU. CPB\$C_NS Reserved to Digital. CPB\$C_QUORUM Quorum required. CPB\$C_HARDAFF Hard affinity. Reserved for diagnostics software.
CPU\$L_TENUSEC*	10-microsecond delay value.
CPU\$L_UBDELAY*	UNIBUS delay counter.
CPU\$L_KERNEL*	Set of seven longwords that tally the processor's clock ticks in kernel mode, in executive mode, in supervisor mode, in user mode, on the interrupt stack, in compatibility mode, and in kernel-mode spin-lock busy-wait state, respectively.
CPU\$L_NULLCPU*	Clock ticks during which the null job has been the current process on this processor.
CPU\$W_UKERNEL*	Reserved to Digital.
CPU\$W_UNULLCPU*	Reserved to Digital.
CPU\$W_CLKUTICS*	Reserved to Digital.
CPU\$W_HARDAFF*	Count of processes with hard affinity for this processor.
CPU\$L_RANK_VEC*	Longword recording the ranks of all spin locks currently held by the processor. Spin lock acquisition code issues a Find First Set (FFS) instruction on this longword to determine if the processor holds any locks that are lower ranked than the one it seeks.
CPU\$L_IPL_VEC*	Vector recording, in inverse order, the IPLs of all spin locks currently held by the processor (that is, bit 0 represents IPL 31).
CPU\$L_IPL_ARRAY*	Array of 32 longwords, corresponding in inverse order to the 32 IPLs (that is, the first longword represents IPL 31). Upon each successful spin lock acquisition by this processor, the IPL vector corresponding to the spin lock's synchronization IPL (SPL\$B_IPL) is incremented.
CPU\$L_TPOINTER*	Address of the sanity timer (CPU\$W_SANITY_TIMER) of the active processor with the next highest CPU ID.
CPU\$W_SANITY_TIMER*	Number of sanity cycles before this processor times out.
CPU\$W_SANITY_TICKS*	Number of clock ticks until the next sanity cycle.
CPU\$L_VP_OWNER*	PCB address of the vector consumer.
CPU\$L_VP_VARIANT_EXIT*	Variant exit address to the disabled fault handler.

(continued on next page)

Data Structures

1.4 Per-CPU Database (CPU)

Table 1–4 (Cont.) Contents of Per-CPU Database

Field	Contents
CPU\$L_VP_FLAGS*	Vector processing flags. The following fields are defined within CPU\$L_VP_FLAGS: CPU\$V_VP_POWERFAIL Powerfail variant CPU\$V_VP_BUGCHECK Bugcheck variant CPU\$V_VP_CTX_INIT Initialization in progress for vector context CPU\$V_VP_CTX_SAVE Save in progress for vector context CPU\$V_VP_CTX_RESTORE Restore in progress for vector context
CPU\$L_VP_CPUTIM*	Scheduled time for a vector consumer.
CPU\$B_FLAGS*	Miscellaneous processor flags. The following fields are defined within CPU\$B_FLAGS: CPU\$V_SCHED Idle loop in wait for CPU scheduler CPU\$V_FOREVER STOP/CPU with /FOREVER qualifier CPU\$V_NEWPRIM Primary-to-be CPU
CPU\$L_INTFLAGS*	Interlocked flags. This word contains one flag bit: CPU\$V_STOPPING for the CPU stopping indicator.

1.5 Channel Request Block (CRB)

The activity of each controller in a configuration is described in a channel request block (CRB). This data structure contains pointers to the wait queue of drivers ready to gain access to a device through the controller. It also stores the entry points to the driver's interrupt service routines and unit/controller initialization routines.

The channel request block is illustrated in Figure 1–6 and described in Table 1–5.

Figure 1–6 Channel Request Block (CRB)

CRB\$L_FQFL			0
CRB\$L_FQBL			4
CRB\$B_FLCK	CRB\$B_TYPE*	CRB\$W_SIZE*	8
CRB\$L_FPC			12
CRB\$L_FR3			16
CRB\$L_FR4			20
CRB\$L_WQFL*			24

(continued on next page)

Data Structures

1.5 Channel Request Block (CRB)

Figure 1-6 (Cont.) Channel Request Block (CRB)

CRB\$L_WQBL*			28
Unused		CRB\$B_TT_TYPE*	32
CRB\$B_UNIT_BRK*	CRB\$B_MASK*	CRB\$W_REFC*	36
CRB\$L_AUXSTRUC			40
CRB\$L_TIMELINK*			44
CRB\$L_DUETIME*			48
CRB\$L_TOUTROUT*			52
CRB\$L_LINK*			56
CRB\$L_DLCK*			60
CRB\$L_BUGCHECK*			64
~	CRB\$L_RTINTD* (12 bytes)		~ 68
~	CRB\$L_INTD* (40 bytes)		~ 80
CRB\$L_BUGCHECK2*			120
~	CRB\$L_RTINTD2* (12 bytes)		~ 124
~	CRB\$L_INTD2* (40 bytes)		~ 136

*A read-only field

Data Structures

1.5 Channel Request Block (CRB)

Table 1–5 Contents of Channel Request Block

Field Name	Contents
CRB\$L_FQFL	Fork queue forward link. The link points to the next entry in the fork queue. Controller initialization routines write this field when they must drop IPL to utilize certain executive routines, such as those that allocate memory, that must be called at a lower IPL. The CRB timeout mechanism also uses the CRB fork block to lower IPL prior to calling the CRB timeout routine.
CRB\$L_FQBL	Fork queue backward link. The link points to the previous entry in the fork queue.
CRB\$W_SIZE*	Size of CRB. The driver-loading procedure writes this field when it creates the CRB.
CRB\$B_TYPE*	Type of data structure. The driver-loading procedure writes the symbolic constant DYN\$C_CRB into this field when it creates the CRB.
CRB\$B_FLCK	Fork lock at which the controller's fork operations are synchronized. If it must use the CRB fork block, a driver either uses a DPT_STORE macro to initialize this field or explicitly sets its value within the controller initialization routine.
CRB\$L_FPC	Address of instruction at which execution resumes when the VMS fork dispatcher dequeues the fork block. EXE\$FORK writes this field when called to suspend driver execution.
CRB\$L_FR3	Value of R3 at the time that the executing code requests VMS to create a fork block. EXE\$FORK writes this field when called to suspend driver execution.
CRB\$L_FR4	Value of R4 at the time that the executing code requests VMS to create a fork block. EXE\$FORK writes this field when called to suspend driver execution.
CRB\$L_WQFL*	Controller data channel wait queue forward link. IOC\$REQxCHANy and IOC\$RELxCHAN insert and remove driver fork block addresses in this field. A channel wait queue contains addresses of driver fork blocks that record the context of suspended drivers waiting to gain control of a controller data channel. If a channel is busy when a driver requests access to the channel, IOC\$REQxCHANy suspends the driver by saving the driver's context in the device's UCB fork block and inserting the fork block address in the channel wait queue. When a driver releases a channel because an I/O operation no longer needs the channel, IOC\$RELxCHAN dequeues a driver fork block, allocates the channel to the driver, and reactivates the suspended driver fork process. If no drivers are awaiting the channel, IOC\$RELxCHAN clears the channel busy bit.
CRB\$L_WQBL*	Controller channel wait queue backward link. IOC\$REQxCHANy and IOC\$RELxCHAN read and write this field.
CRB\$B_TT_TYPE*	Type of controller (for instance, DZ11 or DZ32) for terminals. A terminal port driver fills in this field.
CRB\$W_REFC*	UCB reference count. The driver-loading procedure increases the value in this field each time it creates a UCB for a device attached to the controller.
CRB\$B_MASK*	Mask that describes controller status. The following fields are defined in CRB\$B_MASK: CRB\$V_BSY Busy bit. IOC\$REQxCHANy reads the busy bit to determine whether the controller is free and sets this bit when it allocates the controller data channel to a driver. IOC\$RELxCHAN clears the busy bit if no driver is waiting to acquire the channel.

(continued on next page)

Data Structures

1.5 Channel Request Block (CRB)

Table 1-5 (Cont.) Contents of Channel Request Block

Field Name	Contents
	CRB\$V_UNINIT Indication, when set, that the VMS adapter initialization routine has created a CRB for a generic VAXBI device, but has not yet called its controller initialization routine. SYSGEN reads this bit to determine whether to call the controller initialization routine and clears it when the initialization routine completes. This facilitates SYSGEN's processing of multiunit generic VAXBI devices.
CRB\$B_UNIT_BRK*	Break bits for terminal lines. Used by VMS terminal port drivers.
CRB\$L_AUXSTRUC	Address of auxiliary data structure used by device driver to store special controller information. A device driver requiring such a structure generally allocates a block of nonpaged dynamic memory in its controller initialization routine and places a pointer to it in this field.
CRB\$L_TIMELINK*	Forward link in queue of CRBs waiting for periodic wakeups. This field points to the CRB\$L_TIMELINK field of the next CRB in the list. The CRB\$L_TIMELINK field of the last CRB in the list contains zero. The listhead for this queue is IOC\$GL_CRBTMOUT. Use of this field is reserved to Digital.
CRB\$L_DUETIME*	Time in seconds, relative to EXE\$GL_ABSTIM, at which next periodic wakeup associated with the CRB is to be delivered. Compute this value by raising IPL to IPL\$_POWER, adding the desired number of seconds to the contents of EXE\$GL_ABSTIM, and storing the result in this field. Use of this field is reserved to Digital.
CRB\$L_TOUTROUT*	Address of routine to be called at fork IPL (holding a corresponding fork lock if necessary) when a periodic wakeup associated with CRB becomes due. The routine must compute and reset the value in CRB\$L_DUETIME if another periodic wakeup request is desired. Use of this field is reserved to Digital.
CRB\$L_LINK*	Address of secondary CRB (for MASSBUS devices only). This field is written by the driver-loading procedure and read by IOC\$REQSCHANx and IOC\$RELSCHAN.
CRB\$L_DLCK*	Address of controller's device lock. The driver-loading procedure initializes this field and propagates it to each UCB it creates for the device units associated with the controller.
CRB\$L_BUGCHECK*	Bugcheck data used to issue an ILLQBUSCFG bugcheck when the multilevel interrupt dispatching code (at CRB\$L_RTINTD) determines that a Q22 bus is illegally configured.
CRB\$L_RTINTD*	Portion of interrupt transfer vector created at system initialization when a MicroVAX system implements multilevel device interrupt dispatching. The code stored in this 12-byte field implements a conditional lowering to device IPL. See Section 1.5.1 for a description of the contents of the interrupt transfer vector.

(continued on next page)

Table 1–5 (Cont.) Contents of Channel Request Block

Field Name	Contents
CRB\$_INTD*	<p>Interrupt transfer vector. This 10-longword field (described in Section 1.5.1) stores executable code, driver entry points, and I/O adapter information. It contains pointers to the driver's controller and unit initialization routines, the interrupt dispatch block (IDB), and the adapter control block (ADP). It may also contain fields that describe the disposition of a controller's data paths and map registers. The interrupt transfer routine is located at the top of the interrupt transfer vector.</p> <p>Although certain of the symbolic offsets defined in the data structure definition macro \$VECDEF have negative values, driver code can uniformly refer to the contents of the VEC structure in the following form:</p> $\text{CRB}\$_\text{INTD} + \text{VEC}\$_x_symbol.$
CRB\$_BUGCHECK2*	Bugcheck data used to issue an ILLQBUSCFG bugcheck when the multilevel interrupt dispatching code (at CRB\$_RTINTD2) determines that the Q22 bus is illegally configured.
CRB\$_RTINTD2*	Portion of second interrupt transfer vector initialized and used if multilevel interrupt dispatching is enabled in a MicroVAX system. See Section 1.5.1 for a description of the contents of the interrupt transfer vector.
CRB\$_INTD2*	Second interrupt transfer vector for devices with multiple interrupt vectors. The data structure definition macro \$CRBDEF supplies symbolic offsets for only the first two interrupt transfer vector structures.

1.5.1 Interrupt Transfer Vector Block (VEC)

VMS creates the appropriate number of interrupt transfer vector blocks (VEC) (shown in Figure 1–7) within a CRB if a driver specifies that the addresses of additional interrupt service routines be loaded into these structures. For example:

```
DPT_STORE,CRB,CRB$_INTD2+VEC$_ISR,D,isr_for_vec2
DPT_STORE,CRB,CRB$_INTD+<2*VEC$_K_LENGTH>+VEC$_ISR,D,isr_for_vec3
```

The offset of the *n*th vector located within the CRB is equal to the result of the following formula:

$$\text{CRB}\$_\text{INTD} + (n * \text{VEC}\$_\text{K_LENGTH})$$

VMS automatically initializes the interrupt dispatching instructions and the data structure locations from information located in the primary vector. The number of device vectors and vector structures actually created can be overridden by the value specified in the /NUMVEC qualifier to the SYSGEN command CONNECT. For a description of the fields in VEC, see Table 1–6.

Data Structures

1.5 Channel Request Block (CRB)

Figure 1-7 Interrupt Transfer Vector Block (VEC)

VEC\$L_BUGCHECK*			0
VEC\$L_RTINTD* (12 bytes)			4
VEC\$L_INTD*			16
VEC\$L_ISR			20
VEC\$L_IDB*			24
VEC\$L_INITIAL			28
VEC\$B_DATAPATH	VEC\$B_NUMREG	VEC\$W_MAPREG	32
VEC\$L_ADP*			36
VEC\$L_UNITINIT*			40
VEC\$L_START*			44
VEC\$L_UNITDISC*			48
VEC\$W_NUMALT		VEC\$W_MAPALT	52

*A read-only field

Table 1-6 Contents of Interrupt Transfer Vector Block (VEC)

Field Name	Contents
VEC\$L_BUGCHECK*	Bugcheck data used to issue an ILLQBUSCFG bugcheck when the multilevel interrupt dispatching code determines that the Q22 bus is illegally configured.
VEC\$L_RTINTD*	Portion of interrupt transfer vector created at system initialization when a MicroVAX system implements multilevel device interrupt dispatching. The code stored in this 12-byte field implements a conditional lowering to device IPL, as follows: <pre> CMPZV #PSL\$V_IPL, #PSL\$S_IPL, - 4(SP), S^#DIPL BGEQ BUGCHECK SETIPL S^#DIPL </pre>

(continued on next page)

Data Structures

1.5 Channel Request Block (CRB)

Table 1–6 (Cont.) Contents of Interrupt Transfer Vector Block (VEC)

Field Name	Contents								
VEC\$L_INTD*	<p>Interrupt dispatching code, written by the driver-loading procedure as follows:</p> <pre style="margin-left: 40px;">PUSHR #^M<R0, R1, R2, R3, R4, R5> JSB @#</pre> <p>The destination of the JSB instruction is the driver's interrupt service routine, as indicated at offset VEC\$L_ISR. Under normal operations, direct-vector UNIBUS or Q22-bus adapters—as well as VAXBI system interrupt dispatching—transfer control to CRB\$L_INTD. The code located here causes the processor to execute the PUSHHR instruction to save R0 through R5 on the stack and execute a JSB instruction to transfer control to the driver's interrupt service routine.</p> <p>In dispatching interrupts from non-direct-vector UNIBUS adapters, the UNIBUS adapter interrupt service routine transfers control to CRB\$L_INTD+2, which contains the JSB instruction to the driver's interrupt service routine. Because the UNIBUS adapter's interrupt service routine has already saved R0 through R5, interrupt dispatching bypasses the PUSHHR instruction in these instances.</p> <p>This field, plus VEC\$L_ISR, is also known as VEC\$Q_DISPATCH.</p>								
VEC\$L_ISR	<p>The DPT in every driver for an interrupting device specifies the address of a driver interrupt service routine.</p>								
VEC\$L_IDB*	<p>Address of IDB for controller. The driver-loading procedure creates an IDB for each CRB and loads the address of the IDB in this field. Device drivers use the IDB address to obtain the virtual addresses of device registers.</p> <p>When a driver's interrupt service routine gains control, the top of the stack contains a pointer to this field.</p>								
VEC\$L_INITIAL	<p>Address of controller initialization routine. If a device controller requires initialization at driver-loading time and during recovery from a power failure, the driver specifies a value for this field in the DPT.</p> <p>The driver-loading procedure calls this routine each time the procedure loads the driver. The VMS power failure recovery procedure also calls this routine to initialize a controller after a power failure.</p>								
VEC\$W_MAPREG	<p>The following bits are defined within VEC\$W_MAPREG:</p> <table style="width: 100%; border: none;"> <tr> <td style="width: 30%; vertical-align: top;">VEC\$V_MAPREG</td> <td>Number of first standard map register allocated to the driver that owns controller data channel.</td> </tr> <tr> <td></td> <td>IOC\$REQMAPREG writes this field when the routine allocates a set of standard map registers to a driver fork process for a DMA transfer. IOC\$RELMAPREG reads the field to deallocate a set of map registers.</td> </tr> <tr> <td></td> <td>Device drivers read this field in calculating the starting address of a UNIBUS or MicroVAX/Q22-bus transfer.</td> </tr> <tr> <td style="vertical-align: top;">VEC\$V_MAPLOCK</td> <td>Map register set is permanently allocated (when set).</td> </tr> </table>	VEC\$V_MAPREG	Number of first standard map register allocated to the driver that owns controller data channel.		IOC\$REQMAPREG writes this field when the routine allocates a set of standard map registers to a driver fork process for a DMA transfer. IOC\$RELMAPREG reads the field to deallocate a set of map registers.		Device drivers read this field in calculating the starting address of a UNIBUS or MicroVAX/Q22-bus transfer.	VEC\$V_MAPLOCK	Map register set is permanently allocated (when set).
VEC\$V_MAPREG	Number of first standard map register allocated to the driver that owns controller data channel.								
	IOC\$REQMAPREG writes this field when the routine allocates a set of standard map registers to a driver fork process for a DMA transfer. IOC\$RELMAPREG reads the field to deallocate a set of map registers.								
	Device drivers read this field in calculating the starting address of a UNIBUS or MicroVAX/Q22-bus transfer.								
VEC\$V_MAPLOCK	Map register set is permanently allocated (when set).								
VEC\$B_NUMREG	<p>Number of UNIBUS adapter or MicroVAX Q22-bus standard map registers allocated to driver. IOC\$REQMAPREG writes this 15-bit field when the routine allocates a set of standard map registers. IOC\$RELMAPREG reads this field to deallocate a set of standard map registers.</p>								

(continued on next page)

Data Structures

1.5 Channel Request Block (CRB)

Table 1-6 (Cont.) Contents of Interrupt Transfer Vector Block (VEC)

Field Name	Contents
VEC\$B_DATAPATH	Data path specifier. The bits that make up this field are used as follows:
VEC\$V_DATAPATH	Number of data path used in DMA transfer. The routine IOC\$REQDATAP writes this 5-bit field when a buffered data path is allocated and clears the field when the data path is released. The routine IOC\$LOADUBAMAP copies the contents of this field into UNIBUS adapter map registers. These bits also serve as implicit input to the IOC\$PURGDATAP routine.
VEC\$V_LWAE	Longword access enable (LWAE) bit. Drivers set this bit when they wish to limit the data path to longword-aligned, random-access mode. The routine IOC\$LOADUBAMAP copies the value in this field to the UNIBUS adapter map registers.
<6>	Reserved to Digital.
VEC\$V_PATHLOCK	Buffered data path allocation indicator. Drivers set this bit to specify that the buffered data path is permanently allocated.
VEC\$L_ADP*	Address of ADP. The SYSGEN command CONNECT must specify the nexus number of the UNIBUS adapter used by a controller. The driver-loading procedure writes the address of the ADP for the specified UBA into the VEC\$L_ADP field. IOC\$REQMAPREG, IOC\$REQALTMAP, and IOC\$RELMAPREG read and write fields in the ADP to allocate and deallocate map registers.
VEC\$L_UNITINIT*	Address of device driver's unit initialization routine. If a device unit requires initialization at driver-loading time and during recovery from a power failure, the driver specifies a value for this field in the DPT. The driver-loading procedure calls this routine for each device unit each time the procedure loads the driver. The VMS power failure recovery procedure also calls this routine to initialize device units after a power failure. MASSBUS drivers that support mixed device types must not use this field. Instead, they should specify the unit initialization routine in the unit initialization field of the DDT (DDT\$L_UNITINIT). Other drivers can use either field.
VEC\$L_START*	Address of VMS start protocol routine. Use of this field is reserved to Digital.
VEC\$L_UNITDISC*	Address of unit disconnect routine. Use of this field is reserved to Digital.
VEC\$W_MAPALT	The following bits are defined within VEC\$W_MAPALT:
VEC\$V_MAPALT	Number of first Q22-bus alternate map register allocated to driver that owns controller data channel. IOC\$REQALTMAP writes this field when the routine allocates a set of Q22-bus alternate map registers to a driver fork process for a DMA transfer. IOC\$RELMAPREG reads the field to deallocate a set of map registers. Device drivers read this 15-bit field in calculating the starting address of a MicroVAX Q22-bus transfer that uses a set of alternate map registers.
VEC\$V_ALTLOCK	Alternate map register set is permanently allocated (when set).

(continued on next page)

Data Structures

1.5 Channel Request Block (CRB)

Table 1–6 (Cont.) Contents of Interrupt Transfer Vector Block (VEC)

Field Name	Contents
VEC\$W_NUMALT	Number of Q22-bus alternate map registers allocated to driver. IOC\$REQALTMAP writes this field when allocating a set of alternate map registers. IOC\$RELMAPREG reads this field to deallocate a set of alternate map registers.

1.6 Device Data Block (DDB)

The device data block (DDB) is a block that identifies the generic device/controller name and driver name for a set of devices attached to a single controller. The driver-loading procedure creates a DDB for each controller during autoconfiguration at system startup and dynamically creates additional DDBs for new controllers as they are added to the system using the SYSGEN command CONNECT. The procedure initializes all fields in the DDB. All the DDBs in the I/O database are linked in a singly linked list. The contents of IOC\$GL_DEVLIST point to the first entry in the list.

VMS routines and device drivers refer to the DDB.

The device data block is illustrated in Figure 1–8 and described in Table 1–7.

Figure 1–8 Device Data Block (DDB)


DDB\$L_LINK*			0
DDB\$L_UCB*			4
Unused	DDB\$B_TYPE*	DDB\$W_SIZE*	8
DDB\$L_DDT			12
DDB\$L_ACPD			16
⋈	DDB\$T_NAME* (16 bytes)		⋈ 20
⋈	DDB\$T_DRVNAME* (16 bytes)		⋈ 36
DDB\$L_SB*			52
DDB\$L_CONLINK*			56
DDB\$L_ALLOCLS*			60

(continued on next page)

Data Structures

1.6 Device Data Block (DDB)

Figure 1–8 (Cont.) Device Data Block (DDB)



*A read-only field

Table 1–7 Contents of Device Data Block

Field Name	Contents
DDB\$L_LINK*	Address of next DDB. A zero indicates that this is the last DDB in the DDB chain.
DDB\$L_UCB*	Address of UCB for first unit attached to controller.
DDB\$W_SIZE*	Size of DDB.
DDB\$B_TYPE*	Type of data structure. The driver-loading procedure writes the constant DYN\$C_DDB into this field when the procedure creates the DDB.
DDB\$L_DDT	Address of DDT. VMS can transfer control to a device driver only through addresses listed in the DDT, the CRB, and the UCB fork block. The DPT of every device driver must specify a value for this field.
DDB\$L_ACPD	Name of default ACP (or XQP) for controller. ACPs that control access to file-structured devices (or the XQP) use the high-order byte of this field, DDB\$B_ACPCLASS, to indicate the class of the file-structured device. If the ACP_MULTIPLE system parameter is set, the initialization procedure creates a unique ACP for each class of file-structured device. Drivers initialize DDB\$B_ACPCLASS by invoking a DPT_STORE macro. Values for DDB\$B_ACPCLASS are as follows: DDB\$K_PACK Standard disk pack DDB\$K_CART Cartridge disk pack DDB\$K_SLOW Floppy disk DDB\$K_TAPE Magnetic tape that simulates file-structured device
DDB\$T_NAME*	Generic name for the devices attached to controller. The first byte of this field is the number of characters in the generic name. The remainder of the field consists of a string of up to 15 characters that, suffixed by a device unit number, identifies devices on the controller.
DDB\$T_DRVNAME*	Name of device driver for controller. The first byte of this field is the number of characters in the driver name. The remainder of the field contains a string of up to 15 characters taken from the DPT in the driver.
DDB\$L_SB*	Address of system block.
DDB\$L_CONLINK*	Address of next DDB in the connection subchain.
DDB\$L_ALLOCLS*	Allocation class of device.
DDB\$L_2P_UCB*	Address of the first UCB on the secondary path. Another name for this field is DDB\$L_DP_UCB.

1.7 Driver Dispatch Table (DDT)

Each device driver contains a driver dispatch table (DDT). The DDT lists entry points in the driver that VMS routines call, for instance, the entry point for the driver start-I/O routine.

A device driver creates a DDT by invoking the VMS macro DDTAB. The fields in the driver dispatch table are illustrated in Figure 1-9 and described in Table 1-8.

Figure 1-9 Driver Dispatch Table (DDT)

DDT\$L_START		0
DDT\$L_UN SOLINT		4
DDT\$L_FDT		8
DDT\$L_CANCEL		12
DDT\$L_REGDUMP		16
DDT\$W_ERRORBUF	DDT\$W_DIAGBUF	20
DDT\$L_UNITINIT		24
DDT\$L_ALTSTART		28
DDT\$L_MNTVER		32
DDT\$L_CLONEDUCB		36
Unused	DDT\$W_FDTSIZE*	40
DDT\$L_MNTV_SSSC*		44
DDT\$L_MNTV_FOR*		48
DDT\$L_MNTV_SQD*		52
DDT\$L_AUX_STORAGE*		56
DDT\$L_AUX_ROUTINE*		60

*A read-only field

Data Structures

1.7 Driver Dispatch Table (DDT)

Table 1-8 Contents of Driver Dispatch Table

Field Name	Contents
DDT\$L_START	<p>Entry point to the driver's start-I/O routine. Every driver must specify this address in the start argument to the DDTAB macro.</p> <p>When a device unit is idle and an I/O request is pending for that unit, IOC\$INITIATE transfers control to the address contained in this field.</p>
DDT\$L_UN SOLINT	<p>Entry point to a MASSBUS driver's unsolicited-interrupt service routine. The driver specifies this address in the unsolic argument to the DDTAB macro.</p> <p>This field contains the address of a routine that analyzes unexpected interrupts from a device. The standard interrupt service routine, the address of which is stored in the CRB, determines whether an interrupt was solicited by a driver. If the interrupt is unsolicited, the interrupt service routine can call the unsolicited-interrupt service routine.</p>
DDT\$L_FDT	<p>Address of the driver's FDT. Every driver must specify this address in the functb argument to the DDTAB macro.</p> <p>EXE\$QIO refers to the FDT to validate I/O function codes, decide which functions are buffered, and call FDT routines associated with function codes.</p>
DDT\$L_CANCEL	<p>Entry point to the driver's cancel-I/O routine. The driver specifies this address in the cancel argument to the DDTAB macro.</p> <p>Some devices require special cleanup processing when a process or a VMS routine cancels an I/O request before the I/O operation completes or when the last channel is deassigned. The \$DASSGN, \$DALLOC, and \$CANCEL system services cancel I/O requests.</p>
DDT\$L_REGDUMP	<p>Entry point to the driver's register dumping routine. The driver specifies this address in the regdmp argument to the DDTAB macro.</p> <p>IOC\$DIAGBUFILL, ERL\$DEVICERR, and ERL\$DEVICTMO call the address contained in this field to write device register contents into a diagnostic buffer or error message buffer.</p>
DDT\$W_DIAGBUF	<p>Size of diagnostic buffer. The driver specifies this value in the diagbf argument to the DDTAB macro. The value is the size in bytes of a diagnostic buffer for the device.</p> <p>When EXE\$QIO preprocesses an I/O request, it allocates a system buffer of the size recorded in this field (if it contains a nonzero value) if the process requesting the I/O has DIAGNOSE privilege and specifies a diagnostic buffer in the I/O request. IOC\$DIAGBUFILL fills the buffer after the I/O operation completes.</p>
DDT\$W_ERRORBUF	<p>Size of error message buffer. The driver specifies this value in the erlgbf argument to the DDTAB macro. The value is the size in bytes of an error message buffer for the device.</p> <p>If error logging is enabled and an error occurs during an I/O operation, the driver calls ERL\$DEVICERR or ERL\$DEVICTMO to allocate and write error-logging data into the error message buffer. IOC\$INITIATE and IOC\$REQCOM write values into the buffer if an error has occurred.</p>
DDT\$L_UNITINIT	<p>Address of the device's unit initialization routine, if one exists. Drivers for MASSBUS devices use this field rather than CRB\$_INTD+VEC\$_UNITINIT. Drivers for UNIBUS, VAXBI, and Q22 devices can use either field.</p>
DDT\$L_ALTSTART	<p>Address of a driver's alternate start-I/O routine. EXE\$ALTQUEPKT transfers control to the alternate start-I/O routine at this address.</p>

(continued on next page)

Data Structures

1.7 Driver Dispatch Table (DDT)

Table 1–8 (Cont.) Contents of Driver Dispatch Table

Field Name	Contents
DDT\$L_MNTVER	Address of the VMS routine (IOC\$MNTVER) called at the beginning and end of mount verification operation. The mntver argument to the DPTAB macro defaults to this routine. Use of the mntver argument to call any routine other than IOC\$MNTVER is reserved to Digital.
DDT\$L_CLONEDUCB	Address of routine to call when UCB is cloned.
DDT\$W_FDTSIZE*	Number of bytes in FDT. The driver-loading procedure uses this field to relocate addresses in the FDT to system virtual addresses.
DDT\$L_MNTV_SSSC*	Address of routine to call when performing mount verification for a shadow-set state change. Use of this field is reserved to Digital.
DDT\$L_MNTV_FOR*	Address of routine to call when performing mount verification for a foreign device. Use of this field is reserved to Digital.
DDT\$L_MNTV_SQD*	Address of routine to call when performing mount verification for a sequential device. Use of this field is reserved to Digital.
DDT\$L_AUX_STORAGE*	Address of auxiliary storage area. Use of this field is reserved to Digital.
DDT\$L_AUX_ROUTINE*	Address of auxiliary routine. Use of this field is reserved to Digital.

1.8 Driver Prologue Table (DPT)

When loading a device driver and its database into virtual memory, the driver-loading procedure finds the basic description of the driver and its device in a driver prologue table (DPT). The DPT provides the length, name, adapter type, and loading and reloading specifications for the driver.

A device driver creates a DPT by invoking the VMS macros DPTAB and DPT_STORE. The driver prologue table is illustrated in Figure 1–10 and described in Table 1–9.

Figure 1–10 Driver Prologue Table (DPT)

DPT\$L_FLINK*			0
DPT\$L_BLINK*			4
DPT\$B_REFC*	DPT\$B_TYPE*	DPT\$W_SIZE	8
DPT\$W_UCBSIZE		Unused	DPT\$B_ADPTYPE
DPT\$L_FLAGS			16
DPT\$W_REINITTAB		DPT\$W_INITTAB	
DPT\$W_MAXUNITS		DPT\$W_UNLOAD	
DPT\$W_DEFUNITS		DPT\$W_VERSION*	
			28

(continued on next page)

Data Structures

1.8 Driver Prologue Table (DPT)

Figure 1-10 (Cont.) Driver Prologue Table (DPT)

DPT\$W_VECTOR	DPT\$W_DELIVER	32
DPT\$T_NAME (12 bytes)		36
DPT\$Q_LINKTIME*		48
DPT\$L_ECOLEVEL*		56
DPT\$L_UCODE*		60
DPT\$Q_LMF_1*		64
DPT\$Q_LMF_2*		72
DPT\$Q_LMF_3*		80
DPT\$Q_LMF_4*		88
DPT\$Q_LMF_5*		96
DPT\$Q_LMF_6*		104
DPT\$Q_LMF_7*		112
DPT\$Q_LMF_8*		120
DPT\$W_DECW_SNAME*		

*A read-only field

Data Structures

1.8 Driver Prologue Table (DPT)

Table 1–9 Contents of Driver Prologue Table

Field Name	Contents																				
DPT\$L_FLINK*	Forward link to next DPT. The driver-loading procedure writes this field. The procedure links all DPTs in the system in a doubly linked list.																				
DPT\$L_BLINK*	Backward link to previous DPT. The driver-loading procedure writes this field.																				
DPT\$W_SIZE	Size in bytes of the driver. The DPTAB macro writes this field by subtracting the address of the beginning of the DPT from the address specified as the end argument to the DPTAB macro. The driver-loading procedure uses this value to determine the space needed in nonpaged system memory to load the driver.																				
DPT\$B_TYPE*	Type of data structure. The DPTAB macro always writes the symbolic constant DYN\$C_DPT into this field.																				
DPT\$B_REFC*	Number of DDBs that refer to the driver. The driver-loading procedure increments the value in this field each time the procedure creates another DDB that points to the driver's DDT.																				
DPT\$B_ADPTYPE	Type of adapter used by the devices using this driver. Every driver must specify the string "UBA", "MBA", "GENBI", "NULL", or "DR" as the value of the adapter argument to the DPTAB macro. Q22-bus drivers should specify "UBA" as the adapter type. The macro writes the value AT\$_UBA, AT\$_MBA, or AT\$_GENBI in this field.																				
DPT\$W_UCBSIZE	Size in bytes of the unit control block for a device that uses this driver. Every driver must specify a value for this field in the ucbsize argument to the DPTAB macro. The driver-loading procedure allocates blocks of nonpaged system memory of the specified size when creating UCBs for devices associated with the driver.																				
DPT\$L_FLAGS	Driver-loading flags. This field is also known as DPT\$B_FLAGS. The driver can specify any of a set of flags as the value of the flags argument to the DPTAB macro. The driver-loading procedure modifies its loading and reloading algorithm based on the settings of these flags. Flags defined in the flag field include the following: <table style="width: 100%; border: none;"> <tr> <td style="padding-left: 2em;">DPT\$V_SUBCNTRL</td> <td>Device is a subcontroller.</td> </tr> <tr> <td style="padding-left: 2em;">DPT\$V_SVP</td> <td>Device requires permanent system page to be allocated during driver loading.</td> </tr> <tr> <td style="padding-left: 2em;">DPT\$V_NOUNLOAD</td> <td>Driver cannot be reloaded.</td> </tr> <tr> <td style="padding-left: 2em;">DPT\$V_SCS</td> <td>SCS code must be loaded with this driver.</td> </tr> <tr> <td style="padding-left: 2em;">DPT\$V_DUSHADOW</td> <td>Driver is the shadowing disk class driver.</td> </tr> <tr> <td style="padding-left: 2em;">DPT\$V_SCSCI</td> <td>Common SCS/CI subroutines must be loaded with this driver.</td> </tr> <tr> <td style="padding-left: 2em;">DPT\$V_BVPSUBS</td> <td>Common BVP subroutines must be loaded with this driver.</td> </tr> <tr> <td style="padding-left: 2em;">DPT\$V_UCODE</td> <td>Driver has an associated microcode image.</td> </tr> <tr> <td style="padding-left: 2em;">DPT\$V_SMPMOD</td> <td>Driver has been designed to run in a VMS multiprocessing environment.</td> </tr> <tr> <td style="padding-left: 2em;">DPT\$V_DECW_DECODE</td> <td>Driver is a decoding class driver.</td> </tr> </table>	DPT\$V_SUBCNTRL	Device is a subcontroller.	DPT\$V_SVP	Device requires permanent system page to be allocated during driver loading.	DPT\$V_NOUNLOAD	Driver cannot be reloaded.	DPT\$V_SCS	SCS code must be loaded with this driver.	DPT\$V_DUSHADOW	Driver is the shadowing disk class driver.	DPT\$V_SCSCI	Common SCS/CI subroutines must be loaded with this driver.	DPT\$V_BVPSUBS	Common BVP subroutines must be loaded with this driver.	DPT\$V_UCODE	Driver has an associated microcode image.	DPT\$V_SMPMOD	Driver has been designed to run in a VMS multiprocessing environment.	DPT\$V_DECW_DECODE	Driver is a decoding class driver.
DPT\$V_SUBCNTRL	Device is a subcontroller.																				
DPT\$V_SVP	Device requires permanent system page to be allocated during driver loading.																				
DPT\$V_NOUNLOAD	Driver cannot be reloaded.																				
DPT\$V_SCS	SCS code must be loaded with this driver.																				
DPT\$V_DUSHADOW	Driver is the shadowing disk class driver.																				
DPT\$V_SCSCI	Common SCS/CI subroutines must be loaded with this driver.																				
DPT\$V_BVPSUBS	Common BVP subroutines must be loaded with this driver.																				
DPT\$V_UCODE	Driver has an associated microcode image.																				
DPT\$V_SMPMOD	Driver has been designed to run in a VMS multiprocessing environment.																				
DPT\$V_DECW_DECODE	Driver is a decoding class driver.																				

(continued on next page)

Data Structures

1.8 Driver Prologue Table (DPT)

Table 1–9 (Cont.) Contents of Driver Prologue Table

Field Name	Contents
	DPT\$V_TPALLOC Select the tape allocation class parameter.
	DPT\$V_SNAPSHOT Driver is certified for system snapshot.
	DPT\$V_NO_IDB_ Do not select IDB\$L_UCBLST for UCB vectors. DISPATCH
DPT\$W_INITTAB	Offset to driver initialization table. Every driver must specify a list of data structure fields and values to be written into the fields at the time that the driver-loading procedure creates the driver's data structures and loads the driver. The driver invokes the VMS macro DPT_STORE to specify these fields and their values.
DPT\$W_REINITTAB	Offset to driver-reinitialization table. Every driver must specify a list of data structure fields and values to be written into these fields at the time that the driver-loading procedure creates the driver's data structures and loads the driver or the driver is reloaded. The driver invokes the VMS macro DPT_STORE to specify these fields and their values.
DPT\$W_UNLOAD	Relative address of driver routine to be called when driver is reloaded. The driver specifies this field with the value of the unload argument to the DPTAB macro. The driver-loading procedure calls the driver unloading routine before reinitializing all device units associated with the driver.
DPT\$W_MAXUNITS	Maximum number of units on controller that this driver supports. Specify this value in the maxunits argument to the DPTAB macro. If no value is specified, the default is eight units.
DPT\$W_VERSION*	Version number that identifies format of DPT. The DPTAB macro automatically inserts a value in this field. SYSGEN checks its copy of the version number against the value stored in this field. If the values do not match, an error is generated. To correct the error, reassemble and relink the driver.
DPT\$W_DEFUNITS	Number of UCBs that the VMS autoconfiguration facility will automatically create. Drivers specify this number with the defunits argument to the DPTAB macro. If the driver also gives a value to DPT\$W_DELIVER, this field is also the number of times that the autoconfiguration facility calls the unit delivery routine.
DPT\$W_DELIVER	Relative address of the unit delivery routine that the VMS autoconfiguration facility calls for the number of UCBs specified in DPT\$W_DEFUNITS. The driver supplies the address of the unit delivery routine in the deliver argument to the DPTAB macro.
DPT\$W_VECTOR	Relative address of a driver-specific vector. A terminal class or port driver stores the address of its class or port entry vector table in this field.
DPT\$T_NAME	Name of the device driver. Field is 12 bytes. One byte records the length of the name string; the name string can be up to 11 characters. Drivers specify this field as the value of the name argument to the DPTAB macro. The driver-loading procedure compares the name of a driver to be loaded with the values in this field in all DPTs already loaded into system memory to ensure that it loads only one copy of a driver at a time.
DPT\$Q_LINKTIME*	Time and date at which driver was linked, taken from its image header.
DPT\$L_ECOLEVEL*	ECO level of driver, taken from its image header.

(continued on next page)

Data Structures

1.8 Driver Prologue Table (DPT)

Table 1–9 (Cont.) Contents of Driver Prologue Table

Field Name	Contents
DPT\$L_UCODE*	Address of associated microcode image, if DPT\$V_UCODE is set in DPT\$L_FLAGS. Use of this field is reserved to Digital.
DPT\$Q_LMF_1*	First of eight quadwords reserved to Digital for the use of the VMS license management facility. (The others are DPT\$Q_LMF_2, DPT\$Q_LMF_3, DPT\$Q_LMF_4, DPT\$Q_LMF_5, DPT\$Q_LMF_6, DPT\$Q_LMF_7, and DPT\$Q_LMF_8.)
DPT\$W_DECW_SNAME*	Offset to counted ASCII string used by decoding drivers.

1.9 Interrupt Dispatch Block (IDB)

The interrupt dispatch block (IDB) records controller characteristics. The driver-loading procedure creates and initializes this block when the procedure creates a CRB. The IDB points to the physical controller by storing the virtual address of the CSR. The CSR is the indirect pointer to all device unit registers.

The interrupt dispatch block is illustrated in Figure 1–11 and described in Table 1–10.

Figure 1–11 Interrupt Dispatch Block (IDB)

IDB\$L_CSR*				0
IDB\$L_OWNER				4
IDB\$B_VECTOR*	IDB\$B_TYPE*	IDB\$W_SIZE*		8
IDB\$B_COMBO_CSR*	IDB\$B_TT_ENABLE*	IDB\$W_UNITS*		12
Unused		IDB\$B_FLAGS*	IDB\$B_COMBO_VEC*	16
IDB\$L_SPL*				20
IDB\$L_ADP*				24
IDB\$L_UCBLST* (32 bytes)				28

*A read-only field

Data Structures

1.9 Interrupt Dispatch Block (IDB)

Table 1–10 Contents of Interrupt Dispatch Block

Field Name	Contents
IDB\$L_CSR*	<p>Address of CSR. The SYSGEN command CONNECT specifies the address of a device's CSR. The driver-loading procedure writes the system virtual equivalent of this address into the IDB\$L_CSR field. Device drivers set and clear bits in device registers by referencing all device registers at fixed offsets from the CSR address.</p> <p>The driver-loading procedure tests the value of this field. If the value is not a CSR address, it sets IDB\$V_NO_CSR in IDB\$L_FLAGS and places the device offline by clearing UCB\$V_ONLINE in UCB\$L_STS. In this event, it does not call the driver's controller and unit initialization routines.</p>
IDB\$L_OWNER	<p>Address of UCB of device that owns controller data channel. IOC\$REQx CHANy writes a UCB address into this field when the routine allocates a controller data channel to a driver. IOC\$RELx CHAN confirms that the proper driver fork process is releasing a channel by comparing the driver's UCB with the UCB stored in the IDB\$L_OWNER field. If the UCB addresses are the same, IOC\$RELx CHAN allocates the channel to a waiting driver by writing a new UCB address into the field. If no driver fork processes are waiting for the channel, IOC\$RELxCHAN clears the field.</p> <p>If the controller is a single-unit controller, the unit or controller initialization routine should write the UCB address of the single device into this field.</p>
IDB\$W_SIZE*	<p>Size of IDB. The driver-loading procedure writes the constant IDB\$K_LENGTH into this field when the procedure creates the IDB.</p>
IDB\$B_TYPE*	<p>Type of data structure. The driver-loading procedure writes the symbolic constant DYN\$C_IDB into this field when the procedure creates the IDB.</p>
IDB\$B_VECTOR*	<p>Interrupt vector number of the device, right-shifted by two bits. SYSGEN writes a value into this field using either the autoconfiguration database or the value specified in the /VECTOR qualifier to the CONNECT command. Drivers for devices that define the interrupt vector address through a device register must use this field to load that register during unit initialization and reinitialization after a power failure.</p>
IDB\$W_UNITS*	<p>Maximum number of units connected to the controller. The maximum number of units is specified in the DPT and can be overridden at driver-loading time.</p>
IDB\$B_TT_ENABLE*	<p>Reserved for use by the VMS terminal driver.</p>
IDB\$B_COMBO_CSR*	<p>Address of the start of CSRs for a multicontroller device such as the DMF32. (The name of this field is IDB\$B_COMBO_CSR_OFFSET.)</p>
IDB\$B_COMBO_VEC*	<p>Address of the start of interrupt vectors for a multicontroller device. (The name of this field is IDB\$B_COMBO_VECTOR_OFFSET.)</p>
IDB\$B_FLAGS*	<p>Flags associated with the IDB. The only flag currently defined is IDB\$V_NO_CSR. The driver loading procedure sets this flag if IDB\$L_CSR does not contain the address of a CSR.</p>
IDB\$L_SPL*	<p>Address of the device lock that—in a VMS multiprocessing environment—synchronizes access to device registers and those fields in the UCB accessed at device IPL.</p>

(continued on next page)

Data Structures

1.9 Interrupt Dispatch Block (IDB)

Table 1–10 (Cont.) Contents of Interrupt Dispatch Block

Field Name	Contents
IDB\$L_ADP*	Address of the adapter's ADP. The SYSGEN CONNECT command must specify the nexus number of the I/O adapter used by a device. The driver-loading procedure writes the address of the ADP for the specified I/O adapter into the IDB\$L_ADP field.
IDB\$L_UCBLST*	List of UCB addresses. The size of this field is the maximum number of units supported by the controller, as defined in the DPT. The maximum specified in the DPT can be overridden at driver load time. The driver-loading procedure writes a UCB address into this field every time the routine creates a new UCB associated with the controller.

1.10 I/O Request Packet (IRP)

When a user process queues a valid I/O request by issuing a \$QIO or \$QIOW system service, the service creates an I/O request packet (IRP). The IRP contains a description of the request and receives the status of the I/O processing as it proceeds.

The I/O request packet is illustrated in Figure 1–12 and described in Table 1–11. Note that the standard IRP contains space for fields required by VMS multiprocessing and the VMS class drivers. Under no circumstances should a driver not supplied by Digital use these fields.

Figure 1–12 I/O Request Packet (IRP)

IRP\$L_IOQFL			0
IRP\$L_IOQBL			4
IRP\$B_RMOD*	IRP\$B_TYPE*	IRP\$W_SIZE*	8
IRP\$L_PID*			12
IRP\$L_AST*			16
IRP\$L_ASTPRM*			20
IRP\$L_WIND*			24
IRP\$L_UCB*			28
IRP\$B_PRI*	IRP\$B_EFN*	IRP\$W_FUNC	32
IRP\$L_IOSB*			36
IRP\$W_STS		IRP\$W_CHAN*	40
IRP\$L_SVAPTE			44

(continued on next page)

Data Structures

1.10 I/O Request Packet (IRP)

Figure 1–12 (Cont.) I/O Request Packet (IRP)

	IRP\$L_BCNT	IRP\$W_BOFF	48
↪	IRP\$W_STS2	IRP\$L_BCNT	52
	IRP\$L_IOST1		56
	IRP\$L_IOST2		60
	IRP\$L_ABCNT		64
	IRP\$L_OBCNT		68
	IRP\$L_SEGVBN		72
	IRP\$L_DIAGBUF*		76
	IRP\$L_SEQNUM*		80
	IRP\$L_EXTEND		84
	IRP\$L_ARB*		88
	IRP\$L_KEYDESC*		92
⋈	Reserved (72 bytes)		96

*A read-only field

Table 1–11 Contents of an I/O Request Packet

Field Name	Contents
IRP\$L_IOQFL	I/O queue forward link. EXE\$INSERTIRP reads and writes this field when the routine inserts IRPs into a pending-I/O queue. IOC\$REQCOM reads and writes this field when the routine dequeues IRPs from a pending-I/O queue in order to send an IRP to a device driver.
IRP\$L_IOQBL	I/O queue backward link. EXE\$INSERTIRP and IOC\$REQCOM read and write these fields.
IRP\$W_SIZE*	Size of IRP. EXE\$QIO writes the symbolic constant IRP\$C_LENGTH into this field when the routine allocates and fills an IRP.
IRP\$B_TYPE*	Type of data structure. EXE\$QIO writes the symbolic constant DYN\$C_IRP into this field when the routine allocates and fills an IRP.

(continued on next page)

Data Structures

1.10 I/O Request Packet (IRP)

Table 1–11 (Cont.) Contents of an I/O Request Packet

Field Name	Contents
IRP\$B_RMOD*	Information used by I/O postprocessing. This field contains the same bit fields as the ACB\$B_RMOD field of an AST control block. For instance, the two bits defined at ACB\$V_MODE indicate the access mode of the process at time of the I/O request. EXE\$QIO obtains the processor access mode from the PSL and writes the value into this field.
IRP\$L_PID*	Process identification of the process that issued the I/O request. EXE\$QIO obtains the process identification from the PCB and writes the value into this field.
IRP\$L_AST*	Address of AST routine, if specified by the process in the I/O request. (This field is otherwise clear.) If the process specifies an AST routine address in the \$QIO call, EXE\$QIO writes the address in this field. During I/O postprocessing, the special kernel-mode AST routine queues a user mode AST to the requesting process if this field contains the address of an AST routine.
IRP\$L_ASTPRM*	Parameter sent as an argument to the AST routine specified by the user in the I/O request. If the process specifies an AST routine and a parameter to that AST routine in the \$QIO call, EXE\$QIO writes the parameter in this field. During I/O postprocessing, the special kernel-mode AST routine queues a user mode AST if the IRP\$L_AST field contains an address, and passes the value in IRP\$L_ASTPRM to the AST routine as an argument.
IRP\$L_WIND*	Address of window control block (WCB) that describes the file being accessed in the I/O request. EXE\$QIO writes this field if the I/O request refers to a file-structured device. An ACP or XQP reads this field. When a process gains access to a file on a file-structured device or creates a logical link between a file and a process I/O channel, the device ACP or XQP creates a WCB that describes the virtual-to-logical mapping of the file data on the disk. EXE\$QIO stores the address of this WCB in the IRP\$L_WIND field.
IRP\$L_UCB*	Address of UCB for the device assigned to the process's I/O channel. EXE\$QIO copies this value from the CCB.
IRP\$W_FUNC	I/O function code that identifies the function to be performed for the I/O request. The I/O request call specifies an I/O function code; EXE\$QIO and driver FDT routines map the code value to its most basic level (virtual → logical → physical) and copy the reduced value into this field. Based on this function code, EXE\$QIO calls FDT action routines to preprocess an I/O request. Six bits of the function code describe the basic function. The remaining 10 bits modify the function.
IRP\$B_EFN*	Event flag number and group specified in I/O request. If the I/O request call does not specify an event flag number, EXE\$QIO uses event flag 0 by default. EXE\$QIO writes this field. The I/O postprocessing routine calls SCH\$POSTEF to set this event flag when the I/O operation is complete.
IRP\$B_PRI*	Base priority of the process that issued the I/O request. EXE\$QIO obtains a value for this field from the process's PCB. EXE\$INSERTIRP reads this field to insert an IRP into a priority-ordered pending-I/O queue.

(continued on next page)

Data Structures

1.10 I/O Request Packet (IRP)

Table 1–11 (Cont.) Contents of an I/O Request Packet

Field Name	Contents																																
IRP\$L_IOSB*	<p>Virtual address of the process's I/O status block (IOSB) that receives final status of the I/O request at I/O completion. EXE\$QIO writes a value into this field if the I/O request call specifies an IOSB address. (This field is otherwise clear.) The I/O postprocessing special kernel-mode AST routine writes two longwords of I/O status into the IOSB after the I/O operation is complete.</p> <p>When an FDT routine aborts an I/O request by calling EXE\$ABORTIO, EXE\$ABORTIO fills the IRP\$L_IOSB field with zeros so that I/O postprocessing does not write status into the IOSB.</p>																																
IRP\$W_CHAN*	Index number of process I/O channel for request. EXE\$QIO writes this field.																																
IRP\$W_STS	<p>Status of I/O request. EXE\$QIO initializes this field to 0. EXE\$QIO, FDT routines, and driver fork processes modify this field according to the current status of the I/O request. I/O postprocessing reads this field to determine what sort of postprocessing is necessary (for example, deallocate system buffers and adjust quota usage).</p> <p>Bits in the IRP\$W_STS field describe the type of I/O function, as follows:</p> <table border="0"> <tr> <td>IRP\$V_BUFIO</td> <td>Buffered-I/O function</td> </tr> <tr> <td>IRP\$V_FUNC</td> <td>Read function</td> </tr> <tr> <td>IRP\$V_PAGIO</td> <td>Paging-I/O function</td> </tr> <tr> <td>IRP\$V_COMPLX</td> <td>Complex-buffered-I/O function</td> </tr> <tr> <td>IRP\$V_VIRTUAL</td> <td>Virtual-I/O function</td> </tr> <tr> <td>IRP\$V_CHAINED</td> <td>Chained-buffered-I/O function</td> </tr> <tr> <td>IRP\$V_SWAPIO</td> <td>Swapping-I/O function</td> </tr> <tr> <td>IRP\$V_DIAGBUF</td> <td>Diagnostic buffer is present</td> </tr> <tr> <td>IRP\$V_PHYSIO</td> <td>Physical-I/O function</td> </tr> <tr> <td>IRP\$V_TERMIO</td> <td>Terminal I/O (for priority increment calculation)</td> </tr> <tr> <td>IRP\$V_MBXIO</td> <td>Mailbox-I/O function</td> </tr> <tr> <td>IRP\$V_EXTEND</td> <td>An extended IRP is linked to this IRP</td> </tr> <tr> <td>IRP\$V_FILACP</td> <td>File ACP I/O</td> </tr> <tr> <td>IRP\$V_MVIRP</td> <td>Mount-verification I/O function</td> </tr> <tr> <td>IRP\$V_SRVIO</td> <td>Server-type I/O</td> </tr> <tr> <td>IRP\$V_KEY</td> <td>Encrypted function (encryption key address at IRP\$L_KEYDESC)</td> </tr> </table>	IRP\$V_BUFIO	Buffered-I/O function	IRP\$V_FUNC	Read function	IRP\$V_PAGIO	Paging-I/O function	IRP\$V_COMPLX	Complex-buffered-I/O function	IRP\$V_VIRTUAL	Virtual-I/O function	IRP\$V_CHAINED	Chained-buffered-I/O function	IRP\$V_SWAPIO	Swapping-I/O function	IRP\$V_DIAGBUF	Diagnostic buffer is present	IRP\$V_PHYSIO	Physical-I/O function	IRP\$V_TERMIO	Terminal I/O (for priority increment calculation)	IRP\$V_MBXIO	Mailbox-I/O function	IRP\$V_EXTEND	An extended IRP is linked to this IRP	IRP\$V_FILACP	File ACP I/O	IRP\$V_MVIRP	Mount-verification I/O function	IRP\$V_SRVIO	Server-type I/O	IRP\$V_KEY	Encrypted function (encryption key address at IRP\$L_KEYDESC)
IRP\$V_BUFIO	Buffered-I/O function																																
IRP\$V_FUNC	Read function																																
IRP\$V_PAGIO	Paging-I/O function																																
IRP\$V_COMPLX	Complex-buffered-I/O function																																
IRP\$V_VIRTUAL	Virtual-I/O function																																
IRP\$V_CHAINED	Chained-buffered-I/O function																																
IRP\$V_SWAPIO	Swapping-I/O function																																
IRP\$V_DIAGBUF	Diagnostic buffer is present																																
IRP\$V_PHYSIO	Physical-I/O function																																
IRP\$V_TERMIO	Terminal I/O (for priority increment calculation)																																
IRP\$V_MBXIO	Mailbox-I/O function																																
IRP\$V_EXTEND	An extended IRP is linked to this IRP																																
IRP\$V_FILACP	File ACP I/O																																
IRP\$V_MVIRP	Mount-verification I/O function																																
IRP\$V_SRVIO	Server-type I/O																																
IRP\$V_KEY	Encrypted function (encryption key address at IRP\$L_KEYDESC)																																
IRP\$L_SVAPTE	<p>For a <i>direct-I/O</i> transfer, virtual address of the first page-table entry (PTE) of the I/O-transfer buffer, written here by the FDT routine locking process pages; for <i>buffered-I/O</i> transfer, address of a buffer in system address space, written here by the FDT routine allocating buffer.</p> <p>IOC\$INITIATE copies this field into UCB\$L_SVAPTE before transferring control to a device driver start-I/O routine.</p> <p>I/O postprocessing uses this field to deallocate the system buffer for a buffered-I/O transfer or to unlock pages locked for a direct-I/O transfer.</p>																																

(continued on next page)

Data Structures

1.10 I/O Request Packet (IRP)

Table 1–11 (Cont.) Contents of an I/O Request Packet

Field Name	Contents														
IRP\$W_BOFF	<p>Byte offset into the first page of a direct-I/O transfer. FDT routines calculate this offset and write the field.</p> <p>For buffered-I/O transfers, FDT routines must write the number of bytes to be charged to the process in this field because these bytes are being used for a system buffer.</p> <p>IOC\$INITIATE copies this field into UCB\$W_BOFF before calling a device driver start-I/O routine.</p> <p>I/O postprocessing uses IRP\$W_BOFF in conjunction with IRP\$L_BCNT and IRP\$L_SVAPTE to unlock pages locked for direct I/O. For buffered I/O, I/O postprocessing adds the value of IRP\$W_BOFF to the process byte count quota.</p>														
IRP\$L_BCNT	<p>Byte count of the I/O transfer. FDT routines calculate the count value and write the field. IOC\$INITIATE copies the low-order word of this field into UCB\$W_BCNT before calling a device driver's start-I/O routine.</p> <p>For a buffered-I/O-read function, I/O postprocessing uses IRP\$L_BCNT to determine how many bytes of data to write to the user's buffer.</p> <p>The field IRP\$W_BCNT points to the low-order word of this field to provide compatibility with previous versions of VMS.</p>														
IRP\$W_STS2	<p>Second word of I/O request status. EXE\$QIO initializes this field to 0. EXE\$QIO, FDT routines, and driver fork processes modify this field according to the current status of the I/O request.</p> <p>Bits in the IRP\$W_STS2 field describe the type of I/O function, as follows:</p> <table style="width: 100%; border: none;"> <tr> <td style="padding-left: 2em;">IRP\$V_START_PAST_HWM</td> <td style="padding-left: 2em;">I/O starts past file highwater mark.</td> </tr> <tr> <td style="padding-left: 2em;">IRP\$V_END_PAST_HWM</td> <td style="padding-left: 2em;">I/O ends past file highwater mark.</td> </tr> <tr> <td style="padding-left: 2em;">IRP\$V_ERASE</td> <td style="padding-left: 2em;">Erase I/O function.</td> </tr> <tr> <td style="padding-left: 2em;">IRP\$V_PART_HWM</td> <td style="padding-left: 2em;">Partial file highwater mark update.</td> </tr> <tr> <td style="padding-left: 2em;">IRP\$V_LCKIO</td> <td style="padding-left: 2em;">Locked I/O request, as used by DECnet direct I/O.</td> </tr> <tr> <td style="padding-left: 2em;">IRP\$V_SHDIO</td> <td style="padding-left: 2em;">Shadowing IRP.</td> </tr> <tr> <td style="padding-left: 2em;">IRP\$V_CACHEIO</td> <td style="padding-left: 2em;">I/O using VBN cache buffers.</td> </tr> </table>	IRP\$V_START_PAST_HWM	I/O starts past file highwater mark.	IRP\$V_END_PAST_HWM	I/O ends past file highwater mark.	IRP\$V_ERASE	Erase I/O function.	IRP\$V_PART_HWM	Partial file highwater mark update.	IRP\$V_LCKIO	Locked I/O request, as used by DECnet direct I/O.	IRP\$V_SHDIO	Shadowing IRP.	IRP\$V_CACHEIO	I/O using VBN cache buffers.
IRP\$V_START_PAST_HWM	I/O starts past file highwater mark.														
IRP\$V_END_PAST_HWM	I/O ends past file highwater mark.														
IRP\$V_ERASE	Erase I/O function.														
IRP\$V_PART_HWM	Partial file highwater mark update.														
IRP\$V_LCKIO	Locked I/O request, as used by DECnet direct I/O.														
IRP\$V_SHDIO	Shadowing IRP.														
IRP\$V_CACHEIO	I/O using VBN cache buffers.														
IRP\$L_IOST1	<p>First I/O status longword. IOC\$REQCOM and EXE\$FINISHIO(C) write the contents of R0 into this field. The I/O postprocessing routine copies the contents of this field into the user's IOSB.</p> <p>EXE\$ZEROPARM copies a 0 and EXE\$ONEPARM copies p1 into this field. This field is a good place to put a \$QIO request argument (p1 through p6) or a computed value.</p> <p>This field is also called IRP\$L_MEDIA.</p>														
IRP\$L_IOST2	<p>Second I/O status longword. IOC\$REQCOM, EXE\$FINISHIO, and EXE\$FINISHIOC write the contents of R1 into this field. The I/O postprocessing routine copies the contents of this field into the user's IOSB.</p> <p>The low byte of this field is also known as IRP\$B_CARCON. IRP\$B_CARCON contains carriage control instructions to the driver. EXE\$READ and EXE\$WRITE copy the contents of p4 of the user's I/O request into this field.</p>														

(continued on next page)

Data Structures

1.10 I/O Request Packet (IRP)

Table 1–11 (Cont.) Contents of an I/O Request Packet

Field Name	Contents
IRP\$L_ABCNT	Accumulated bytes transferred in virtual I/O transfer. IOC\$IOPPOST reads and writes this field after a partial virtual transfer. The symbol IRP\$W_ABCNT points to the low-order word of this field to provide compatibility with previous versions of VMS.
IRP\$L_OBCNT	Original transfer byte count in a virtual I/O transfer. IOC\$IOPPOST reads this field to determine whether a virtual transfer is complete, or whether another I/O request is necessary to transfer the remaining bytes. The symbol IRP\$W_OBCNT points to the low-order word of this field to provide compatibility with previous versions of VMS.
IRP\$L_SEGVBN	Virtual block number of the current segment of a virtual I/O transfer. IOC\$IOPPOST writes this field after a partial virtual transfer.
IRP\$L_DIAGBUF*	Address of a diagnostic buffer in system address space. If the I/O request call specifies a diagnostic buffer and if a diagnostic buffer length is specified in the DDT, and if the process has diagnostic privilege, EXE\$QIO copies the buffer address into this field. EXE\$QIO allocates a diagnostic buffer in system address space to be filled by IOC\$DIAGBUFILL during I/O processing. During I/O postprocessing, the special kernel-mode AST routine copies diagnostic data from the system buffer into the process diagnostic buffer.
IRP\$L_SEQNUM*	I/O transaction sequence number. If an error is logged for the request, this field contains the universal error log sequence number.
IRP\$L_EXTEND	Address of an IRPE linked to this IRP. FDT routines write an extension address to this field when a device requires more context than the IRP can accommodate. This field is read by IOC\$IOPPOST. IRP\$V_EXTEND in IRP\$W_STS is set if this extension address is used.
IRP\$L_ARB*	Address of access rights block (ARB). This block is located in the PCB and contains the process privilege mask and UIC, which are set up as follows: ARB\$Q_PRIV Quadword containing process privilege mask SPARE\$L Unused longword ARB\$L_UIC Longword containing process UIC
IRP\$L_KEYDESC	Address of encryption key.

1.11 I/O Request Packet Extension (IRPE)

I/O request packet extensions (IRPEs) hold additional I/O request information for devices that require more context than the standard IRP can accommodate. IRP extensions are also used when more than one buffer (region) must be locked into memory for a direct-I/O operation, or when a transfer requires a buffer that is larger than 64K. An IRPE provides space for two buffer regions, each with a 32-bit byte count.

FDT routines allocate IRPEs by calling EXE\$ALLOCIRP. Driver routines link the IRPE to the IRP, store the IRPE's address in IRP\$L_EXTEND, and set the bit field IRP\$V_EXTEND in IRP\$W_STS to show that an IRPE exists for the IRP. The FDT routine initializes the contents of the IRPE. Any fields within the extension not described in Table 1–12 can store driver-dependent information.

Data Structures

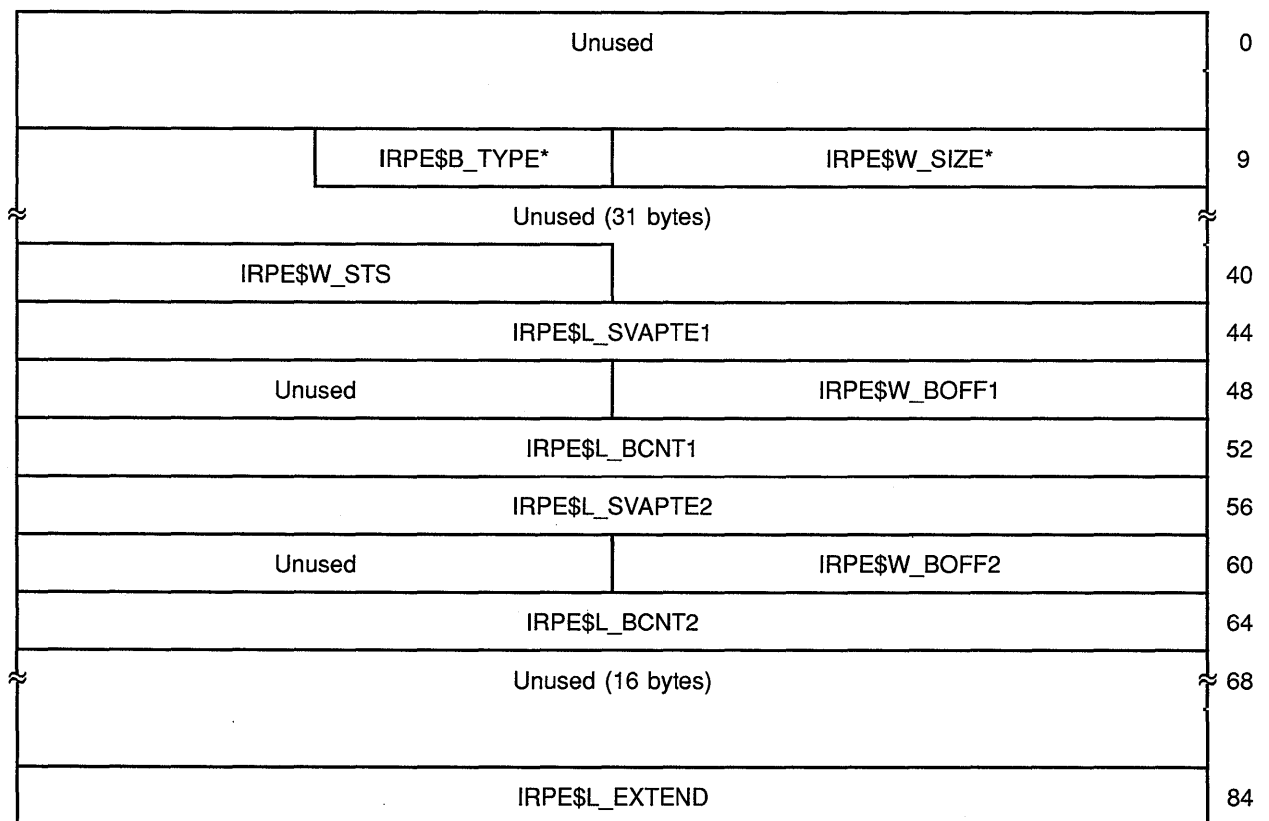
1.11 I/O Request Packet Extension (IRPE)

If the IRP extension specifies additional buffer regions, the FDT routine must use those buffer locking routines that perform coroutine calls back to the driver if the locking procedure fails (EXE\$READLOCKR, EXE\$WRITELOCKR, and EXE\$MODIFYLOCKR). If an error occurs during the locking procedure, the driver must unlock all previously locked regions using MMG\$UNLOCK and deallocate the IRPE before returning to the buffer locking routine.

IOC\$IOPOST automatically unlocks the pages in region 1 (if defined) and region 2 (if defined) for all the IRPEs linked to the IRP undergoing completion processing. IOC\$IOPOST also deallocates all the IRPEs.

The I/O request packet extension is illustrated in Figure 1-13 and described in Table 1-12.

Figure 1-13 I/O Request Packet Extension (IRPE)



*A read-only field

Data Structures

1.11 I/O Request Packet Extension (IRPE)

Table 1–12 Contents of the I/O Request Packet Extension

Field Name	Contents
IRPE\$W_SIZE*	Size of IRPE. EXE\$ALLOCIRP writes the constant IRP\$C_LENGTH to this field.
IRPE\$B_TYPE*	Type of data structure. EXE\$ALLOCIRP writes the constant DYN\$C_IRP to this field.
IRPE\$W_STS	IRPE status field. If bit IRPE\$V_EXTENDIRPE is set, it indicates that another IRPE is linked to this one.
IRPE\$L_SVAPTE1	System virtual address of the page-table entry (PTE) that maps the start of region 1. FDT routines write this field. If the region is not defined, this field is zero.
IRPE\$W_BOFF1	Byte offset of region 1. FDT routines write this field.
IRPE\$L_BCNT1	Size in bytes of region 1. FDT routines write this field.
IRPE\$L_SVAPTE2	System virtual address of the PTE that maps the start of region 2. Set by FDT routines. This field contains a value of zero if region 2 is not defined.
IRPE\$W_BOFF2	Byte offset of region 2. This field is set by FDT routines.
IRPE\$L_BCNT2	Size in bytes of region 2. FDT routines write this field.
IRPE\$L_EXTEND	Address of next IRPE for this IRP, if any.

1.12 Object Rights Block (ORB)

The object rights block (ORB) is a data structure that describes the rights a process must have in order to access the object with which the ORB is associated.

The ORB is usually allocated when the device is connected by means of SYSGEN's CONNECT command. SYSGEN also sets the address of the ORB in UCB\$L_ORB at that time.

The object rights block is illustrated in Figure 1–14 and described in Table 1–13.

Figure 1–14 Object Rights Block (ORB)

ORB\$L_OWNER			0
ORB\$L_ACL_MUTEX			4
ORB\$B_FLAGS	ORB\$B_TYPE*	ORB\$W_SIZE*	8
ORB\$W_REFCOUNT		Unused	12
ORB\$Q_MODE_PROT			16
ORB\$L_SYS_PROT			24
ORB\$L_OWN_PROT			28

(continued on next page)

Data Structures

1.12 Object Rights Block (ORB)

Figure 1–14 (Cont.) Object Rights Block (ORB)

ORB\$L_GRP_PROT	32
ORB\$L_WOR_PROT	36
ORB\$L_ACLFL	40
ORB\$L_ACLBL	44
~ ORB\$K_MIN_CLASS (20 bytes) ~	~ 48 ~
~ ORB\$K_MAX_CLASS (20 bytes) ~	~ 88 ~

*A read-only field

Table 1–13 Contents of Object Rights Block

Field	Contents
ORB\$L_OWNER	UIC of the object's owner.
ORB\$L_ACL_MUTEX	Mutex for the object's ACL, used to control access to the ACL for reading and writing. The driver-loading procedure initializes this field with -1.
ORB\$W_SIZE*	Size in bytes of ORB. The driver-loading procedure writes the symbolic constant ORB\$K_LENGTH into this field when it creates an ORB.
ORB\$B_TYPE*	Type of data structure. The driver-loading procedure writes the symbolic constant DYN\$C_ORB into this field when it creates an ORB.
ORB\$B_FLAGS	Flags needed for interpreting portions of the ORB that can have alternate meanings. The following fields are defined within ORB\$B_FLAGS:
ORB\$V_PROT_16	The driver-loading procedure sets this bit to 1, signifying SOGW protection.
ORB\$V_ACL_QUEUE	This flag represents the existence of an ACL queue. The driver-loading procedure does not set this bit.
ORB\$V_MODE_VECTOR	Use vector mode protection, not byte mode.
ORB\$V_NOACL	This object cannot have an ACL.
ORB\$V_CLASS_PROT	Security classification is valid.
ORB\$W_REFCOUNT	Reference count.
ORB\$Q_MODE_PROT	Mode protection vector. The low byte of this quadword is known as ORB\$B_MODE.
ORB\$L_SYS_PROT	System protection field. The low word of this field is known as ORB\$W_PROT and contains the standard SOGW protection.
ORB\$L_OWN_PROT	Owner protection field.

(continued on next page)

Data Structures

1.12 Object Rights Block (ORB)

Table 1–13 (Cont.) Contents of Object Rights Block

Field	Contents
ORB\$L_GRP_PROT	Group protection field.
ORB\$L_WOR_PROT	World protection field.
ORB\$L_ACLFL	ACL queue forward link. If ORB\$V_ACL_QUEUE is 0, this field should contain 0. This field is also known as ORB\$L_ACL_COUNT and is cleared by the driver-loading procedure.
ORB\$L_ACLBL	ACL queue backward link. If ORB\$V_ACL_QUEUE is 0, this field should contain 0. This field is also known as ORB\$L_ACL_DESC and is cleared by the driver-loading procedure.
ORB\$R_MIN_CLASS	Minimum classification mask.
ORB\$R_MAX_CLASS	Maximum classification mask.

1.13 SCSI Class Driver Request Packet (SCDRP)

The SCSI class driver allocates and builds a SCSI class driver request packet (SCDRP) for each I/O request it services, passing it to the SCSI port driver. The class driver routine initializes the SCDRP with the addresses of the UCB, SCDT, and IRP and copies to it data obtained from the IRP. The SCDRP also contains the addresses of the SCSI command buffer and status buffer.

The SCSI class driver passes the address of the SCDRP to the port driver in the call to SPI\$SEND_COMMAND.

The SCDRP is illustrated in Figure 1–15 and described in Table 1–14.

Figure 1–15 SCSI Class Driver Request Packet (SCDRP)

SCDRP\$L_FQFL		0
SCDRP\$L_FQBL		4
SCDRP\$B_FLCK	SCDRP\$B_CD_TYPE	8
SCDRP\$L_FPC		12
SCDRP\$L_FR3		16
SCDRP\$L_FR4		20
SCDRP\$L_PORT_UCB		24
SCDRP\$L_UCB		28
SCDRP\$W_STS	SCDRP\$W_FUNC	32

(continued on next page)

Data Structures

1.13 SCSI Class Driver Request Packet (SCDRP)

Figure 1–15 (Cont.) SCSI Class Driver Request Packet (SCDRP)

SCDRP\$L_SVAPTE		36
Reserved	SCDRP\$W_BOFF	40
SCDRP\$L_BCNT		44
SCDRP\$L_MEDIA		48
SCDRP\$L_ABCNT		52
SCDRP\$L_SAVD_RTN		56
Reserved		60
SCDRP\$L_CDT		68
Reserved		72
SCDRP\$L_IRP		76
SCDRP\$L_SVA_USER		80
SCDRP\$L_CMD_BUF		84
SCDRP\$L_CMD_BUF_LEN		88
SCDRP\$L_CMD_PTR		92
SCDRP\$L_STS_PTR		96
SCDRP\$L_SCSI_FLAGS		100
SCDRP\$L_DATACHECK		104
SCDRP\$L_SCSI_STK_PTR		108
~	SCDRP\$L_SCSI_STK (32 bytes)	~112
SCDRP\$L_CL_RETRY		144
SCDRP\$L_DMA_TIMEOUT		148
SCDRP\$L_DISCON_TIMEOUT		152
Reserved	SCDRP\$W_PAD_BCNT	156

(continued on next page)

Data Structures

1.13 SCSI Class Driver Request Packet (SCDRP)

Figure 1–15 (Cont.) SCSI Class Driver Request Packet (SCDRP)

SCDRP\$B_TQE* (52 bytes)		2160
SCDRP\$L_TQE_DELAY*		212
SCDRP\$L_SVA_DMA*		216
SCDRP\$L_SVA_CMD*		220
SCDRP\$W_CMD_MAPREG*	SCDRP\$W_MAPREG*	224
SCDRP\$W_CMD_NUMREG*	SCDRP\$W_NUMREG*	228
SCDRP\$L_SVA_SPT*		232
SCDRP\$L_SCSIMSGO_PTR*		236
SCDRP\$L_SCSIMSGI_PTR*		240
SCDRP\$B_SCSIMSGO_BUF*		244
SCDRP\$B_SCSIMSGI_BUF*		248
SCDRP\$L_MSGO_PENDING*		256
SCDRP\$L_MSGI_PENDING*		260
Reserved	SCDRP\$B_LAST_MSGO*	264
SCDRP\$L_DATA_PTR*		268
SCDRP\$L_TRANS_CNT*		272
SCDRP\$L_SAVE_DATA_CNT*		276
SCDRP\$L_SAVE_DATA_PTR*		280
SCDRP\$L_SDP_DATA_CNT*		284
SCDRP\$L_SDP_DATA_PTR*		288
SCDRP\$L_DUETIME*		292
SCDRP\$L_TIMEOUT_ADDR*		296
SCDRP\$W_BUSY_RETRY_CNT*	SCDRP\$W_CMD_BCNT*	300
SCDRP\$W_SEL_RETRY_CNT*	SCDRP\$W_ARB_RETRY_CNT*	304

(continued on next page)

Data Structures

1.13 SCSI Class Driver Request Packet (SCDRP)

Figure 1–15 (Cont.) SCSI Class Driver Request Packet (SCDRP)

SCDRP\$W_SEL_TQE_RETRY_CNT*	SCDRP\$W_CMD_RETRY_CNT*	308
SCDRP\$L_SAVER3*		312
SCDRP\$L_SAVER6*		316
SCDRP\$L_SAVER7*		320
SCDRP\$L_SAVER3CL*		324
SCDRP\$L_SAVEPCCL*		328
SCDRP\$L_ABORTPCCL*		332
SCDRP\$L_PO_STK_PTR*		336
~	SCDRP\$L_PO_STK* (24 bytes)	~340
SCDRP\$L_TAG*		364
~	Reserved (40 bytes)	~368

*A read-only field

Table 1–14 Contents of SCSI Class Driver Request Packet

Field Name	Contents
SCDRP\$L_FQFL	Fork queue forward link. This field points to the next entry in the SCSI adapter's command buffer wait queue (ADP\$L_BVPWAITFL), map register wait queue (ADP\$L_MRQFL), port wait queue (SPDT\$L_PORT_WQFL), or system fork queue.
SCDRP\$L_FQBL	Fork queue backward link. This field points to the previous entry in the SCSI adapter's command buffer wait queue (ADP\$L_BVPWAITFL), map register wait queue (ADP\$L_MRQFL), port wait queue (SPDT\$L_PORT_WQFL), or system fork queue.
SCDRP\$W_SCDRPSIZE	Size of SCDRP. A SCSI class driver, after allocating sufficient nonpaged pool for the SCDRP, writes the constant SCDRP\$C_LENGTH into this field.
SCDRP\$B_CD_TYPE	Class driver type. This field is currently unused.

(continued on next page)

Data Structures

1.13 SCSI Class Driver Request Packet (SCDRP)

Table 1–14 (Cont.) Contents of SCSI Class Driver Request Packet

Field Name	Contents
SCDRP\$B_FLCK	Index of the fork lock that synchronizes access to this SCDRP at fork level. A SCSI class driver, after allocating sufficient nonpaged pool for the SCDRP, copies to this field the value of UCB\$B_FLCK. All devices controlled by a single SCSI adapter and actively competing for shared adapter resources must specify the same value for this field.
SCDRP\$L_FPC	Address of instruction at which processing resumes when SCSI adapter resources become available to satisfy a request stalled in an adapter resource wait queue.
SCDRP\$L_FR3	Value of R3 when the request is stalled to wait for SCSI adapter resources. When the request is satisfied, this value is restored to R3 before the driver resumes execution at SCDRP\$L_FPC.
SCDRP\$L_FR4	Value of R4 when the request is stalled to wait for SCSI adapter resources. When the request is satisfied, this value is restored to R4 before the driver resumes execution at SCDRP\$L_FPC.
SCDRP\$L_PORT_UCB	SCSI adapter's UCB address. The SCSI port driver reads and writes this field in order to manage ownership of the SCSI port across bus reselection.
SCDRP\$L_UCB	SCSI device's UCB address. The SCSI class driver initializes this field to indicate that the SCDRP is active.
SCDRP\$W_FUNC	I/O function code that identifies the function to be performed for the I/O request. The SCSI class driver's start-I/O routine copies the contents of IRP\$W_FUNC to this field.
SCDRP\$W_STS	Status of I/O request. The SCSI class driver's start-I/O routine copies the contents of IRP\$W_STS to this field. Bits in the SCDRP\$W_STS field correspond to the bits in the IRP\$W_STS field that describe the type of I/O function, as follows: IRP\$V_BUFIO Buffered-I/O function IRP\$V_FUNC Read function IRP\$V_PAGIO Paging-I/O function IRP\$V_COMPLX Complex-buffered-I/O function IRP\$V_VIRTUAL Virtual-I/O function IRP\$V_CHAINED Chained-buffered-I/O function IRP\$V_SWAPIO Swapping-I/O function IRP\$V_DIAGBUF Diagnostic buffer present IRP\$V_PHYSIO Physical-I/O function IRP\$V_TERMIO Terminal I/O (for priority increment calculation) IRP\$V_MBXIO Mailbox-I/O function IRP\$V_EXTEND An extended IRP is linked to this IRP IRP\$V_FILACP File ACP I/O

(continued on next page)

Data Structures

1.13 SCSI Class Driver Request Packet (SCDRP)

Table 1–14 (Cont.) Contents of SCSI Class Driver Request Packet

Field Name	Contents
	IRP\$V_MVIRP Mount-verification I/O function
	IRP\$V_SRVIO Server-type I/O
	IRP\$V_KEY Encrypted function (encryption key address at IRP\$L_KEYDESC)
SCDRP\$L_SVAPTE	For a <i>direct-I/O</i> transfer, virtual address of the first page-table entry (PTE) of the I/O transfer buffer. This address is originally written to IRP\$L_SVAPTE by the FDT routine that locks process pages. For a <i>buffered-I/O</i> transfer, address of a buffer in system address space. This address is originally written to IRP\$L_SVAPTE by the class driver FDT routine that allocates the buffer. The class driver's start-I/O routine copies the address from the IRP to this field.
SCDRP\$W_BOFF	For a <i>direct-I/O</i> transfer, byte offset into the first page of the buffer; for a <i>buffered-I/O</i> transfer, number of bytes to be charged to the process requesting the transfer. FDT routines calculate this value and write it to IRP\$W_BOFF. The class driver's start-I/O routine copies the value from the IRP to this field.
SCDRP\$L_BCNT	Byte count of the I/O transfer. Class driver FDT routines calculate this value and write it to IRP\$L_BCNT. The class driver's start-I/O routine copies the value from the IRP to this field.
SCDRP\$L_MEDIA	Spare field.
SCDRP\$L_ABCNT	Accumulated count of bytes transferred. The SCSI class driver maintains this field to accomplish segmented transfers.
SCDRP\$L_SAVD_RTN	Saved return address from Level 1 JSB.
SCDRP\$L_CDT	Address of the SCSI connection descriptor table (SCDT). When the SCSI class driver's unit initialization routine invokes the SPI\$CONNECT macro, the macro returns the address of the SCDT describing the connection it established to the SCSI port. The class driver stores that address in SCDRP\$L_CDT.
SCDRP\$L_IRP	Address of I/O request block. The SCSI class driver copies the address of the IRP to this field.
SCDRP\$L_SVA_USER	System virtual address of a process buffer as mapped in system space (S0 space). The SCSI port driver initializes this field as the result of a class driver call to SPI\$MAP_BUFFER.
SCDRP\$L_CMD_BUF	Address of the port command buffer. The SCSI class driver initializes this field with the address returned from a call to SPI\$ALLOCATE_COMMAND_BUFFER.
SCDRP\$L_CMD_BUF_LEN	Length of SCSI command buffer.
SCDRP\$L_CMD_PTR	Address of the SCSI command descriptor block (its length byte) in the SCSI command buffer allocated by the SCSI port driver. The SCSI class driver initializes this field.

(continued on next page)

Data Structures

1.13 SCSI Class Driver Request Packet (SCDRP)

Table 1–14 (Cont.) Contents of SCSI Class Driver Request Packet

Field Name	Contents
SCDRP\$L_STS_PTR	Address of SCSI status byte in the port command buffer. The SCSI class driver initializes this field.
SCDRP\$L SCSI_FLAGS	SCSI flags. The SCSI class and port drivers use the following bits: SCDRP\$V_S0BUF System buffer mapped. A SCSI class driver sets this bit, before invoking SPI\$MAP_BUFFER, if the data transfer buffer is in system space (S0). SCDRP\$V_BUFFER_MAPPED Data transfer buffer mapped. A SCSI class driver sets this bit, after invoking SPI\$MAP_BUFFER, to indicate that the data transfer buffer (either a system or process space buffer) has been mapped. SCDRP\$V_DISK_SPUN_UP START UNIT command issued. The VMS SCSI disk class sets this bit.
SCDRP\$L_DATACHECK	Address of buffer for datacheck operations. A SCSI class driver maintains this field.
SCDRP\$L SCSI_STK_PTR	Stack pointer of the class driver's return address stack.
SCDRP\$L SCSI_STK	Class driver's return address stack. This stack is 32 bytes long.
SCDRP\$L_CL_RETRY	Retry count.
SCDRP\$L_DMA_TIMEOUT	Maximum number of seconds for a target to change the SCSI bus phase or complete a data transfer. Upon sending the last command byte, the port driver waits this many seconds for the target to change the bus phase lines and assert REQ (indicating a new phase). Or, if the target enters the DATA IN or DATA OUT phase, the transfer must be completed within this interval. A class driver can initialize this field to specify a per-request DMA timeout value.
SCDRP\$L_DISCON_TIMEOUT	Maximum number of seconds, from the time the initiator receives the DISCONNECT message, for a target to reselect the initiator so that it can proceed with the disconnected I/O transfer. A class driver can initialize this field to specify a per-request disconnect timeout value.
SCDRP\$W_PAD_BCNT	Pad byte count. This field contains the number of bytes required to make the size of the user buffer equal to the data length value required by a specific SCSI command. A SCSI class driver uses this field to accommodate SCSI device classes that require that the transfer length be specified in terms of a larger data unit than the count of bytes expressed in the SCDRP\$L_BCNT. If the total amount of data requested in the SCSI command does not match that specified in the SCDRP\$L_BCNT, this field must account for the difference.
SCDRP\$B_TQE*	Timer queue element, used by the port driver to time out pending disconnected I/O transfers. When this TQE expires, the timer thread times out expired pending I/O transfers.

(continued on next page)

Data Structures

1.13 SCSI Class Driver Request Packet (SCDRP)

Table 1–14 (Cont.) Contents of SCSI Class Driver Request Packet

Field Name	Contents
SCDRP\$L_TQE_DELAY*	Delay time for next TQE delay.
SCDRP\$L_SVA_DMA*	System address of the section of the port DMA buffer allocated for the data transfer.
SCDRP\$L_SVA_CMD*	System address of the segment of the port DMA buffer allocated for the port command buffer.
SCDRP\$W_MAPREG*	Page number of the first port DMA buffer page allocated for the data transfer.
SCDRP\$W_CMD_MAPREG*	Page number of the first port DMA buffer page allocated for the port command buffer.
SCDRP\$W_NUMREG*	Number of port DMA buffer pages allocated for the data transfer.
SCDRP\$W_CMD_NUMREG*	Number of port DMA buffer pages allocated for the port DMA buffer.
SCDRP\$L_SVA_SPT*	System virtual address of the system page-table entry that maps the first page of the process buffer in S0 space.
SCDRP\$L_SCSIMSGO_PTR*	SCSI output message pointer.
SCDRP\$L_SCSIMSGI_PTR*	SCSI input message pointer.
SCDRP\$B_SCSIMSGO_BUF*	SCSI output message buffer.
SCDRP\$B_SCSIMSGI_BUF*	SCSI input message buffer.
SCDRP\$L_MSGO_PENDING*	Output message pending flags. One or more of the following bits are set in this longword if the port driver is to send the corresponding message: SCDRP\$V_IDENTIFY IDENTIFY message SCDRP\$V_SYNC_OUT SYNCHRONOUS DATA TRANSFER REQUEST (out) message SCDRP\$V_BUS_DEVICE_RESET BUS DEVICE RESET message SCDRP\$V_MESSAGE_PARITY_ERROR MESSAGE PARITY ERROR message SCDRP\$V_ABORT ABORT message SCDRP\$V_NOP NO OPERATION message SCDRP\$V_MESSAGE_REJECT MESSAGE REJECT message
SCDRP\$L_MSGI_PENDING*	Input message pending flags. The only currently defined bit is SCDRP\$V_SYNC_IN, which is set when the port driver expects to receive a SYNCHRONOUS DATA TRANSFER REQUEST (in) message.
SCDRP\$B_LAST_MSGO*	Last message sent.
SCDRP\$L_DATA_PTR*	Current data pointer address.
SCDRP\$L_TRANS_CNT*	Actual number of bytes sent or received by the port driver. The port driver returns a value in this field to the class driver when it completes a SCSI data transfer.
SCDRP\$L_SAVE_DATA_CNT*	Running count of bytes (in two's-complement form) to be transferred. The port driver maintains this count.

(continued on next page)

Data Structures

1.13 SCSI Class Driver Request Packet (SCDRP)

Table 1–14 (Cont.) Contents of SCSI Class Driver Request Packet

Field Name	Contents
SCDRP\$L_SAVE_DATA_PTR*	Pointer to current port DMA buffer segment. The SCSI port driver maintains this pointer.
SCDRP\$L_SDP_DATA_CNT*	Storage for SDP data count.
SCDRP\$L_SDP_DATA_PTR*	Storage for SDP data pointer.
SCDRP\$L_DUETIME*	Timeout time for a disconnected I/O transfer.
SCDRP\$L_TIMEOUT_ADDR*	Address of timeout routine.
SCDRP\$W_CMD_BCNT*	Command byte count.
SCDRP\$W_BUSY_RETRY_CNT*	Count of remaining busy retries.
SCDRP\$W_ARB_RETRY_CNT*	Count of remaining arbitration retries.
SCDRP\$W_SEL_RETRY_CNT*	Count of remaining selection retries.
SCDRP\$W_CMD_RETRY_CNT*	Count of remaining command retries.
SCDRP\$W_SEL_TQE_RETRY_CNT*	Count of remaining TQE retries.
SCDRP\$L_SAVER3*	Reserved to Digital.
SCDRP\$L_SAVER6*	Reserved to Digital.
SCDRP\$L_SAVER7*	Reserved to Digital.
SCDRP\$L_SAVER3CL*	Reserved to Digital.
SCDRP\$L_SAVEPCCL*	Reserved to Digital.
SCDRP\$L_ABORTPCCL*	Reserved to Digital.
SCDRP\$L_PO_STK_PTR*	Stack pointer of the port driver's return address stack.
SCDRP\$L_PO_STK*	Port driver's return address stack. This stack is 24 bytes long.
SCDRP\$L_TAG*	Reserved to Digital.

1.14 SCSI Connection Descriptor Table (SCDT)

The SCSI connection descriptor table (SCDT) contains information specific to a connection established between a SCSI class driver and the port, such as phase records, timeout values, and error counters. The SCSI port driver creates an SCDT each time a SCSI class driver, by invoking the `SPI$CONNECT` macro, connects to a device on the SCSI bus. The class driver stores the address of the SCDT in the SCSI device's UCB.

The SCSI port driver has exclusive access to the SCDT. A SCSI class driver has no access to this structure.

The SCDT is illustrated in Figure 1–16 and described in Table 1–15.

Data Structures

1.14 SCSI Connection Descriptor Table (SCDT)

Figure 1–16 SCSI Connection Descriptor Table (SCDT)

SCDT\$L_FLINK*		0
Reserved	SCDT\$W_SIZE*	4
SCDT\$B_FLCK*	Reserved	8
SCDT\$L_FPC*		12
SCDT\$L_FR3*		16
SCDT\$L_FR4*		20
SCDT\$L_STS*		24
SCDT\$W_STATE*	SCDT\$W_SCDT_TYPE*	28
SCDT\$L_SPDT*		32
SCDT\$L_SCSI_PORT_ID*		36
SCDT\$L_SCSI_BUS_ID*		40
SCDT\$L_SCSI_LUN*		44
Reserved		48
SCDT\$L_SCDRP_ADDR*		56
SCDT\$L_BUS_PHASE*		60
SCDT\$L_OLD_PHASES*		64
~	SCDT\$W_PHASES* (44 bytes)	~ 68
SCDT\$L_PHASE_STK_PTR*		112
SCDT\$L_PHASE_END_STK_PTR*		116
SCDT\$L_EVENTS_SEEN*		120
SCDT\$L_ARB_FAIL_CNT*		124
SCDT\$L_SEL_FAIL_CNT*		128
SCDT\$L_PARERR_CNT*		132

(continued on next page)

Data Structures

1.14 SCSI Connection Descriptor Table (SCDT)

Figure 1–16 (Cont.) SCSI Connection Descriptor Table (SCDT)

SCDT\$L_MISPHS_CNT*		136
SCDT\$L_BADPHS_CNT*		140
SCDT\$L_RETRY_CNT*		144
SCDT\$L_RST_CNT*		148
SCDT\$L_CTLERR_CNT*		152
SCDT\$L_BUSERR_CNT*		156
SCDT\$L_CMDSENT*		160
SCDT\$L_MSGSENT*		164
SCDT\$L_BYTSENT*		168
SCDT\$L_CON_FLAGS*		172
SCDT\$L_SYNCHRONOUS*		176
SCDT\$W_TRANSFER_PERIOD*	SCDT\$W_REQACK_OFFSET*	180
SCDT\$W_ARB_RETRY_CNT*	SCDT\$W_BUSY_RETRY_CNT*	184
SCDT\$W_CMD_RETRY_CNT*	SCDT\$W_SEL_RETRY_CNT*	188
SCDT\$L_DMA_TIMEOUT*		192
SCDT\$L_DISCON_TIMEOUT*		196
SCDT\$L_SEL_CALLBACK*		200
Reserved (40 bytes)		~204

*A read-only field

1.14 SCSI Connection Descriptor Table (SCDT)

Table 1–15 Contents of SCSI Connection Descriptor Table

Field Name	Contents						
SCDT\$L_FLINK*	SCDT forward link. This field points to the next SCDT in the port's SCDT list (at SPDT\$L_SCDT_VECTOR). The SCSI port driver initializes this field when it creates the SCDT in response to an SPI\$CONNECT call.						
SCDT\$W_SIZE*	Size of SCDT. The port driver, after allocating sufficient nonpaged pool for the SCDT, writes the constant SCDT\$C_LENGTH into this field.						
SCDT\$B_FLCK*	Index of the fork lock that synchronizes access to this SCDT at fork level. The SCSI port driver, when creating the SCDT, initializes this field with SPL\$C_IOLOCK8. The SCDT fork block is used during an ABORT command request on the connection.						
SCDT\$L_FPC*	Address of instruction at which the suspended port driver thread is to be resumed.						
SCDT\$L_FR3*	Value of R3 when the request is stalled during disconnection. The value in R3 is restored before a suspended driver thread is resumed.						
SCDT\$L_FR4*	Value of R4 when the request is stalled during disconnection. The value in R4 is restored before a suspended driver thread is resumed.						
SCDT\$L_STS*	Connection status. This field is a bit map, maintained by the port driver. The only currently defined bit is SCDT\$V_BSY (connection busy).						
SCDT\$W_SCDT_TYPE*	Type of SCDT.						
SCDT\$W_STATE*	SCSI connection state. The VMS SCSI port driver maintains this field, using the following constants: <table style="margin-left: 2em;"> <tr> <td>SCDT\$C_CLOSED</td> <td>Closed</td> </tr> <tr> <td>SCDT\$C_OPEN</td> <td>Open</td> </tr> <tr> <td>SCDT\$C_FAIL</td> <td>Failed</td> </tr> </table>	SCDT\$C_CLOSED	Closed	SCDT\$C_OPEN	Open	SCDT\$C_FAIL	Failed
SCDT\$C_CLOSED	Closed						
SCDT\$C_OPEN	Open						
SCDT\$C_FAIL	Failed						
SCDT\$L_SPDT*	Address of port descriptor table with which this SCDT is associated.						
SCDT\$L_SCSI_PORT_ID*	SCSI port ID of the port to which this connection is established.						
SCDT\$L_SCSI_BUS_ID*	SCSI device ID of the device unit to which this connection is established.						
SCDT\$L_SCSI_LUN*	SCSI logical unit number (LUN) of the device unit to which this connection is established.						
SCDT\$L_SCDRP_ADDR*	Address of SCDRP current on the connection.						

(continued on next page)

Data Structures

1.14 SCSI Connection Descriptor Table (SCDT)

Table 1–15 (Cont.) Contents of SCSI Connection Descriptor Table

Field Name	Contents
SCDT\$L_BUS_PHASE*	Current SCSI bus phase. The VMS SCSI port driver defines the following flags in this longword bit map:
SCDT\$V_DATAOUT	DATA OUT phase
SCDT\$V_DATAIN	DATA IN phase
SCDT\$V_CMD	COMMAND phase
SCDT\$V_STS	STATUS phase
SCDT\$V_INV1	Invalid phase 1
SCDT\$V_INV2	Invalid phase 2
SCDT\$V_MSGOUT	MESSAGE OUT phase
SCDT\$V_MSGIN	MESSAGE IN phase
SCDT\$V_ARB	ARBITRATION phase
SCDT\$V_SEL	SELECTION phase
SCDT\$V_RESEL	RESELECTION phase
SCDT\$V_DISCON	DISCONNECT message seen
SCDT\$V_TMODISCON	Disconnect operation timed out
SCDT\$V_CMD_CMPL	COMMAND COMPLETE message received
SCDT\$V_PND_RESEL	Reselection interrupt pending
SCDT\$V_FREE	BUS FREE phase
SCDT\$L_OLD_PHASES*	Bus phase tracking information.
SCDT\$W_PHASES*	Bus phase tracking information. This field is 44 bytes long.
SCDT\$L_PHASE_STK_PTR*	Address of the top of the bus phase stack. The VMS SCSI port driver uses the bus phase stack to maintain a phase histogram.
SCDT\$L_PHASE_END_STK_PTR*	Address of the bottom of the bus phase stack. The VMS SCSI port driver uses the bus phase stack to maintain a phase histogram.
SCDT\$L_EVENTS_SEEN*	Longword bit mask of bus events seen by the VMS SCSI port driver. VMS defines the following bits:
SCDT\$V_PARERR	Parity error
SCDT\$V_BSYERR	Bus lost during command
SCDT\$V_MISPHS	Missing bus phase
SCDT\$V_BADPHS	Bad phase transition
SCDT\$V_RST	Bus reset during command
SCDT\$V_CTLERR	SCSI controller error
SCDT\$V_BUSERR	SCSI bus error
SCDT\$L_ARB_FAIL_CNT*	Count of arbitration failures.
SCDT\$L_SEL_FAIL_CNT*	Count of selection failures.
SCDT\$L_PARERR_CNT*	Count of parity errors.
SCDT\$L_MISPHS_CNT*	Count of missing phases errors.

(continued on next page)

1.14 SCSI Connection Descriptor Table (SCDT)

Table 1–15 (Cont.) Contents of SCSI Connection Descriptor Table

Field Name	Contents
SCDT\$L_BADPHS_CNT*	Count of bad phase errors.
SCDT\$L_RETRY_CNT*	Count of retries.
SCDT\$L_RST_CNT*	Count of bus resets.
SCDT\$L_CTLERR_CNT*	Count of controller errors.
SCDT\$L_BUSERR_CNT*	Count of bus errors.
SCDT\$L_CMDSENT*	Number of commands sent on this connection.
SCDT\$L_MSGSENT*	Number of messages sent on this connection.
SCDT\$L_BYTSENT*	Number of bytes sent during DATA OUT phase.
SCDT\$L_CON_FLAGS*	Connection-specific flags. The VMS SCSI port driver sets or clears these flags according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro. The following bits are defined: SCDT\$V_ENA_DISCON Enable disconnect SCDT\$V_DIS_RETRY Disable command retry SCDT\$V_TARGET_MODE Enable asynchronous event notification from target
SCDT\$L_SYNCHRONOUS*	Synchronous data transfer enabled field. This longword contains 1 if synchronous data transfers are enabled for this connection; otherwise it contains a 0. The VMS SCSI port driver writes this field according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro.
SCDT\$W_REQACK_OFFSET*	For synchronous data transfers, maximum number of REQs outstanding on the connection before an ACK is transmitted. The VMS SCSI port driver writes this field according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro.
SCDT\$W_TRANSFER_PERIOD*	Number of 4-nanosecond ticks between a REQ and an ACK on this connection. The VMS SCSI port driver writes this field according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro.
SCDT\$W_BUSY_RETRY_CNT*	Remaining number of retries allowed on this connection to successfully send a command to the target device. The VMS SCSI port driver initially writes this field according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro.
SCDT\$W_ARB_RETRY_CNT*	Remaining number of retries allowed on this connection while waiting for the port to win arbitration of the bus. The VMS SCSI port driver initially writes this field according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro.
SCDT\$W_SEL_RETRY_CNT*	Select retry count. Remaining number of retries allowed on this connection while waiting for the port to be selected by the target device. The VMS SCSI port driver initially writes this field according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro.

(continued on next page)

Data Structures

1.14 SCSI Connection Descriptor Table (SCDT)

Table 1–15 (Cont.) Contents of SCSI Connection Descriptor Table

Field Name	Contents
SCDT\$W_CMD_RETRY_CNT*	Remaining number of retries allowed on this connection to successfully send a command to the target device. The VMS SCSI port driver initially writes this field according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro.
SCDT\$L_DMA_TIMEOUT*	Timeout value (in seconds) for a target to change the SCSI bus phase or complete a data transfer. The VMS SCSI port driver initially writes this field according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro.
SCDT\$L_DISCON_TIMEOUT*	Disconnect timeout. Default timeout value (in seconds) for a target to reselect the initiator to proceed with a disconnected I/O transfer. The VMS SCSI port driver initially writes this field according to information the SCSI class driver supplies to the SPI\$SET_CONNECTION_CHAR macro.
SCDT\$L_SEL_CALLBACK*	Address of class driver's asynchronous event notification callback routine.

1.15 SCSI Port Descriptor Table (SPDT)

The SCSI port descriptor table (SPDT) contains information specific to a SCSI port, such as the port driver connection database. The SPDT also includes a set of vectors, corresponding to the SPI macros invoked by SCSI class drivers, that point to service routines within the port driver. The SCSI port driver's unit initialization routine creates an SPDT for each SCSI port defined for a specific MicroVAX/VAXstation system and initializes each SPI vector.

The port driver reads and writes fields in the SPDT. The class driver reads the SPDT indirectly when it invokes an SPI macro.

The SPDT is illustrated in Figure 1–17 and described in Table 1–16.

Figure 1–17 SCSI Port Descriptor Table (SPDT)

SPDT\$L_FLINK*		0
Reserved		4
SPDT\$B_FLCK*	SPDT\$B_SCSI_INT_MSK*	8
SPDT\$L_FPC*		12
SPDT\$L_FR3*		16
SPDT\$L_FR4*		20
SPDT\$L_SCSI_PORT_ID*		24
SPDT\$L_SCSI_BUS_ID*		28

(continued on next page)

Data Structures

1.15 SCSI Port Descriptor Table (SPDT)

Figure 1-17 (Cont.) SCSI Port Descriptor Table (SPDT)

SPDT\$L_STS*	32
SPDT\$L_PORT_WQFL*	36
SPDT\$L_PORT_WQBL*	40
SPDT\$L_MAXBYTECNT*	44
Reserved	48
SPDT\$L_PORT_UCB*	56
SPDT\$L_PORT_CSR*	60
SPDT\$L_PORT_IDB*	64
SPDT\$L_DMA_BASE*	68
SPDT\$L_SPTE_BASE*	72
SPDT\$L_SPTE_SVAPTE*	76
SPDT\$L_ADP*	80
~ SPDT\$L_PORT_RING* (64 bytes) ~	~ 84 ~
SPDT\$L_PORT_RING_PTR*	148
SPDT\$L_OWNERSCDT*	152
~ SPDT\$L_SCDT_VECTOR* (256 bytes) ~	~ 156 ~
SPDT\$L_DLCK*	412
<div style="float: right; border: 1px solid black; padding: 2px;">SPDT\$B_DIPL*</div>	416
Reserved	
SPDT\$L_SEL_SCDRP*	424
SPDT\$L_ENB_SEL_SCDRP*	428
SPDT\$L_MAP_BUFFER*	432
SPDT\$L_UNMAP*	436

(continued on next page)

Data Structures

1.15 SCSI Port Descriptor Table (SPDT)

Figure 1-17 (Cont.) SCSI Port Descriptor Table (SPDT)

SPDT\$L_SEND*		440
SPDT\$L_SET_CONN_CHAR*		444
SPDT\$L_GET_CONN_CHAR*		448
SPDT\$L_RESET*		452
SPDT\$L_CONNECT*		456
SPDT\$L_DISCONNECT*		460
SPDT\$L_ALLOC_COMMAND_BUFFER*		464
SPDT\$L_DEALLOC_COMMAND_BUFFER*		468
SPDT\$L_ABORT*		472
SPDT\$L_SET_PHASE*		476
SPDT\$L_SENSE_PHASE*		480
SPDT\$L_SEND_BYTES*		484
SPDT\$L_RECEIVE_BYTES*		488
SPDT\$L_FINISH_CMD*		492
SPDT\$L_RELEASE_BUS*		496
~	Reserved (52 bytes)	~500
Reserved	BUS_HUNG_VEC*	552
~	SPDT\$B_TQE* (52 bytes)	~556
SPDT\$L_TQE_DELAY*		608
SPDT\$L_BUS_HUNG_CNT*		612
SPDT\$L_TARRST_CNT*		616
SPDT\$L_RETRY_CNT*		620
SPDT\$L_STRAY_INT_CNT*		624
SPDT\$L_UNEXP_INT_CNT*		628

(continued on next page)

Data Structures

1.15 SCSI Port Descriptor Table (SPDT)

Figure 1–17 (Cont.) SCSI Port Descriptor Table (SPDT)

SPDT\$L_NODISCON_CNT*				632
SPDT\$W_DISCON_CNT*		Reserved		636
SPDT\$L_PORT_FLAGS*				640
SPDT\$L_VERSION_CHECK*				644
Reserved (36 bytes)				648
SPDT\$B_EVENT_CNT*	SPDT\$B_MODE*	SPDT\$B_STATUS*	SPDT\$B_CUR_STAT*	684
Reserved (16 bytes)				688

*A read-only field

Table 1–16 Contents of SCSI Port Descriptor Table

Field Name	Contents
SPDT\$L_FLINK*	SPDT forward link. This field points to the next SPDT in the system SPDT list. The SCSI port driver initializes this field when it creates the SPDT.
SPDT\$W_SIZE*	Size of SPDT. The VMS SCSI port driver initializes this field to SPDT\$C_PKNLENGTH or SPDT\$C_PKSLENGTH when creating the SPDT.
SPDT\$W_SPDT_TYPE*	SPDT type. The VMS SCSI port driver initializes this field to SPDT\$C_PKN or SPDT\$C_PKS when creating the SPDT.
SPDT\$B_SCSI_INT_MSK*	Port-specific interrupt mask.
SPDT\$B_FLCK*	Index of the fork lock that synchronizes access to this SPDT at fork level. The SCSI port driver, when creating the SPDT, copies to this field the value of UCB\$B_FLCK. The SPDT fork block is used during reselection and disconnection.
SPDT\$L_FPC*	Address of instruction at which the suspended port driver thread is to be resumed.
SPDT\$L_FR3*	Value of R3 when the request is stalled during disconnection. The value in R3 is restored before a suspended driver thread is resumed.
SPDT\$L_FR4*	Value of R4 when the request is stalled during disconnection. The value in R4 is restored before a suspended driver thread is resumed.
SPDT\$L_SCSI_PORT_ID*	SCSI port ID, an alphabetic value from A to Z.
SPDT\$L_SCSI_BUS_ID*	SCSI device ID of the port, a numeric value from 0 to 7.

(continued on next page)

Data Structures

1.15 SCSI Port Descriptor Table (SPDT)

Table 1–16 (Cont.) Contents of SCSI Port Descriptor Table

Field Name	Contents
SPDT\$L_STS*	Port device status. This field is a bit map maintained by the port driver. The following bits are defined: SPDT\$V_ONLINE Online SPDT\$V_TIMEOUT Timed out SPDT\$V_ERLOGIP Error log in progress SPDT\$V_CANCEL Cancel I/O SPDT\$V_POWER Power failed while unit busy SPDT\$V_BSY Busy SPDT\$V_FAILED Port failed operation or initialization
SPDT\$L_PORT_WQFL*	Port wait queue forward link. This field points to the first SCDRP waiting for the port to be free.
SPDT\$L_PORT_WQBL*	Port wait queue backward link. This field points to the last SCDRP waiting for the port to be free.
SPDT\$L_MAXBYTECNT*	Maximum byte count for a transfer using this port.
SPDT\$L_PORT_UCB*	Address of port UCB.
SPDT\$L_PORT_CSR*	Address of the port hardware's CSR.
SPDT\$L_PORT_IDB*	Address of the port IDB.
SPDT\$L_DMA_BASE*	Base address of the port's DMA buffer.
SPDT\$L_SPTE_BASE*	System virtual address of the system page-table entry mapping the first page of the port's DMA buffer.
SPDT\$L_SPTE_SVAPTE*	System virtual address of the system page-table entry that double-maps the data transfer buffer.
SPDT\$L_ADP*	Address of the adapter control block managing port resources.
SPDT\$L_PORT_RING*	64-byte field recording the PCs of port channel request and release transactions.
SPDT\$L_PORT_RING_PTR*	Pointer to the current port channel ring buffer entry.
SPDT\$L_OWNERSCDT*	Address of the SCDT of the connection that currently owns the port.
SPDT\$L_SCDT_VECTOR*	256-byte vector, recording the SCDT addresses associated with connections active for a given SCSI device ID (0 through 7).
SPDT\$L_DLCK*	Address of device lock that—in a VMS multiprocessing environment—synchronizes access to device registers and those fields at the SPDT accessed at device IPL. The port driver initializes this field from UCB\$L_DLCK when it creates the SPDT.
SPDT\$B_DIPL*	Interrupt priority level (IPL) at which the device requests hardware interrupts. The port driver initializes this field from UCB\$L_DLCK when it creates the SPDT.
SPDT\$L_SEL_SCDRP*	SCDRP used during selection interrupt.
SPDT\$L_ENB_SEL_SCDRP*	SCDRP used to enable selection.
SPDT\$L_MAP_BUFFER*	Address of the port driver routine that executes in response to a class driver's SPI\$MAP_BUFFER macro call. The port driver initializes this field.

(continued on next page)

1.15 SCSI Port Descriptor Table (SPDT)

Table 1–16 (Cont.) Contents of SCSI Port Descriptor Table

Field Name	Contents
SPDT\$L_UNMAP*	Address of the port driver routine that executes in response to a class driver's SPI\$UNMAP_BUFFER macro call. The port driver initializes this field.
SPDT\$L_SEND*	Address of the port driver routine that executes in response to a class driver's SPI\$SEND_COMMAND macro call. The port driver initializes this field.
SPDT\$L_SET_CONN_CHAR*	Address of the port driver routine that executes in response to a class driver's SPI\$SET_CONNECTION_CHAR macro call. The port driver initializes this field.
SPDT\$L_GET_CONN_CHAR*	Address of the port driver routine that executes in response to a class driver's SPI\$GET_CONNECTION_CHAR macro call. The port driver initializes this field.
SPDT\$L_RESET*	Address of the port driver routine that executes in response to a class driver's SPI\$RESET macro call. The port driver initializes this field.
SPDT\$L_CONNECT*	Address of the port driver routine that executes in response to a class driver's SPI\$CONNECT macro call. The port driver initializes this field.
SPDT\$L_DISCONNECT*	Address of the port driver routine that executes in response to a class driver's SPI\$DISCONNECT macro call. The port driver initializes this field.
SPDT\$L_ALLOC_COMMAND_BUFFER*	Address of the port driver routine that executes in response to a class driver's SPI\$ALLOCATE_COMMAND_BUFFER macro call. The port driver initializes this field.
SPDT\$L_DEALLOC_COMMAND_BUFFER*	Address of the port driver routine that executes in response to a class driver's SPI\$DEALLOCATE_COMMAND_BUFFER macro call. The port driver initializes this field.
SPDT\$L_ABORT*	Address of the port driver routine that executes in response to a class driver's SPI\$ABORT_COMMAND macro call. The port driver initializes this field.
SPDT\$L_SET_PHASE*	Address of the port driver asynchronous event notification (AEN) routine that executes in response to a class driver's SPI\$SET_PHASE macro call. The port driver initializes this field.
SPDT\$L_SENSE_PHASE*	Address of the port driver AEN routine that executes in response to a class driver's SPI\$SENSE_PHASE macro call. The port driver initializes this field.
SPDT\$L_SEND_BYTES*	Address of the port driver AEN routine that executes in response to a class driver's SPI\$SEND_BYTES macro call. The port driver initializes this field.
SPDT\$L_RECEIVE_BYTES*	Address of the port driver AEN routine that executes in response to a class driver's SPI\$RECEIVE_BYTES macro call. The port driver initializes this field.
SPDT\$L_FINISH_CMD*	Address of the port driver AEN routine that executes in response to a class driver's SPI\$FINISH_COMMAND macro call. The port driver initializes this field.
SPDT\$L_RELEASE_BUS*	Address of the port driver routine that executes in response to a class driver's SPI\$RELEASE_BUS macro call. The port driver initializes this field.

(continued on next page)

Data Structures

1.15 SCSI Port Descriptor Table (SPDT)

Table 1–16 (Cont.) Contents of SCSI Port Descriptor Table

Field Name	Contents
SPDT\$B_BUS_HUNG_VEC*	Vector of suspected hung connections.
SPDT\$B_TQE*	Timer queue element (52 bytes long), used by the port driver to time out pending disconnected I/O transfers. When this TQE expires, the timer thread times out expired pending I/O transfers.
SPDT\$L_TQE_DELAY*	Delay time for next TQE delay.
SPDT\$L_BUS_HUNG_CNT*	Count of detected bus hangs.
SPDT\$L_TARRST_CNT*	Count of target-initiated bus resets.
SPDT\$L_RETRY_CNT*	Total of retry attempts.
SPDT\$L_STRAY_INT_CNT*	Count of interrupts occurring when channel is unowned.
SPDT\$L_UNEXP_INT_CNT*	Count of unexpected interrupts occurring when channel is owned.
SPDT\$L_NODISCON_CNT*	Count of reselections when port is not disconnected.
SPDT\$W_DISCON_CNT*	Count of outstanding disconnects.
SPDT\$L_PORT_FLAGS*	Port-specific flags. The following bits are defined: SPDT\$V_SYNCH Port supports synchronous mode data transfers. SPDT\$V_ASYNC Port supports asynchronous mode data transfers. SPDT\$V_MAPPING_REG Port supports map registers. SPDT\$V_BUF_DMA Port supports buffered DMA transfers. SPDT\$V_DIR_DMA Port supports direct DMA transfers. SPDT\$V_AEN Port supports asynchronous event notification. SPDT\$V_LUNS Port supports logical unit numbers.
SPDT\$L_VERSION_CHECK*	Value used to check driver versions.
SPDT\$B_CUR_STAT*	Copy of CUR_STAT register.
SPDT\$B_STATUS*	Copy of STATUS register.
SPDT\$B_MODE*	Copy of MODE register.
SPDT\$B_EVENT_CNT*	Count of events while servicing current interrupt.

1.16 Spin Lock Data Structure (SPL)

The spin lock data structure records all information necessary to properly grant, release, and record the ownership of a spin lock. Each static system spin lock (including the fork locks) and device lock uses an SPL to record the IPL required for spin lock acquisition, its rank, and its owner. The spin lock structure also maintains a history of spin lock use and a variety of counters used in accounting and debugging.

Static system spin locks are assembled from module LDAT and are located from a vector of longword addresses starting at SMP\$AR_SPNLKVEC. UCB\$L_DLCK contains the address of the device lock for the corresponding device unit.

Data Structures

1.16 Spin Lock Data Structure (SPL)

The fields described in the spin lock data structure are illustrated in Figure 1–18 and described in Table 1–17.

Figure 1–18 Spin Lock Data Structure (SPL)

SPL\$B_VEC_INX*	SPL\$B_RANK*	SPL\$B_IPL*	SPL\$B_SPINLOCK*	0
SPL\$W_WAIT_CPUS*		SPL\$W_OWN_CNT*		4
SPL\$B_SUBTYPE*	SPL\$B_TYPE*	SPL\$W_SIZE*		8
SPL\$L_OWN_CPU*				12
SPL\$L_OWN_PC_VEC* (32 bytes)				16
SPL\$L_WAIT_PC*				48
SPL\$Q_ACQ_COUNT*				52
SPL\$L_BUSY_WAITS*				60
SPL\$Q_SPINS*				64
SPL\$L_TIMO_INT*				72
SPL\$L_RLS_PC*				76

*A read-only field

Table 1–17 Contents of the Spin Lock Data Structure

Field	Contents
SPL\$B_SPINLOCK*	The following fields are defined within SPL\$B_SPINLOCK: SPL\$V_INTERLOCK Spin lock access interlock. When set, this bit signifies that the spin lock is owned. <7:1> Reserved to Digital.
SPL\$B_IPL*	IPL required for spin lock acquisition.

(continued on next page)

Data Structures

1.16 Spin Lock Data Structure (SPL)

Table 1–17 (Cont.) Contents of the Spin Lock Data Structure

Field	Contents
SPL\$B_RANK*	Spin lock rank. Note that the internal value of a spin lock's rank, as stored in this field, is the inverse of the spin lock's logical rank, as displayed by the System Dump Analyzer. For instance, the structure of a spin lock with a logical rank of 0 contains the value 31 in this field.
SPL\$B_VEC_INX*	Index of the next entry to be written in the spin lock PC vector index (SPL\$L_OWN_PCVEC). SPL\$B_VEC_INX is updated upon each successful acquisition or release of the spin lock.
SPL\$W_OWN_CNT*	Ownership count. This field is –1 if the spin lock is unowned, zero or positive if owned. When a processor initially acquires a spin lock, this field goes from –1 to zero. A positive ownership count signifies concurrent acquisitions by a single processor.
SPL\$W_WAIT_CPUS*	Number of processors waiting to obtain the spin lock.
SPL\$W_SIZE*	Size of spin lock data structure (SPL\$C_LENGTH).
SPL\$B_TYPE*	Type of data structure. VMS writes the value DYN\$C_SPL in this field when it creates the SPL data structure.
SPL\$B_SUBTYPE*	Spin lock subtype. This field can contain the following values: SPL\$C_SPL_SPINLOCK Static system spin lock SPL\$C_SPL_FORKLOCK Fork lock SPL\$C_SPL_DEVICELOCK Device lock (dynamic spin lock)
SPL\$L_OWN_CPU*	Physical ID of owner CPU. This field is initialized to –1. Upon a successful acquisition, VMS copies the physical ID of the acquiring processor from CPU\$L_PHY_CPUID to this field.
SPL\$L_OWN_PC_VEC*	Last eight calling PCs of acquirers and releasers of the spin lock. SPL\$B_VEC_INX serves as the index of the next vector to be written in this array.
SPL\$L_WAIT_PC*	Last busy-wait PC.
SPL\$Q_ACQ_COUNT*	Count of successful acquisitions.
SPL\$L_BUSY_WAITS*	Count of failed acquisitions.
SPL\$Q_SPINS*	Count of number of spins.
SPL\$L_TIMO_INT*	Timeout interval before a spin lock acquisition attempt fails.
SPL\$L_RLS_PC*	PC of the last unconditional release of a set of nested acquisitions of the spin lock.

1.17 Unit Control Block (UCB)

The unit control block (UCB) is a variable-length block that describes a single device unit. Each device unit on the system has its own UCB. The UCB describes or provides pointers to the device type, controller, driver, device status, and current I/O activity.

During autoconfiguration, the driver-loading procedure creates one UCB for each device unit in the system. A privileged system user can request the driver-loading procedure to create UCBs for additional devices with the SYSGEN command CONNECT. The procedure creates UCBs of the length specified in the DPT. The driver uses UCB storage located beyond the standard UCB fields for device-specific data and temporary driver storage.

Data Structures

1.17 Unit Control Block (UCB)

The driver-loading procedure initializes some static UCB fields when it creates the block. VMS and device drivers can read and modify all nonstatic fields of the UCB. The UCB fields that are present for all devices are illustrated in Figure 1-20 and described in Table 1-19. The length of the basic UCB is defined by the symbol `UCB$K_LENGTH`.

UCBs are variable in length depending on the type of device and whether the driver performs error logging for the device. VMS defines a number of UCB extensions in the data structure definition macro `$UCBDEF` and defines a terminal device extension in `$TTYUCBDEF`. Table 1-18 lists those extensions that are most often used by device drivers, indicating where each is described in this chapter. Note that use of the dual-path extension is reserved to Digital; its contents should remain zero.

Table 1-18 UCB Extensions and Sizes Defined in `$UCBDEF`

Extension	Used by	Size	Figure	Table
Base UCB	All devices	<code>UCB\$K_SIZE</code>	1-20	1-19
Error log extension	All disk and tape devices	<code>UCB\$K_ERL_LENGTH</code>	1-21	1-20
Dual-path extension	Reserved to Digital	<code>UCB\$K_DP_LENGTH</code> (<code>UCB\$K_2P_LENGTH</code>)	—	—
Local tape extension	All tape devices	<code>UCB\$K_LCL_TAPE_LENGTH</code>	1-22	1-21
Local disk extension	All disk devices	<code>UCB\$K_LCL_DISK_LENGTH</code>	1-23	1-22
Terminal extension ¹	Terminal class and port drivers	<code>UCB\$K_TT_LENGTH</code>	1-24 ²	1-23

¹The terminal UCB extension is defined by the data structure definition macro, `$TTYUCBDEF`.

²Fields marked by asterisks may be written only by the VMS terminal class driver (`TTDRIVER.EXE`); a port driver may only read these fields.

In order to use an extended UCB, a device driver must specify its length in the **ucbsize** argument to the `DPTAB` macro. For instance:

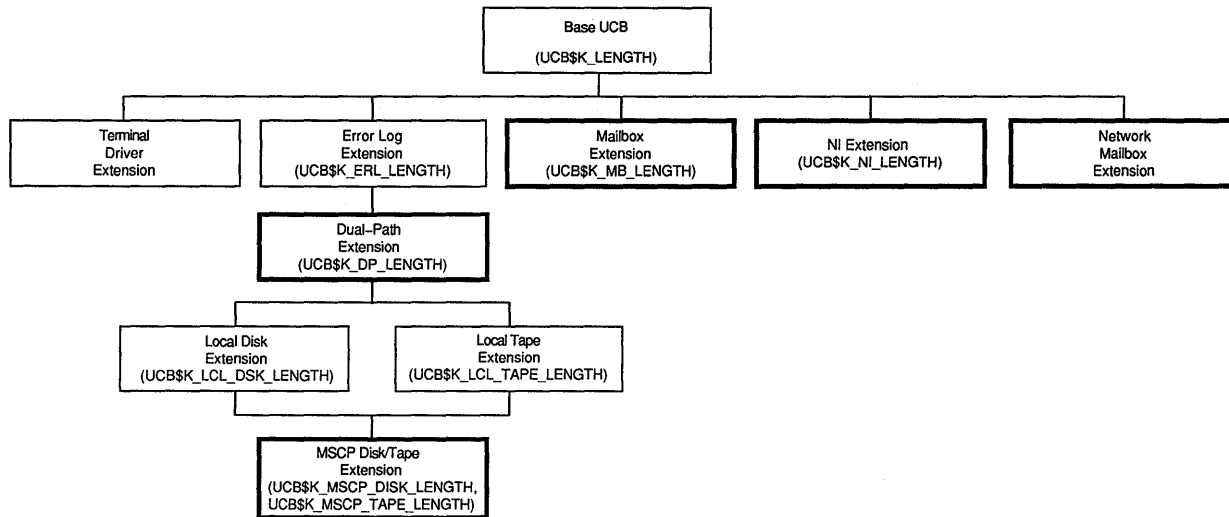
```
DPTAB  -,
      .
      .
      .
      .   UCBSIZE=UCB$K_LCL_TAPE_LENGTH, -
      .
      .
      .
```


As represented in Figure 1-19, each UCB extension used in a disk or tape driver builds upon the base UCB structure and any extension `$UCBDEF` defines earlier in the structure. (Note that UCB extensions shown in bold boxes are reserved to Digital.) For instance, if you specify a UCB size of `UCB$K_LCL_TAPE_LENGTH`, the size of the resulting UCB can accommodate the base UCB, the error log extension, the dual-path extension, and the local tape extension.

Data Structures

1.17 Unit Control Block (UCB)

Figure 1-19 Composition of Extended Unit Control Blocks



Legend:
 Bold boxes indicate UCB extensions that are reserved to Digital.

ZK-6620-GE

A device driver can further extend a UCB by using the \$DEFINI, \$DEF, \$DEFEND, and _VIELD macros. For instance:

```

$DEFINI UCB
.=UCB$K_LCL_DISK_LENGTH
$DEF   UCB$W_XX_FIELD1   .BLKW 1
$DEF   UCB$W_XX_FIELD2   .BLKW 1
$DEF   UCB$L_XX_FLAGS     .BLKL 1
      _VIELD UCB, 0, <-
      <XX_BIT1,,M>, -
      <XX_BIT2,,M>, -
      >
$DEF   UCB$K_XX_LENGTH
$DEFEND UCB
  
```

In this case, too, the driver must ensure that it specifies the length of the extended UCB in the **ucbsize** argument of the DPTAB macro:

```

DPTAB  -,
      .
      .
      .
      UCSIZE=UCB$K_XX_LENGTH, -
      .
      .
  
```

Data Structures

1.17 Unit Control Block (UCB)

Figure 1–20 Unit Control Block (UCB)

UCB\$L_FQFL*			0
UCB\$L_FQBL*			4
UCB\$B_FLCK	UCB\$B_TYPE*	UCB\$W_SIZE*	8
UCB\$L_FPC			12
UCB\$L_FR3			16
UCB\$L_FR4			20
UCB\$W_INIQUO*		UCB\$W_BUFQUO*	24
UCB\$L_ORB*			28
UCB\$L_LOCKID*			32
UCB\$L_CRB*			36
UCB\$L_DLCK*			40
UCB\$L_DDB*			44
UCB\$L_PID*			48
UCB\$L_LINK*			52
UCB\$L_VCB*			56
UCB\$L_DEVCHAR			60
UCB\$L_DEVCHAR2			64
UCB\$L_AFFINITY*			68
UCB\$L_XTRA			72
UCB\$W_DEVBUFSIZ	UCB\$B_DEVTYPE	UCB\$B_DEVCLASS	76
UCB\$Q_DEVDEPEND			80
UCB\$Q_DEVDEPEND2			88
UCB\$L_IOQFL*			96

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Figure 1-20 (Cont.) Unit Control Block (UCB)

UCB\$L_IOQBL*			100	
UCB\$W_CHARGE*		UCB\$W_UNIT*		104
UCB\$L_IRP			108	
UCB\$B_AMOD*	UCB\$B_DIPL	UCB\$W_REFC*		112
UCB\$L_AMB*			116	
UCB\$L_STS			120	
UCB\$W_QLEN*		UCB\$W_DEVSTS		124
UCB\$L_DUETIM*			128	
UCB\$L_OPCNT*			132	
UCB\$L_SVPN*			136	
UCB\$L_SVAPTE*			140	
UCB\$W_BCNT		UCB\$W_BOFF		144
UCB\$W_ERRCNT		UCB\$B_ERTMAX	UCB\$B_ERTCNT	148
UCB\$L_PDT*			152	
UCB\$L_DDT*			156	
UCB\$L_MEDIA_ID*			160	

*A read-only field

Table 1-19 Contents of Unit Control Block

Field Name	Contents
UCB\$L_FQFL*	Fork queue forward link. The link points to the next entry in the fork queue. EXE\$IOFORK and VMS resource management routines write this field. The queue contains addresses of UCBs that contain driver fork process context of drivers waiting to continue I/O processing.
UCB\$L_FQBL*	Fork queue backward link. The link points to the previous entry in the fork queue. EXE\$IOFORK and VMS resource management routines write this field.

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–19 (Cont.) Contents of Unit Control Block

Field Name	Contents
UCB\$W_SIZE*	Size of UCB. The DPT of every driver must specify a value for this field. The driver-loading procedure uses the value to allocate space for a UCB and stores the value in each UCB created. Extra space beyond the standard bytes in a UCB (UCB\$K_LENGTH) is for device-specific data and temporary storage.
UCB\$B_TYPE*	Type of data structure. The driver-loading procedure writes the constant DYN\$C_UCB into this field when the procedure creates the UCB.
UCB\$B_FLCK	<p>Index of the fork lock that synchronizes access to this UCB at fork level. The DPT of every driver must specify a value for this field. The driver-loading procedure writes the value in the UCB when the procedure creates the UCB. All devices that are attached to a single I/O adapter and actively compete for shared adapter resources and/or a controller data channel must specify the same value for this field.</p> <p>When VMS creates a driver fork process to service an I/O request for a device, the fork process gains control at the IPL associated with the fork lock, holding the fork lock itself in a VMS multiprocessing environment. When the driver creates a fork process after an interrupt, VMS inserts the fork block into a processor-specific fork queue based on this fork IPL. A VMS fork dispatcher, executing at fork IPL, obtains the fork lock (if necessary), dequeues the fork block, and restores control to the suspended driver fork process.</p> <p>This field is also known as UCB\$B_FIPL. Drivers designed to execute exclusively in a VMS uniprocessing environment store the fork IPL associated with the UCB in this field.</p>
UCB\$L_FPC	<p>Fork process driver PC address. When a VMS routine saves driver fork context in order to suspend driver execution, the routine stores the address of the next driver instruction to be executed in this field. A VMS routine that reactivates a suspended driver transfers control to the saved PC address.</p> <p>VMS routines that suspend driver processing include EXE\$IOFORK, IOC\$REQxCHANY, IOC\$REQMAPREG, IOC\$REQALTMAP, IOC\$REQDATAP, and IOC\$WFIKPCH. Routines that reactivate suspended drivers include IOC\$RELCHAN, IOC\$RELMAPREG, IOC\$RELALTMAP, IOC\$RELDATAP, EXE\$FORKDSPTH, and driver interrupt service routines.</p> <p>When a driver interrupt service routine determines that a device is expecting an interrupt, the routine restores control to the saved PC address in the device's UCB.</p>
UCB\$L_FR3	Value of R3 at the time that a VMS routine suspends a driver fork process. The value of R3 is restored just before a suspended driver regains control.
UCB\$L_FR4	Value of R4 at the time that a VMS routine suspends a driver fork process. The value of R4 is restored just before a suspended driver regains control.
UCB\$W_BUFQUO*	Buffered-I/O quota if the UCB represents a mailbox.
UCB\$W_INIQUO*	Initial buffered-I/O quota if the UCB represents a mailbox.
UCB\$L_ORB*	Address of ORB associated with the UCB. SYSGEN places the address in this field when you use SYSGEN's CONNECT command.
UCB\$L_LOCKID*	Lock management lock ID of device allocation lock. A lock management lock is used for device allocation so that device allocation functions properly for cluster-accessible devices in a VAXcluster (DEV\$V_CLU set within UCB\$L_DEVCHAR2).

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–19 (Cont.) Contents of Unit Control Block

Field Name	Contents																						
UCB\$_CRB*	Address of primary CRB associated with the device. The driver-loading procedure writes this field after it creates the associated CRB. Driver fork processes read this field to gain access to device registers. VMS routines use UCB\$_CRB to locate interrupt-dispatching code and the addresses of driver unit and controller initialization routines.																						
UCB\$_DLCK*	Address of device lock that—in a VMS multiprocessing environment—synchronizes access to device registers and those fields in the UCB accessed at device IPL. The driver-loading routine copies the address of the device lock in the CRB (CRB\$_DLCK) to this field as it creates a UCB for each device on a controller.																						
UCB\$_DDB*	Address of DDB associated with device. The driver-loading procedure writes this field when the procedure creates the associated UCB. VMS routines generally read the DDB field in order to locate device driver entry points, the address of a driver FDT, or the ACP associated with a given device.																						
UCB\$_PID*	Process identification number of the process that has allocated the device. Written by the \$ALLOC system service.																						
UCB\$_LINK*	Address of next UCB in the chain of UCBs attached to a single controller and associated with a DDB. The driver-loading procedure writes this field when the procedure adds the next UCB. Any VMS routine that examines the status of all devices on the system reads this field. Such routines include EXE\$TIMEOUT, IOC\$SEARCHDEV, and power failure recovery routines.																						
UCB\$_VCB*	Address of volume control block (VCB) that describes the volume mounted on the device. This field is written by the device's ACP and read by EXE\$QIOACPPKT, ACPs, and the XQP.																						
UCB\$_DEVCHAR	<p>First longword of device characteristics bits. The DPT of every driver should specify symbolic constant values (defined by the \$DEVDEF macro in SYS\$LIBRARY:STARLET.MLB) for this field. The driver-loading procedure writes the field when the procedure creates the UCB. The \$QIO system service reads the field to determine whether a device is spooled, file structured, shared, has a volume mounted, and so on.</p> <p>The system defines the following device characteristics:</p> <table><tbody><tr><td>DEV\$_REC</td><td>Record-oriented device</td></tr><tr><td>DEV\$_CCL</td><td>Carriage control device</td></tr><tr><td>DEV\$_TRM</td><td>Terminal device</td></tr><tr><td>DEV\$_DIR</td><td>Directory-structured device</td></tr><tr><td>DEV\$_SDI</td><td>Single directory-structured device</td></tr><tr><td>DEV\$_SQD</td><td>Sequential block-oriented device (magnetic tape, for example)</td></tr><tr><td>DEV\$_SPL</td><td>Device spooled</td></tr><tr><td>DEV\$_OPR</td><td>Operator device</td></tr><tr><td>DEV\$_RCT</td><td>Device contains RCT</td></tr><tr><td>DEV\$_NET</td><td>Network device</td></tr><tr><td>DEV\$_FOD</td><td>File-oriented device (disk and magnetic tape, for example)</td></tr></tbody></table>	DEV\$_REC	Record-oriented device	DEV\$_CCL	Carriage control device	DEV\$_TRM	Terminal device	DEV\$_DIR	Directory-structured device	DEV\$_SDI	Single directory-structured device	DEV\$_SQD	Sequential block-oriented device (magnetic tape, for example)	DEV\$_SPL	Device spooled	DEV\$_OPR	Operator device	DEV\$_RCT	Device contains RCT	DEV\$_NET	Network device	DEV\$_FOD	File-oriented device (disk and magnetic tape, for example)
DEV\$_REC	Record-oriented device																						
DEV\$_CCL	Carriage control device																						
DEV\$_TRM	Terminal device																						
DEV\$_DIR	Directory-structured device																						
DEV\$_SDI	Single directory-structured device																						
DEV\$_SQD	Sequential block-oriented device (magnetic tape, for example)																						
DEV\$_SPL	Device spooled																						
DEV\$_OPR	Operator device																						
DEV\$_RCT	Device contains RCT																						
DEV\$_NET	Network device																						
DEV\$_FOD	File-oriented device (disk and magnetic tape, for example)																						

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–19 (Cont.) Contents of Unit Control Block

Field Name	Contents
	DEV\$_DUA Dual-ported device
	DEV\$_SHR Shareable device (used by more than one program simultaneously)
	DEV\$_GEN Generic device
	DEV\$_AVL Device available for use
	DEV\$_MNT Device mounted
	DEV\$_MBX Mailbox device
	DEV\$_DMT Device marked for dismount
	DEV\$_ELG Error logging enabled
	DEV\$_ALL Device allocated
	DEV\$_FOR Device mounted as foreign (not file structured)
	DEV\$_SWL Device software write-locked
	DEV\$_IDV Device capable of providing input
	DEV\$_ODV Device capable of providing output
	DEV\$_RND Device allowing random access
	DEV\$_RTM Real-time device
	DEV\$_RCK Read-checking enabled
	DEV\$_WCK Write-checking enabled
UCB\$_DEVCHAR2	<p>Second longword of device characteristics. The DPT of every driver should specify symbolic constant values (defined by the \$DEVDEF macro in SYS\$LIBRARY:STARLET.MLB) for this field. The driver-loading procedure writes the field when the procedure creates the UCB.</p> <p>The system defines the following device characteristics:</p> <p>DEV\$_CLU Device available clusterwide</p> <p>DEV\$_DET Detached terminal</p> <p>DEV\$_RTT Remote-terminal UCB extension</p> <p>DEV\$_CDP Dual-pathed device with two UCBs</p> <p>DEV\$_2P Two paths known to device</p> <p>DEV\$_MSCP Disk or tape accessed using MSCP</p> <p>DEV\$_SSM Shadow set member</p> <p>DEV\$_SRV Served by MSCP server</p> <p>DEV\$_RED Redirected terminal</p> <p>DEV\$_NNM Device name has a prefix of the format "node\$"</p> <p>DEV\$_WBC Device supports write-back caching</p> <p>DEV\$_WTC Device supports write-through caching</p> <p>DEV\$_HOC Device supports host caching</p>
UCB\$_AFFINITY*	<p>Bit mask of the CPU-IDs of processors in a VMS multiprocessing system that have physical connectivity to the device. Such processors can thereby access the device's registers and initiate I/O operations on the device.</p>

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–19 (Cont.) Contents of Unit Control Block

Field Name	Contents																						
UCB\$L_XTRA	SMP alternate STARTIO wait.																						
UCB\$B_DEVCLASS	<p>Device class. The DPT of every driver should specify a symbolic constant (defined by the \$DCDEF macro) for this field. The driver-loading procedure writes this field when it creates the UCB.</p> <p>Drivers with set mode and device characteristics functions can rewrite the value in this field with data supplied in the characteristics buffer, the address of which is passed in the I/O request.</p> <p>VMS defines the following device classes:</p> <table><tr><td>DC\$_DISK</td><td>Disk</td></tr><tr><td>DC\$_TAPE</td><td>Tape</td></tr><tr><td>DC\$_SCOM</td><td>Synchronous communications</td></tr><tr><td>DC\$_CARD</td><td>Card reader</td></tr><tr><td>DC\$_TERM</td><td>Terminal</td></tr><tr><td>DC\$_LP</td><td>Line printer</td></tr><tr><td>DC\$_WORKSTATION</td><td>Workstation</td></tr><tr><td>DC\$_REALTIME</td><td>Real time</td></tr><tr><td>DC\$_BUS</td><td>Bus</td></tr><tr><td>DC\$_MAILBOX</td><td>Mailbox</td></tr><tr><td>DC\$_MISC</td><td>Miscellaneous</td></tr></table> <p>Note that the definition of a device as a real-time device (DC\$_REALTIME) is somewhat subjective; it implies no special treatment by VMS.</p>	DC\$_DISK	Disk	DC\$_TAPE	Tape	DC\$_SCOM	Synchronous communications	DC\$_CARD	Card reader	DC\$_TERM	Terminal	DC\$_LP	Line printer	DC\$_WORKSTATION	Workstation	DC\$_REALTIME	Real time	DC\$_BUS	Bus	DC\$_MAILBOX	Mailbox	DC\$_MISC	Miscellaneous
DC\$_DISK	Disk																						
DC\$_TAPE	Tape																						
DC\$_SCOM	Synchronous communications																						
DC\$_CARD	Card reader																						
DC\$_TERM	Terminal																						
DC\$_LP	Line printer																						
DC\$_WORKSTATION	Workstation																						
DC\$_REALTIME	Real time																						
DC\$_BUS	Bus																						
DC\$_MAILBOX	Mailbox																						
DC\$_MISC	Miscellaneous																						
UCB\$B_DEVTYPE	<p>Device type. The DPT of every driver should specify a symbolic constant (defined by the \$DCDEF macro) for this field. The driver-loading procedure writes the field when it creates the UCB.</p> <p>Drivers for devices with set mode and set characteristics functions can rewrite the value in this field with data supplied in the characteristics buffer, the address of which is passed in the I/O request.</p>																						
UCB\$W_DEVBUFSIZ	<p>Default buffer size. The DPT can specify a value for this field if relevant. The driver-loading procedure writes the field when it creates the UCB.</p> <p>Drivers for devices with set mode and set characteristics functions can rewrite the value in this field with data supplied in the characteristics buffer, the address of which is passed in the I/O request. This field is used by RMS for record I/O on nonfile devices.</p>																						
UCB\$Q_DEVDEPEND	<p>Device-descriptive data interpreted by the device driver itself. The DPT can specify a value for this field. The driver-loading procedure writes this field when it creates the UCB.</p> <p>Drivers for devices with set mode and set characteristics functions can rewrite the value in this field with data supplied in the characteristics buffer, the address of which is passed in the I/O request.</p>																						
UCB\$Q_DEVDEPN2	Second longword for device-dependent status. This field is an extension of UCB\$Q_DEVDEPEND.																						

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–19 (Cont.) Contents of Unit Control Block

Field Name	Contents				
UCB\$L_IOQFL*	<p>Pending-I/O queue listhead forward link. The queue contains the addresses of IRPs waiting for processing on a device. EXE\$INSERTIRP inserts IRPs into the pending-I/O queue when a device is busy. IOC\$REQCOM dequeues IRPs when the device is idle.</p> <p>The queue is a priority queue that has the highest priority IRPs at the front of the queue. Priority is determined by the base priority of the requesting process. IRPs with the same priority are processed first-in/first-out.</p>				
UCB\$L_IOQBL*	<p>Pending-I/O queue listhead backward link. EXE\$INSERTIRP and IOC\$REQCOM modify the pending-I/O queue.</p>				
UCB\$W_UNIT*	<p>Number of the physical device unit; stored as a binary value. The driver-loading procedure writes a value into this field when it creates the UCB. Drivers for multiunit controllers read this field during unit initialization to identify a unit to the controller.</p>				
UCB\$W_CHARGE*	<p>Mailbox byte count quota charge, if the device is a mailbox.</p>				
UCB\$L_IRP	<p>Address of IRP currently being processed on the device unit by the driver fork process. IOC\$INITIATE writes the address of an IRP into this field before the routine creates a driver fork process to handle an I/O request. From this field, a driver fork process obtains the address of the IRP being processed.</p> <p>The value contained in this field is not valid if the UCB\$V_BSY bit in UCB\$L_STS is clear.</p>				
UCB\$W_REFC*	<p>Reference count of processes that currently have process I/O channels assigned to the device. The \$ASSIGN and \$ALLOC system services increment this field. The \$DASSGN and \$DALLOC system services decrement this field.</p>				
UCB\$B_DIPL	<p>Interrupt priority level (IPL) at which the device requests hardware interrupts. The DPT of every driver must specify a value for this field. The driver-loading procedure writes this field when the procedure creates the UCB. When the driver-loading procedure subsequently creates the device lock's spin lock structure (SPL), it moves the contents of this field into SPL\$B_IPL.</p> <p>In a VMS uniprocessing environment, device drivers raise IPL to device IPL before reading or writing device registers or accessing other fields in the UCB synchronized at device IPL. In a VMS multiprocessing environment, drivers obtain the device lock at UCB\$L_DLCK, thereby also raising IPL to device IPL in the process.</p>				
UCB\$B_AMOD*	<p>Access mode at which allocation occurred, if the device is allocated. Written by the \$ALLOC and \$DALLOC system services.</p>				
UCB\$L_AMB*	<p>Associated mailbox UCB pointer. A spooled device uses this field for the address of its associated device. Devices that are nonshareable and not file oriented can use this field for the address of an associated mailbox.</p>				
UCB\$L_STS	<p>Device unit status (formerly UCB\$W_STS). Written by drivers, IOC\$REQCOM, IOC\$CANCELIO, IOC\$INITIATE, IOC\$WFIKPCH, IOC\$WFIRLCH, EXE\$INSIOQ, and EXE\$TIMEOUT. This field is read by drivers, the \$QIO system service routines, IOC\$REQCOM, IOC\$INITIATE, and EXE\$TIMEOUT.</p> <p>This longword includes the following bits:</p> <table style="width: 100%; border: none;"> <tr> <td style="width: 50%;">UCB\$V_TIM</td> <td>Timeout enabled.</td> </tr> <tr> <td>UCB\$V_INT</td> <td>Interrupts expected.</td> </tr> </table>	UCB\$V_TIM	Timeout enabled.	UCB\$V_INT	Interrupts expected.
UCB\$V_TIM	Timeout enabled.				
UCB\$V_INT	Interrupts expected.				

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–19 (Cont.) Contents of Unit Control Block

Field Name	Contents
UCB\$V_ERLOGIP	Error log in progress.
UCB\$V_CANCEL	Cancel I/O on unit.
UCB\$V_ONLINE	Device is on line.
UCB\$V_POWER	Power has failed while unit was busy.
UCB\$V_TIMEOUT	Unit is timed out.
UCB\$V_INTTYPE	Receiver interrupt.
UCB\$V_BSY	Unit is busy.
UCB\$V_MOUNTING	Device is being mounted.
UCB\$V_DEADMO	Deallocate device at dismount.
UCB\$V_VALID	Volume appears valid to software.
UCB\$V_UNLOAD	Unload volume at dismount.
UCB\$V_TEMPLATE	Template UCB from which other UCBs for this device are made. The \$ASSIGN system service checks this bit in the requested UCB and, if the bit is set, creates a UCB from the template. The new UCB is assigned instead.
UCB\$V_MNTVERIP	Mount verification in progress.
UCB\$V_WRONGVOL	Volume name does not match name in the VCB.
UCB\$V_DELETEUCB	Delete this UCB when the value in UCB\$W_REFC becomes zero.
UCB\$V_LCL_VALID	The volume on this device is valid on the local node.
UCB\$V_SUPMMSG	Suppress mount-verification messages if they indicate success.
UCB\$V_MNTVERPND	Mount verification is pending on the device and the device is busy.
UCB\$V_DISMOUNT	Dismount in progress.
UCB\$V_CLUTRAN	VAXcluster state transition in progress.
UCB\$V_WRTLOCKMV	Write-locked mount verification in progress.
UCB\$V_SVPN_END	Last byte used from page is mapped by a system virtual page number.
UCB\$W_DEVSTS	Device-dependent status. Read and written by device drivers. The system defines the following status bits:
UCB\$V_JOB	Job controller has been notified.
UCB\$V_TEMPL_BSY	Template UCB is busy.
UCB\$V_PRMMBX	Device is a permanent mailbox.
UCB\$V_DELMBX	Mailbox is marked for deletion.
UCB\$V_SHMMBS	Device is shared-memory mailbox.
	Disk drivers use bits in UCB\$W_DEVSTS as follows:
UCB\$V_ECC	ECC correction made.

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–19 (Cont.) Contents of Unit Control Block

Field Name	Contents
	UCB\$V_DIAGBUF Diagnostic buffer is specified.
	UCB\$V_NOCNVRT No logical block number to media address conversion.
	UCB\$V_DX_WRITE Console floppy write operation.
	UCB\$V_DATACACHE Data blocks are being cached.
UCB\$W_QLEN*	Length of pending-I/O queue (pointed to by UCB\$L_IOQFL).
UCB\$L_DUETIM*	Due time for I/O completion. Stored as the low-order 32-bit absolute time (time in seconds since the operating system was booted) at which the device will time out. IOC\$WFIKPC and IOC\$WFIRLCH write this value when they suspend a driver to wait for an interrupt or timeout. EXE\$TIMEOUT examines this field in each UCB in the I/O database once per second. If the timeout has occurred and timeouts are enabled for the device, EXE\$TIMEOUT calls the device driver timeout handler.
UCB\$L_OPCNT*	Count of operations completed on device unit since last bootstrap of VMS system. IOC\$REQCOM writes this field every time the routine inserts an IRP into the I/O postprocessing queue.
UCB\$L_SVFN*	Index to the virtual address of the system PTE that the driver loading procedure has permanently allocated to the device. The system virtual address of the page described by this index can be calculated by the following formula: $(\text{index} * 200_{16}) + 80000000_{16}$ If a DPT specifies DPT\$M_SVP in the flags argument to the DPTAB macro, the driver-loading procedure allocates a page of nonpaged system memory to the device. The procedure writes the system PTE's index into UCB\$L_SVFN when the procedure creates the UCB.
	Disk drivers use this field for ECC error correction.
UCB\$L_SVAPTE	For a <i>direct-I/O</i> transfer, the virtual address of the system PTE for the first page to be used in the transfer; for a <i>buffered-I/O</i> transfer, the virtual address of the system buffer used in the transfer. IOC\$INITIATE writes this field from IRP\$L_SVAPTE before calling a driver start-I/O routine. Drivers read this value to compute the starting address of a transfer.
UCB\$W_BOFF	For a <i>direct-I/O</i> transfer, the byte offset in the first page of the transfer buffer; for a <i>buffered-I/O</i> transfer, the number of bytes charged to the process for the transfer. IOC\$INITIATE copies this field from the IRP. Drivers read the field in calculating the starting address of a DMA transfer. If only part of a DMA transfer succeeds, the driver adjusts the value in this field to be the byte offset in the first page of the data that was not transferred.
UCB\$W_BCNT	Count of bytes in the I/O transfer. IOC\$INITIATE copies this field from the IRP. Drivers read this field to determine how many bytes to transfer in an I/O operation.
UCB\$B_ERTCNT	Error retry count of the current I/O transfer. The driver sets this field to the maximum retry count each time it begins I/O processing. Before each retry, the driver decreases the value in this field. During error logging, IOC\$REQCOM copies the value into the error message buffer.

(continued on next page)

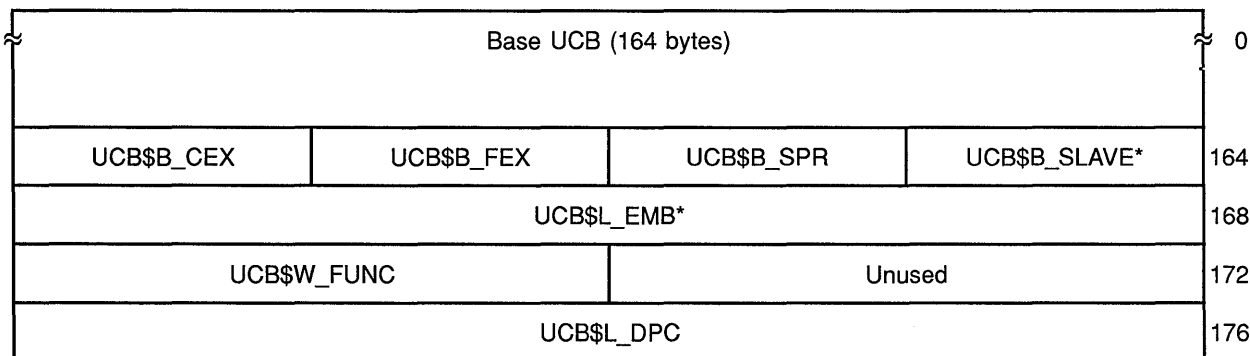
Data Structures

1.17 Unit Control Block (UCB)

Table 1–19 (Cont.) Contents of Unit Control Block

Field Name	Contents
UCB\$B_ERTMAX	Maximum error retry count allowed for single I/O transfer. The DPT of some drivers specifies a value for this field. The driver-loading procedure writes the field when the procedure creates the UCB. During error logging, IOC\$REQCOM copies the value into the error message buffer.
UCB\$W_ERRCNT	Number of errors that have occurred on the device since VMS booted. The driver-loading procedure initializes the field to 0 when the procedure creates the UCB. ERL\$DEVICERR and ERL\$DEVICTMO increment the value in the field and copy the value into an error message buffer. The DCL command SHOW DEVICE displays in its error count column the value contained in this field.
UCB\$L_PDT*	Address of port descriptor table (PDT). This field is reserved for VMS SCS port drivers.
UCB\$L_DDT*	Address of DDT for unit. The driver load procedure writes the contents of DDB\$L_DDT for the device controller to this field when it creates the UCB.
UCB\$L_MEDIA_ID*	Bit-encoded media name and type, used by MSCP devices.

Figure 1–21 UCB Error-Log Extension



*A read-only field

Table 1–20 UCB Error-Log Extension

Field Name	Contents
UCB\$B_SLAVE*	Unit number of slave controller.
UCB\$B_SPR	Spare byte. This field is reserved for driver use. MASSBUS adapter drivers use this field to store a fixed offset to the MASSBUS adapter registers for the unit.
UCB\$B_FEX	Device-specific field. This field is reserved for driver use. Certain VMS disk drivers (such as DLDRIVER in one of the appendixes to the <i>VMS Device Support Manual</i>) use this field to store an index in a hardware function dispatch table.

(continued on next page)

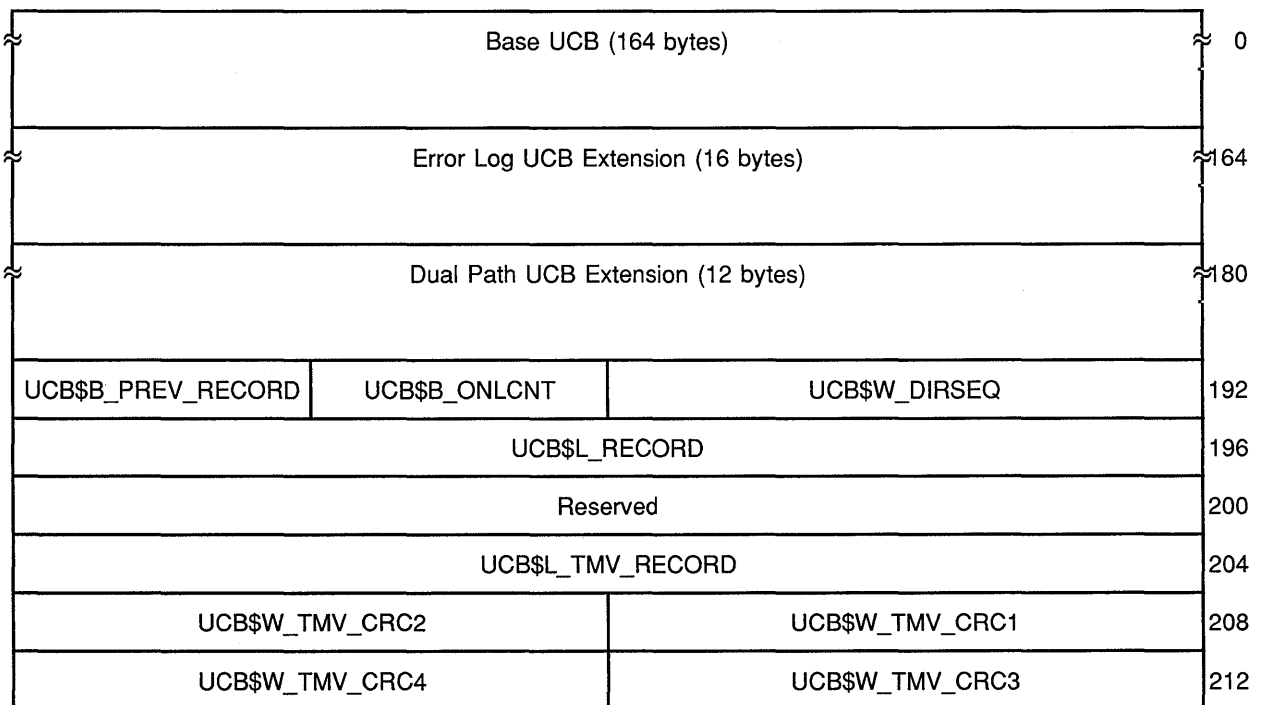
Data Structures

1.17 Unit Control Block (UCB)

Table 1–20 (Cont.) UCB Error-Log Extension

Field Name	Contents
UCB\$B_CEX	Device-specific field. This field is reserved for driver use. Certain VMS disk drivers (such as DLDRIVER in one of the appendixes to the <i>VMS Device Support Manual</i>) use this field to store an index into a software function case table.
UCB\$L_EMB*	Address of error message buffer. If error logging is enabled and a device/controller error or timeout occurs, the driver calls ERL\$DEVICERR or ERL\$DEVICTMO to allocate an error message buffer and copy the buffer address into this field. IOC\$REQCOM writes final device status, error counters, and I/O request status into the buffer specified by this field.
UCB\$W_FUNC	I/O function modifiers. This field is read and written by drivers that log errors.
UCB\$L_DPC	Device-specific field. This field is reserved for driver use. Certain VMS disk drivers (such as DLDRIVER in one of the appendixes to the <i>VMS Device Support Manual</i>) use this field to store the driver's return PC across a dispatch to a hardware function routine.

Figure 1–22 UCB Local Tape Extension



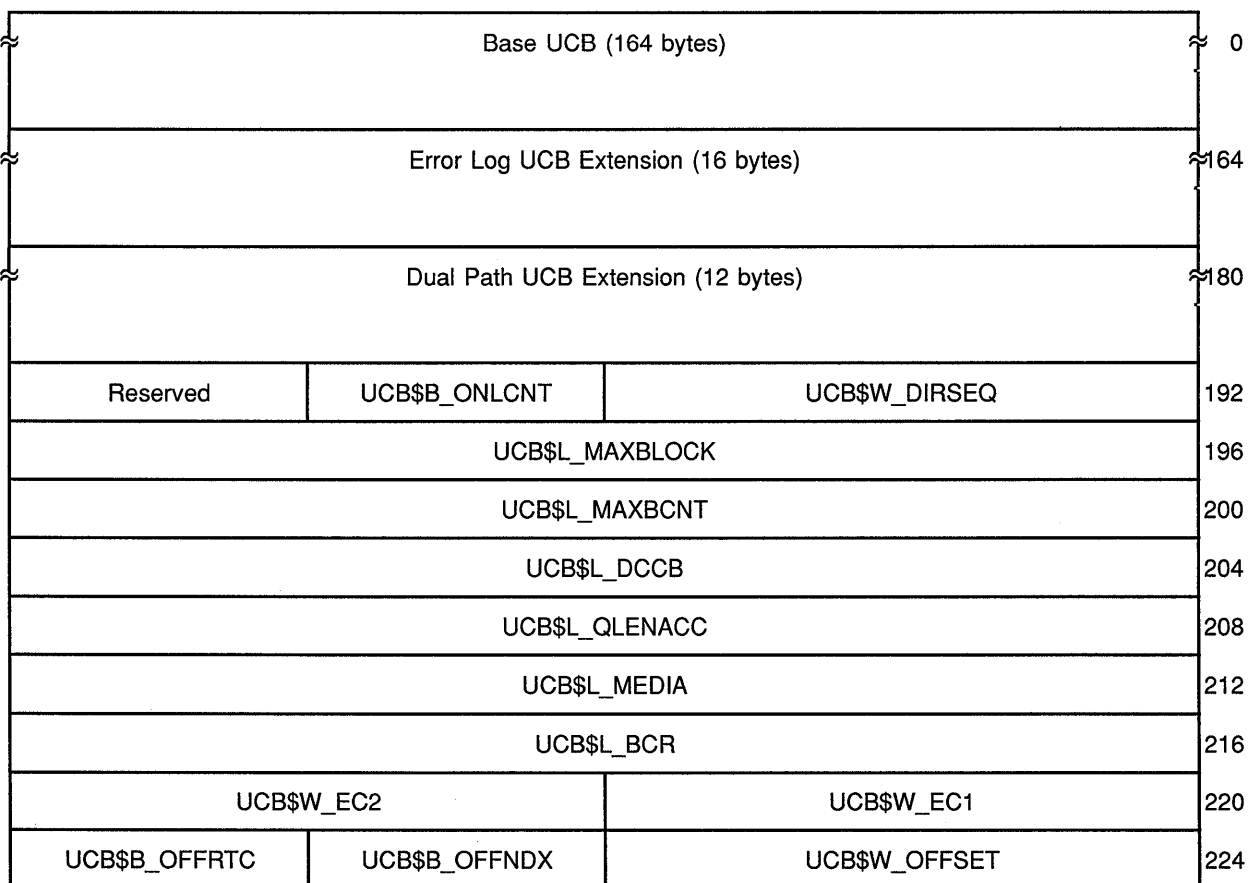
Data Structures

1.17 Unit Control Block (UCB)

Table 1–21 UCB Local Tape Extension

Field Name	Contents
UCB\$W_DIRSEQ	Directory sequence number. If the high-order bit of this word, UCB\$V_AST_ARMED, is set, it indicates that the requesting process is blocking ASTs.
UCB\$B_ONLCNT	Number of times the device has been placed on line since VMS was last bootstrapped.
UCB\$B_PREV_RECORD	Tape position prior to the start of the last I/O operation.
UCB\$L_RECORD	Current tape position or frame counter.
UCB\$L_TMV_RECORD	Position following last guaranteed successful I/O operation.
UCB\$W_TMV_CRC1	First CRC for mount verification's media validation.
UCB\$W_TMV_CRC2	Second CRC for mount verification's media validation.
UCB\$W_TMV_CRC3	Third CRC for mount verification's media validation.
UCB\$W_TMV_CRC4	Fourth CRC for mount verification's media validation.

Figure 1–23 UCB Local Disk Extension



(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Figure 1–23 (Cont.) UCB Local Disk Extension

UCB\$L_DX_BUF			228
UCB\$L_DX_BFPNT			232
UCB\$L_DX_RXDB			236
Unused	UCB\$B_DX_SCTCNT	UCB\$W_DX_BCR	240

Table 1–22 UCB Local Disk Extension

Field Name	Contents
UCB\$W_DIRSEQ	Directory sequence number. If the high-order bit of this word, UCB\$V_AST_ARMED, is set, it indicates that the requesting process is blocking ASTs.
UCB\$B_ONLCNT	Number of times device has been placed on line since VMS was last bootstrapped.
UCB\$L_MAXBLOCK	Maximum number of logical blocks on random-access device. This field is written by a disk driver during unit initialization and power recovery.
UCB\$L_MAXBCNT	Maximum number of bytes that can be transferred. A disk driver writes this field during unit initialization and power recovery.
UCB\$L_DCCB	Pointer to cache control block.
UCB\$L_QLENACC	Queue length accumulator.
UCB\$L_MEDIA	Media address.
UCB\$L_BCR	Byte-count register. Some disk drivers use this field as an internal count of the number of bytes left to be transferred in an I/O request. The symbol UCB\$W_BCR points to the low-order word of this field.
UCB\$W_EC1	ECC position register. This field records the starting bit number of an error burst. Disk driver register dumping routines copy the contents of this field into an error message or diagnostic buffer. The VMS correction routine IOC\$APPLYECC reads the contents of this field to locate the beginning of an error burst in a disk block.
UCB\$W_EC2	ECC position register. Records the exclusive OR correction pattern. Disk driver register dumping routines copy the contents of this field into an error message or diagnostic buffer. The VMS ECC correction routine IOC\$APPLYECC reads the contents of this field to correct disk data.
UCB\$W_OFFSET	Current offset register contents.
UCB\$B_OFFNDX	Current offset table index. When a disk driver transfer ends in an error, the disk driver can retry the transfer a number of times with different offsets of the disk head from the centerline. This field is an index into a driver table of offset positions.
UCB\$B_OFFRTC	Current offset retry count. This field records the number of times to try a particular offset setting in a disk transfer retry.

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–22 (Cont.) UCB Local Disk Extension

Field Name	Contents
UCB\$L_DX_BUF	Address of sector buffer (used by floppy-disk drivers).
UCB\$L_DX_BFPNT	Pointer to current sector (used by floppy-disk drivers).
UCB\$L_DX_RXDB	Address of saved receiver-data buffer (used by floppy-disk drivers).
UCB\$W_DX_BCR	Current floppy byte count (used by floppy-disk drivers).
UCB\$B_DX_SCTCNT	Current sector byte count (used by floppy-disk drivers).

Figure 1–24 UCB Terminal Extension

Base UCB (164 bytes)	0
UCB\$L_TL_CTRLY	164
UCB\$L_TL_CTRLC	168
UCB\$L_TL_OUTBAND	172
UCB\$L_TL_BANDQUE	176
UCB\$L_TL_PHYUCB	180
UCB\$L_TL_CTLPID	184
UCB\$Q_TL_BRKTHRU	188
UCB\$L_TT_RDUE	196
UCB\$L_TT_RTIMOU	200
UCB\$L_TT_STATE1	204
UCB\$L_TT_STATE2	208
UCB\$L_TT_LOGUCB	212
UCB\$L_TT_DECHAR	216
UCB\$L_TT_DECHA1	220
UCB\$L_TT_DECHA2	224
UCB\$L_TT_DECHA3	228

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Figure 1-24 (Cont.) UCB Terminal Extension

UCB\$L_TT_WFLINK				232
UCB\$L_TT_WBLINK				236
UCB\$L_TT_WRTBUF				240
UCB\$L_TT_MULTI				244
UCB\$W_TT_SMLTLEN		UCB\$W_TT_MULTILEN		248
UCB\$L_TT_SMLT				252
UCB\$B_TT_DELFF	UCB\$B_TT_DECRF	UCB\$W_TT_DESPÉE		256
Unused			UCB\$B_TT_DEPARI	260
Reserved	UCB\$W_TT_DESIZE		UCB\$B_TT_DETYPÉ	264
UCB\$B_TT_LFFILL	UCB\$B_TT_CRFILL	UCB\$B_TT_RSPEED	UCB\$B_TT_TSPEED	268
Unused			UCB\$B_TT_PARITY	272
UCB\$L_TT_TYPAHD				276
UCB\$B_TT_LASTC	UCB\$B_TT_LINE	UCB\$W_TT_CURSOR		280
UCB\$B_TT_ESC	UCB\$B_TT_FILL	UCB\$W_TT_BSPLÉN		284
UCB\$W_TT_UNITBIT		UCB\$B_TT_INTCNT	UCB\$B_TT_ESC_O	288
UCB\$B_TT_OUTYPE	UCB\$B_TT_PREMPT	UCB\$W_TT_HOLD		292
UCB\$L_TT_GETNXT				296
UCB\$L_TT_PUTNXT				300
UCB\$L_TT_CLASS				304
UCB\$L_TT_PORT				308
UCB\$L_TT_OUTADR				312
UCB\$W_TT_PRTCTL		UCB\$W_TT_OUTLEN		316
UCB\$W_TT_DS_ST		UCB\$B_TT_DS_TX	UCB\$B_TT_DS_RCV	320
UCB\$B_TT_OLD	UCB\$B_TT_MAINT	UCB\$W_TT_DS_TIM		324
UCB\$L_TT_FBK				328
UCB\$L_TT_RDVERIFY				332

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Figure 1-24 (Cont.) UCB Terminal Extension

UCB\$L_TT_CLASS1		336
UCB\$L_TT_CLASS2		340
UCB\$L_TT_ACCPORNAM		344
UCB\$L_TP_MAP		348
Unused	UCB\$B_TP_STAT	352

Table 1-23 UCB Terminal Extension

Field Name	Contents
UCB\$L_TL_CTRLY*	Listhead of CTRL/Y AST control blocks (ACBs).
UCB\$L_TL_CTRLC*	Listhead of CTRL/C ACBs.
UCB\$L_TL_OUTBAND*	Out-of-band character mask.
UCB\$L_TL_BANDQUE*	Listhead of out-of-band ACBs.
UCB\$L_TL_PHYUCB*	Address of physical UCB.
UCB\$L_TL_CTLPID*	Process ID of controlling process (used with SPAWN).
UCB\$Q_TL_BRKTHRU*	Facility broadcast bit mask.
UCB\$L_TT_RDUE*	Absolute time at which a read timeout is due.
UCB\$L_TT_RTIMOU*	Address of read timeout routine.
UCB\$L_TT_STATE1*	First longword of terminal state information. The following fields are defined within UCB\$L_TT_STATE1:
	TTY\$V_ST_POWER Power failure
	TTY\$V_ST_CTRL Class output
	TTY\$V_ST_FILL Fill mode
	TTY\$V_ST_CURSOR Cursor
	TTY\$V_ST_SENDF Forced line feed
	TTY\$V_ST_BACKSPACE Backspace
	TTY\$V_ST_MULTI Multi-echo
	TTY\$V_ST_WRITE Write in progress
	TTY\$V_ST_EOL End of line
	TTY\$V_ST_EDITREAD Editing read in progress
	TTY\$V_ST_RDVERIFY Read verify in progress
	TTY\$V_ST_RECALL Command recall
	TTY\$V_ST_READ Read in progress

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–23 (Cont.) UCB Terminal Extension

Field Name	Contents
UCB\$L_TT_STATE2*	<p>Second longword of terminal state information.</p> <p>The following fields are defined within UCB\$L_TT_STATE2:</p> <p>TTY\$V_ST_CTRL0 Output enable</p> <p>TTY\$V_ST_DEL Delete</p> <p>TTY\$V_ST_PASALL Pass-all mode</p> <p>TTY\$V_ST_NOECHO No echo</p> <p>TTY\$V_ST_WRTALL Write-all mode</p> <p>TTY\$V_ST_PROMPT Prompt</p> <p>TTY\$V_ST_NOFLTR No control-character filtering</p> <p>TTY\$V_ST_ESC Escape sequence</p> <p>TTY\$V_ST_BADESC Bad escape sequence</p> <p>TTY\$V_ST_NL New line</p> <p>TTY\$V_ST_REFRSH Refresh</p> <p>TTY\$V_ST_ESCAPE Escape mode</p> <p>TTY\$V_ST_TYPFUL Type-ahead buffer full</p> <p>TTY\$V_ST_SKIPLF Skip line feed</p> <p>TTY\$V_ST_ESC_O Output escape</p> <p>TTY\$V_ST_WRAP Wrap enable</p> <p>TTY\$V_ST_OVFLO Overflow condition</p> <p>TTY\$V_ST_AUTOP Autobaud pending</p> <p>TTY\$V_ST_CTRLR Clock prompt and data string from read buffer</p> <p>TTY\$V_ST_SKIPCRLF Skip line feed following a carriage return</p> <p>TTY\$V_ST_EDITING Editing operation</p> <p>TTY\$V_ST_TABEXPAND Expand tab characters</p> <p>TTY\$V_ST_QUOTING Quote character</p> <p>TTY\$V_ST_OVERSTRIKE Overstrike mode</p> <p>TTY\$V_ST_TERMNORM Standard terminator mask</p> <p>TTY\$V_ST_ECHAES Alternate echo string</p> <p>TTY\$V_ST_PRE Pre-type-ahead mode</p> <p>TTY\$V_ST_NINTMULTI Noninterrupt multi-echo mode</p> <p>TTY\$V_ST_RECONNECT Reconnect operation</p> <p>TTY\$V_ST_CTSLOW Clear-to-send low</p> <p>TTY\$V_ST_TABRIGHT Check for tabs to the right of the current position</p>
UCB\$L_TT_LOGUCB*	<p>Address of logical UCB, if the redirect bit is set (DEV\$V_RED in UCB\$L_DEVCHAR2). If this UCB describes the logical UCB, the contents of UCB\$L_TT_LOGUCB are zero.</p>
UCB\$L_TT_DECHAR*	<p>First longword of default device characteristics.</p>

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–23 (Cont.) UCB Terminal Extension

Field Name	Contents
UCB\$L_TT_DECHA1*	Second longword of default device characteristics.
UCB\$L_TT_DECHA2*	Third longword of default device characteristics.
UCB\$L_TT_DECHA3*	Fourth longword of default device characteristics.
UCB\$L_TT_WFLINK*	Write queue forward link.
UCB\$L_TT_WBLINK*	Write queue backward link.
UCB\$L_TT_WRTBUF*	Current write buffer block.
UCB\$L_TT_MULTI*	Address of current multi-echo buffer.
UCB\$W_TT_MULTILEN*	Length of multi-echo string to be written.
UCB\$W_TT_SMLTLEN*	Saved length of multi-echo string.
UCB\$L_TT_SMLT*	Saved address of multi-echo buffer.
UCB\$W_TT_DESPEE*	Default speed.
UCB\$B_TT_DECRF*	Default carriage-return fill.
UCB\$B_TT_DELFF*	Default line-feed fill.
UCB\$B_TT_DEPARI*	Default parity/character size.
UCB\$B_TT_DETTYPE*	Default terminal type.
UCB\$W_TT_DESIZE*	Default line size.
UCB\$W_TT_SPEED*	Terminal line speed. This field is read and written by the class driver, and read by the port driver. It contains the following byte fields: UCB\$B_TT_TSPEED Transmit speed UCB\$B_TT_RSPEED Receive speed
UCB\$B_TT_CRFILL*	Number of fill characters to be output for carriage return.
UCB\$B_TT_LFFILL*	Number of fill characters to be output for line feed.
UCB\$B_TT_PARITY*	Parity, frame and stop bit information to be set when the PORT_SET_LINE service routine is called. This field is read and written by the class driver, and read by the port driver. It contains the following bit fields: UCB\$V_TT_XXPARITY Reserved to Digital. UCB\$V_TT_DISPARERR Reserved to Digital. UCB\$V_TT_USERFRAME Reserved to Digital. UCB\$V_TT_LEN Two bits signifying character length (not counting start, stop, and parity bits), as follows: 00 ₂ = 5 bits; 01 ₂ = 6 bits; 10 ₂ = 7 bits; and 11 ₂ = 8 bits. UCB\$V_TT_STOP Number of stop bits: clear if one stop bit; set if two stop bits. UCB\$V_TT_PARITY Parity checking. This bit is set if parity checking is enabled. UCB\$V_TT_ODD Parity type: clear if even parity; set if odd parity.
UCB\$L_TT_TYPAHD*	Address of type-ahead buffer.
UCB\$W_TT_CURSOR*	Current cursor position.

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–23 (Cont.) UCB Terminal Extension

Field Name	Contents
UCB\$B_TT_LINE*	Current line position on page.
UCB\$B_TT_LASTC*	Last formatted output character.
UCB\$W_TT_BSPLN*	Number of back spaces to output for non-ANSI terminals.
UCB\$B_TT_FILL*	Current fill character count.
UCB\$B_TT_ESC*	Current read escape syntax state.
UCB\$B_TT_ESC_O*	Current write escape syntax state.
UCB\$B_TT_INTCNT*	Number of characters in interrupt string.
UCB\$W_TT_UNITBIT*	Enable and disable modem control.
UCB\$W_TT_HOLD	Port driver's internal flags and unit holding tank. This is read and written by the port driver, and is not accessed by the class driver. It contains the following subfields: TTY\$B_TANK_CHAR Character. TTY\$V_TANK_PREMPT Send preempt character. TTY\$V_TANK_STOP Stop output. TTY\$V_TANK_HOLD Character stored in TTY\$B_TANK_CHAR. TTY\$V_TANK_BURST Burst is active. TTY\$V_TANK_DMA DMA transfer is active.
UCB\$B_TT_PREMPT	Preempt character.
UCB\$B_TT_OUTYPE*	Amount of data to be written on a callback from the class driver. When negative, this field indicates that there is a burst of data ready to be returned; when zero, it signifies that no data is to be written; and when 1, it indicates that a single character is to be written. This field is written by the class driver and read by the port driver.
UCB\$L_TT_GETNXT*	Address of the class driver's input routine. This field is read by the port driver.
UCB\$L_TT_PUTNXT*	Address of the class driver's output routine. This field is read by the port driver.
UCB\$L_TT_CLASS*	Address of the class driver's vector table. This field is initialized by the CLASS_CTRL_INIT macro. The port driver reads UCB\$L_TT_CLASS whenever it must call the class driver at an entry point other than UCB\$L_TT_GETNXT or UCB\$L_TT_PUTNXT.
UCB\$L_TT_PORT	Address of the port driver's vector table.
UCB\$L_TT_OUTADR	Address of the first character of a burst of data to be written. This field is only valid when UCB\$B_TT_OUTYPE contains –1. It is read and written by the port driver, and written by the class driver.
UCB\$W_TT_OUTLEN	Number of characters in a burst of data to be written. This field is only valid when UCB\$B_TT_OUTYPE contains –1. It is read and written by the port driver, and written by the class driver.

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–23 (Cont.) UCB Terminal Extension

Field Name	Contents
UCB\$W_TT_PRTCTL	Port driver control flags. The bits in this field indicate features that are available to the port; the class driver specifies which of these features are to be enabled. The following fields are defined within UCB\$W_TT_PRTCTL.
TTY\$V_PC_NOTIME	No timeout. If set, the terminal class driver is not to set up timers for output.
TTY\$V_PC_DMAENA	DMA enabled. If set, DMA transfers are currently enabled on this port.
TTY\$V_PC_DMAAVL	DMA supported. If set, DMA transfers are supported for this port.
TTY\$V_PC_PRMMAP	Permanent map registers. If set, the port driver is to permanently allocate UNIBUS/Q22-bus map registers.
TTY\$V_PC_MAPAVL	Map registers available. If set, the port driver has currently allocated map registers.
TTY\$V_PC_XOFFAVL	Auto XOFF supported. If set, auto XOFF is supported for this port.
TTY\$V_PC_XOFFENA	Auto XOFF enabled. If set, auto XOFF is currently enabled on this port.
TTY\$V_PC_NOCTRL	No auto line feed. If set, a line feed is not generated following a carriage return.
TTY\$V_PC_BREAK	Break. If set, the port driver should generate break character; if clear, the port should turn off the break feature.
TTY\$V_PC_PORTFDT	FDT routine. If set, the port driver contains FDT routines.
TTY\$V_PC_NOMODEM	No modem. If set, the port cannot support modem operations.
TTY\$V_PC_NODISCONNECT	No disconnect. If set, the device cannot support virtual terminal operations.
TTY\$V_PC_SMART_READ	Smart read. If set, the port contains additional read capabilities.
TTY\$V_PC_ACCPORNAM	Access port name. If set, the port supports an access port name.
TTY\$V_PC_MULTISESSION	Multisession terminal. If set, the port is part of a multisession terminal.
UCB\$B_TT_DS_RCV	Current receive modem.
UCB\$B_TT_DS_TX	Current transmit modem.
UCB\$W_TT_DS_ST*	Current modem state.
UCB\$W_TT_DS_TIM*	Current modem timeout.

(continued on next page)

Data Structures

1.17 Unit Control Block (UCB)

Table 1–23 (Cont.) UCB Terminal Extension

Field Name	Contents												
UCB\$B_TT_MAINT*	<p>Maintenance functions. This field is used as the argument to the port driver's PORT_MAINT routine. It is written by the class driver and read by the port driver.</p> <p>It contains several bits that allow the following maintenance functions:</p> <table style="width: 100%; border: none;"> <tr> <td style="width: 30%;">IO\$M_LOOP</td> <td>Set loopback mode.</td> </tr> <tr> <td>IO\$M_UNLOOP</td> <td>Reset loopback mode.</td> </tr> <tr> <td>IO\$M_AUTXOF_ENA</td> <td>Enable the use of auto XON/XOFF on this line. This is the default.</td> </tr> <tr> <td>IO\$M_AUTXOF_DIS</td> <td>Disable the use of auto XON/XOFF on this line.</td> </tr> <tr> <td>IO\$M_LINE_OFF</td> <td>Disable interrupts on this line.</td> </tr> <tr> <td>IO\$M_LINE_ON</td> <td>Reenable interrupts on this line.</td> </tr> </table> <p>Reference these bits by using the mask, shifted as follows:</p> <pre style="margin-left: 40px;"> BITB #IO\$M_LOOP@- 7,UCB\$B_TT_MAINT(R5) ;Set loopback mode </pre> <p>UCB\$B_TT_MAINT also defines the bit UCB\$V_TT_DSBL that, when set, indicates that the line has been disabled.</p>	IO\$M_LOOP	Set loopback mode.	IO\$M_UNLOOP	Reset loopback mode.	IO\$M_AUTXOF_ENA	Enable the use of auto XON/XOFF on this line. This is the default.	IO\$M_AUTXOF_DIS	Disable the use of auto XON/XOFF on this line.	IO\$M_LINE_OFF	Disable interrupts on this line.	IO\$M_LINE_ON	Reenable interrupts on this line.
IO\$M_LOOP	Set loopback mode.												
IO\$M_UNLOOP	Reset loopback mode.												
IO\$M_AUTXOF_ENA	Enable the use of auto XON/XOFF on this line. This is the default.												
IO\$M_AUTXOF_DIS	Disable the use of auto XON/XOFF on this line.												
IO\$M_LINE_OFF	Disable interrupts on this line.												
IO\$M_LINE_ON	Reenable interrupts on this line.												
UCB\$B_OLD*	The full name of this field is UCB\$B_TT_OLDPCPZORG; it currently serves as a filler byte.												
UCB\$L_TT_FBK*	Address of fallback block.												
UCB\$L_TT_RDVERIFY*	Address of read/verify table. Reserved for future use.												
UCB\$L_TT_CLASS1*	First class driver longword.												
UCB\$L_TT_CLASS2*	Second class driver longword.												
UCB\$L_TT_ACCPORNAM	Address of counted string.												
UCB\$L_TP_MAP*	UNIBUS/Q22-bus map registers.												
UCB\$B_TP_STAT	<p>DMA port-specific status.</p> <p>The following fields are defined within UCB\$B_TP_STAT.</p> <table style="width: 100%; border: none;"> <tr> <td style="width: 30%;">TTY\$V_TP_ABORT</td> <td>DMA abort requested on this line.</td> </tr> <tr> <td>TTY\$V_TP_ALLOC</td> <td>Allocate map fork in progress.</td> </tr> <tr> <td>TTY\$V_TP_DLLOC</td> <td>Deallocate map fork in progress.</td> </tr> </table>	TTY\$V_TP_ABORT	DMA abort requested on this line.	TTY\$V_TP_ALLOC	Allocate map fork in progress.	TTY\$V_TP_DLLOC	Deallocate map fork in progress.						
TTY\$V_TP_ABORT	DMA abort requested on this line.												
TTY\$V_TP_ALLOC	Allocate map fork in progress.												
TTY\$V_TP_DLLOC	Deallocate map fork in progress.												

2

VMS Macros Invoked by Drivers

This chapter describes VMS macros frequently used by device drivers. When referring to the macro descriptions contained herein, you should be aware of the following conventions:

- If an argument is enclosed in brackets, you can choose to include that argument or omit it.
- VMS assigns values by default to certain arguments. If you omit one of these arguments, the macro behaves as if you specified the argument with its default value. In the macro descriptions contained in this chapter, the format signifies such arguments by an equal sign (=) separating the argument from its keyword. For example:

SETIPL [ipl=31]

- If an argument takes a keyword value, you should specify the keyword value using all uppercase letters. For example:

**preserve=YES
condition=RESTORE**

General information about the structure of macros and their arguments appears in the *VAX MACRO and Instruction Set Reference Manual*.

VMS Macros Invoked by Drivers

ADPDISP

DESCRIPTION ADPDISP dispatches upon the following adapter characteristics:

<i>select</i>	Possible Value of <i>flag</i> in <i>addrlist</i>	Definition
ADAP_TYPE	UBA, MBA, GENBI, DR, or NULL. (See those symbols prefixed with AT\$ defined by the \$DCDEF macro in SYS\$LIBRARY:STARLET.MLB.)	Adapter type.
ADDR_BITS	18 or 22	Number of adapter address bits.
ADAP_MAPPING	YES or NO	Does adapter support mapping?
AUTOPURGE_DP	YES or NO	Does adapter support autopurging datapaths?
BUFFERED_DP	YES or NO	Does adapter support buffered datapaths?
DIRECT_VECTOR	YES or NO	Does adapter directly vector device interrupts?
ODD_XFER_BDP	YES or NO	Does adapter support odd-aligned transfers over its buffered data paths?
ODD_XFER_DDP	YES or NO	Does adapter support odd-aligned transfers over its direct data paths?
EXTENDED_MAPREG	YES or NO	Does adapter support extended set (8192) map registers?
QBUS	YES or NO	Is this a Q22-bus device?

Specification of **select=ADAP_TYPE** causes ADPDISP to generate a CASEW instruction using ADP\$W_ADPTYPE as an index into the case table. Specification of **select=ADDR_BITS** similarly causes ADPDISP to dispatch from the contents of ADP\$B_ADDR_BITS (16 or 22 bits). If any of the other conditions is specified for **select**, ADPDISP issues a BBC or BBS instruction on the contents of bit field ADP\$V_select in ADP\$W_ADPDISP_FLAGS.

You cannot use a single invocation of ADPDISP to dispatch on more than one adapter characteristic. For example, if an autopurging datapath that supports direct vectoring is being sought, you must use the ADPDISP macro twice.

ADPDISP requires that the address of an ADP, CRB, UCB, or ECRB be specified. If anything other than an ADP is specified, the **scratch** register is used in determining the ADP address.

VMS Macros Invoked by Drivers

ADPDISP

EXAMPLES

1

```
ADPDISP -  
  SELECT=ADAP_MAPPING, -  
  ADDRLIST=<<NO, 10$>, <YES, 20$>>, -  
  ADPADDR=R3
```

ADPDISP transfers control to the instruction at 10\$ if the adapter does not support mapping, or to 20\$ if it does. ADPDISP uses the value in R3 to locate the ADP.

2

```
ADPDISP -  
  SELECT=ADAP_TYPE, -  
  ADDRLIST=<<CI, 10$>, <MBA, 20$>, <UBA, 30$>>, -  
  UCBADDR=R5, -  
  SCRATCH=R1
```

ADPDISP transfers control to 10\$ if the adapter is a CI, 20\$ if the adapter is a MASSBUS adapter, and 30\$ if it is a UNIBUS adapter. ADPDISP determines the location of the ADP from a chain of pointers starting at the UCB address specified in R5. In doing so, it destroys the contents of scratch register R1.

3

```
ADPDISP -  
  SELECT=ADDR_BITS, -  
  ADDRLIST=<<18, 10$>, <22, 20$>>, -  
  ADPADDR=R3
```

ADPDISP transfers control to 10\$ for all adapters using an 18-bit address and 20\$ for all using a 22-bit address. The ADP address is supplied in R3.

BI_NODE_RESET

Initiates BIIC self-test on the specified VAXBI node.

FORMAT **BI_NODE_RESET** *csr*

PARAMETERS ***csr***
General purpose register that contains the address of the VAXBI node's control and status register (CSR).

DESCRIPTION The BI_NODE_RESET macro uses the recommended instruction sequence to disable arbitration on the specified VAXBI node, and sets the node reset and self-test status bits in the BIIC CSR. The use of any instruction sequence other than that defined by the BI_NODE_RESET macro to perform these actions may cause an undefined condition on the VAXBI bus.

VMS Macros Invoked by Drivers

CASE

CASE

Generates a CASE instruction and its associated table.

FORMAT **CASE** *src* ,*displist* [*,type=W*] [*,limit=#0*] [*,nmode=S^#*]

PARAMETERS ***src***
Source of the index value to be used with the CASE instruction.

displist
List of destinations to which control is to be dispatched, depending on the value of the index.

[type=W]
Data type of ***src*** (B, W, or L).

[limit=#0]
Lower limit of the value of ***src***.

[nmode=S^#]
Addressing mode used to reference the case-table entries; the default, short-literal mode, is good for up to 63 entries.

EXAMPLE

```
10$:     CASE -  
          src=ITEMC,  
          displist=<FIRST,SECOND,THIRD,FOURTH>
```

This invocation of the CASE macro expands to the following code:

```
          CASEW    ITEMC, #0, S^#<<30001$-30000$>/2>-1  
30000$:            .SIGNED_WORD       FIRST-30000$  
                  .SIGNED_WORD       SECOND-30000$  
                  .SIGNED_WORD       THIRD-30000$  
                  .SIGNED_WORD       FOURTH-30000$  
30001$:
```

CLASS_CTRL_INIT

Generates the common code that must be executed by the controller initialization routine of all terminal port drivers.

FORMAT **CLASS_CTRL_INIT** *dpt, vector*

PARAMETERS ***dpt***
Symbolic name of the port driver's driver prologue table.

vector
Address of the port driver vector table.

DESCRIPTION A terminal port driver's controller initialization routine invokes the CLASS_CTRL_INIT macro to relocate the class and port driver vector tables and perform other required initialization.

To use the CLASS_CTRL_INIT macro, the driver must include an invocation of the \$TTYMACS definition macro (from SYS\$LIBRARY:LIB.MLB).

VMS Macros Invoked by Drivers

CLASS_UNIT_INIT

CLASS_UNIT_INIT

Generates the common code that must be executed by the unit initialization routine of all terminal port drivers.

FORMAT CLASS_UNIT_INIT

DESCRIPTION A terminal port driver's unit initialization routine invokes the CLASS_UNIT_INIT macro to perform initialization tasks common to all port drivers. To use the CLASS_UNIT_INIT macro, the driver must include an invocation of the \$TTYMACS definition macro (from SYS\$LIBRARY:LIB.MLB).

The CLASS_UNIT_INIT macro binds the terminal port and class driver into a single, complete driver by initializing the following UCB fields as indicated:

Field	Contents
UCB\$L_TT_CLASS	Class driver vector table address
UCB\$L_TT_PORT	Port driver vector table address
UCB\$L_TT_GETNXT	Address of the class driver's get-next-character routine (CLASS_GETNXT)
UCB\$L_TT_PUTNXT	Address of the class driver's put-next-character routine (CLASS_PUTNXT)
UCB\$L_DDT	Address of the terminal class driver's driver dispatch table

Prior to invoking this macro, the unit initialization should place in R0 the address of the port driver vector table.

CPUDISP

Causes a branch to a specified address according to the CPU type of the VAX processor executing the macro code.

FORMAT **CPUDISP** *addrlist* [,*environ=VMS*] ,*continue=NO*

PARAMETERS ***addrlist***

List containing one or more pairs of arguments in the following format:

<CPU-type, destination>

The **CPU-type** parameter identifies the type or subtype of a VAX processor for which the macro is to generate a case table entry. The CPUDISP macro identifies the following VAX systems by type alone:

CPU Type	VAX System
9AQ	VAX 9000-2xx/9000-4xx
9RR	VAX 6000-4xx
9CC	VAX 6000-2xx/6000-3xx/62xx/63xx
8PS	VAX 8810/8820/8830
8NN	VAX 8530/8550/8700/8800
790	VAX 8600/8650
8SS	VAX 8200/8250/8300/8350
780	VAX-11/780 and VAX-11/785 ¹
785	VAX-11/785
750	VAX-11/750
730	VAX-11/730
670	VAX 4000-300
650	MicroVAX 3400/3600/3900-series system
520	VAX 3000FT
420	VAXstation 3100/MicroVAX 3100
410	VAXstation 2000/MicroVAX 2000
60	VAXstation 3520/3540
UV2	MicroVAX II

¹Because the VAX-11/785 has the same CPU type as the VAX-11/780, the CPUDISP macro contains special code to distinguish between the two processors. This code tests a bit within the processor's system identification register (PR\$_SID) that indicates whether it is a VAX-11/785.

VMS Macros Invoked by Drivers

CPUDISP

The CPUDISP macro identifies the following VAX systems by type and subtype:

CPU Type	Subtype	VAX System
UV		MicroVAX II processor-based system
	UV2	MicroVAX II
	410	VAXstation 2000/MicroVAX 2000
CV		CVAX processor-based system
	420	VAXstation 3100/MicroVAX 3100
	520	VAX 3000FT
	650	MicroVAX 3400/3600/3900-series system
	9CC	VAX 6200/6300-series system
	60	VAXstation 3520/3540
RV		CVAX-Rigel processor-based system
	9RR	VAX 6000-4xx
	670	VAX 4000-300

You can supply any combination of generic type and subtype in a single invocation of the CPUDISP macro. Should the CPUDISP macro code be executed on the appropriate processor, the following transfers of control are possible:

- If you specify a generic type but no subtype, CPUDISP causes the branch designated for the generic type to be taken for all of its subtypes.
- If you specify one or more subtypes but not the generic type, CPUDISP causes the branch designated for each subtype to be taken.
- If you specify *both* the generic type and one or more subtypes, CPUDISP causes the branch designated for each specified subtype to be taken. For those subtypes that you do not specify, CPUDISP causes the branch designated for the generic type to be taken.

The **destination** parameter contains the address to which the code generated by the invocation of the CPUDISP macro passes control to continue with CPU-specific processing.

[environ=VMS]

Identification of the run-time environment of the code generated by the CPUDISP macro. There is no need to change the default value of this argument.

continue=NO

Specifies whether execution should continue at the line immediately after the CPUDISP macro if the value at EXE\$GB_CPUATYPE does not correspond to any of the values specified as the **CPU-type** in the **addrlist** argument. A fatal bugcheck of UNSUPRTCPU occurs if the dispatching code does not find the executing processor identified in the **addrlist** and the value of **continue** is NO.

DESCRIPTION

The CPUDISP macro provides a means for transferring control to a specified destination depending on the CPU type of the executing processor. For those processors that do not have a unique CPU type, CPUDISP also provides the means to dispatch on a particular CPU subtype.

To accomplish this, CPUDISP builds one or two case tables. The first CASEB instruction uses words in the first case table to set up a transfer based on each **CPU-type** specified in the **addrlist** argument. CPUDISP constructs the second case table in the event it encounters a CPU subtype in the **addrlist**.

CPUDISP constructs appropriate symbolic constants for each **CPU-type** listed in **addrlist**, and compares them against the contents of EXE\$GB_CPUTYPE. These constants have the form PR\$_SID_TYPCPU-*type*.

For each CPU subtype it encounters in the **addrlist** argument, CPUDISP also constructs symbolic constants of the form PR\$_XSID_XX_YYY, where *xx* is the generic CPU type (for example CV) and *yyy* is the CPU subtype (420, 520, 650, 9CC, or 60 for CV). It compares the value of PR\$_XSID_XX_YYY against the contents of EXE\$GB_CPUDATA+15.

[mntver=+IOC\$MNTVER]

Address of the VMS routine that is called at the beginning and end of a mount verification operation. The default, IOC\$MNTVER, is suitable for all single-stream disk drives. Use of this field to call any other routine is reserved to Digital.

[cloneducb=+IOC\$RETURN]

Address of routine called when a UCB is cloned by the \$ASSIGN system service.

DESCRIPTION

The DDTAB macro creates a driver dispatch table (DDT). The table has a label of **devnam\$DDT**. Just preceding the table, DDTAB generates the driver code program section with the following statement:

```
.PSECT $$$115_DRIVER
```

The DDTAB macro writes the address of the VMS universal executive routine vector IOC\$RETURN into routine address fields of the DDT that are not supplied in the macro invocation (with the exception of the **mntver** argument). IOC\$RETURN simply executes an RSB instruction.

A plus sign (+) precedes the address of any specified routine that is part of VMS: that is, it is an address that is not relative to the location of the driver. No plus sign precedes the address of a routine (such as a start-I/O routine) that is part of the driver module.

EXAMPLE

```
DDTAB      -                ;DDT-creation macro
DEVNAM=XX, -                ;Name of device
START=XX_START,-          ;Start-I/O routine
FUNCTB=XX_FUNCABLE,-     ;FDT address
CANCEL=+IOC$CANCELIO,-   ;Cancel-I/O routine
REGDMP=XX_REGDUMP,-      ;Register dumping routine
DIAGBF=<<15*4>+<<3+5+1>*4>>,- ;Diagnostic buffer size
ERLGBF=<<15*4>+<1*4>+<EMB$L_DV_REGSAV>> ;Error message buffer size
```

This code excerpt uses the DDTAB macro to create a driver dispatch table for the XX device type. Note that because the cancel-I/O routine is part of VMS, its address is preceded by a plus sign (+).

VMS Macros Invoked by Drivers

\$DEF

\$DEF

Defines a data-structure field within the context of a \$DEFINI macro.

FORMAT **\$DEF** *sym* [*,alloc*] [*,siz*]

PARAMETERS *sym*
Name of the symbol by which the field is to be accessed.

[alloc]
Block-storage-allocation directives, one of the following: .BLKB, .BLKW, .BLKL, .BLKQ, or .BLKO.

[siz]
Number of block storage units to allocate.

DESCRIPTION See the descriptions of the \$DEFINI, \$DEFEND, _VIELD, and \$EQLST macros for additional information on defining symbols for data structure fields.

You can define a second symbolic name for a single field, using the \$DEF macro a second time immediately following the first definition, leaving the **alloc** argument blank in the first definition. The following example does this, equating SYNONYM2 with LABEL2:

```
$DEFINI JLB                                    ;Start structure definition
$DEF LABEL1 .BLKL 1                         ;First JLB field
$DEF SYNONYM2                               ;Synonym for LABEL2 field
$DEF LABEL2 .BLKL 1                         ;Second JLB field
$DEF LABEL3 .BLKL 1                         ;Third JLB field
$DEFEND JLB                                 ;End of JLB structure
```

For another example of the use of the \$DEF macro, see the description of the \$DEFINI macro.

\$DEFEND

Ends the scope of the \$DEFINI macro, thereby completing the definition of fields within a data structure.

FORMAT **\$DEFEND** *struc*

PARAMETERS *struc*
Name of the structure that is being defined.

DESCRIPTION See the descriptions of the \$DEFINI, _VIELD, and \$EQLST macros for additional information on defining symbols for data structure fields.

VMS Macros Invoked by Drivers

\$DEFINI

\$DEFINI

Begins the definition of a data structure.

FORMAT **\$DEFINI** *struc* [*,gbl=LOCAL*] [*,dot=0*]

PARAMETERS ***struc***
Name of the data structure that is being defined.

[gbl=LOCAL]

Specifies whether the symbols defined for this data structure are to be local or global symbols. The default is to make them local.

To make the definitions of symbols global, you must specify **GLOBAL** for the value of the **gbl** argument.

[dot=0]

Offset from the beginning of the data structure of the first field to be defined. The \$DEFINI macro moves this value into the current location counter (.).

DESCRIPTION The \$DEF macro defines fields within the structure specified by the invocation of the \$DEFINI macro, and the \$DEFEND macro ends the definition. See the descriptions of the _VIELD and \$EQLST macros for additional information on defining symbols for data structure fields.

EXAMPLE

```
$DEFINI UCB,,UCB$K_LCL_DISK_LENGTH
;Start UCB extension, begin definitions
; at end of local disk UCB extension
$DEF  UCB_W_DL_PBCR  .BLKW 1 ;Partial byte count
$DEF  UCB_W_DL_CS   .BLKW 1 ;Control status register
$DEF  UCB_W_DL_BA   .BLKW 1 ;Bus address register
$DEF  UCB_A_DL_BUF_PA .BLKL 1 ;Physical buffer physical address
$DEF  UCB_K_DL_LEN  .BLKW 1 ;Length of extended UCB
$DEFEND UCB
```

This code excerpt, when assembled in VMS Version 5.0, produces the following symbol listing:

```
.
.
.
UCB_A_DL_BUF_PA      000000D2
UCB_K_DL_LEN        000000D6
UCB$K_LCL_DISK_LENGTH = 000000CC
UCB_W_DL_BA         000000D0
UCB_W_DL_CS         000000CE
UCB_W_DL_PBCR       000000CC
```


VMS Macros Invoked by Drivers

DEVICELOCK

- Calls either SMP\$ACQUIREL or SMP\$ACQNOIPL, depending upon the presence of **condition=NOSETIPL**. SMP\$ACQUIREL raises IPL to device IPL prior to obtaining the lock, determining appropriate IPL from the device lock's data structure (SPL\$B_IPL).

In both processing environments, the DEVICELOCK macro performs the following tasks:

- Preserves the current IPL at the specified location (if **savipl** is specified)
- Sets the SMP-modified bit in the driver prologue table (DPT\$V_SMPMOD in DPT\$L_FLAGS)

EXAMPLE

```
DEVICELOCK -
    LOCKADDR=UCB$L_DLCK(R5),- ;Lock device access
    LOCKIPL=UCB$B_DIPL(R5),- ;Raise IPL
    SAVIPL=- (SP),- ;Save current IPL
    PRESERVE=YES ;Save R0
SETIPL #31 ;Disable all interrupts
BBC #UCB$V_POWER,- ;If clear - no power failure
    UCB$W_STS(R5),L1 ;...
    ;Service power failure!
.
.
.
DEVICEUNLOCK -
    LOCKADDR=UCB$L_DLCK(R5),- ;Unlock device access
    NEWIPL=(SP)+,- ;Restore IPL
    PRESERVE=YES ;Save R0
BRW RETREG ;Exit
L1: ;Return for no power failure
.
.
.
WFIKPCH RETREG,#2 ;Wait for interrupt
```

The start-I/O routine of DLDRIVER invokes the DEVICELOCK macro to synchronize access to the device's registers and UCB fields. Thus synchronized at device IPL, and holding the device lock in a VMS multiprocessing environment, the routine raises IPL to IPL\$_POWER (IPL 31) to check for a power failure on the local processor. If a power failure has occurred, the routine releases the device lock and pops the saved IPL from the stack before servicing the failure. If a power failure has not occurred, the routine branches to set up the I/O request. Note that, in this instance, it is the wait-for-interrupt routine, invoked by the WFIKPCH macro, that issues the DEVICEUNLOCK macro and pops the saved IPL from the stack.

VMS Macros Invoked by Drivers

DEVICEUNLOCK

EXAMPLE

```
DEVICELOCK -
    LOCKADDR=UCB$L_DLCK(R5),- ;Lock device access
    CONDITION=NOSETIPL,- ;Do not set IPL
    PRESERVE=NO ;Do not preserve R0
.
.
.
20$: MOVQ   UCB$L_FR3(R5),R3 ;Restore driver context
    JSB    @UCB$L_FPC(R5) ;Call driver at interrupt return address
40$: DEVICEUNLOCK -
    LOCKADDR=UCB$L_DLCK(R5),- ;Unlock device access
    PRESERVE=NO ;Do not preserve R0
```

When the device interrupts, DLDRIVER's interrupt service routine immediately obtains the device lock so that it can examine device registers and preserve their contents. It then calls the driver's start-I/O routine at the location in which it initiated device activity. The routine forks and returns control to the interrupt service routine, which releases the device lock.

DPTAB

Generates a driver prologue table (DPT) in a program section called \$\$\$105_PROLOGUE.

FORMAT **DPTAB** *end ,adapter ,[flags=0] ,ucbsize ,[unload] ,[maxunits=8] ,[defunits=1] ,[deliver] ,[vector] ,name [,psect=\$\$\$105_PROLOGUE] [,smp=NO] [,decode]*

PARAMETERS

end
Address of the end of the driver.

adapter

Type of adapter (as indicated by the symbols prefixed by AT\$ defined by the \$DCDEF macro in SYS\$LIBRARY:STARLET.MLB). The adapter type can be any of the following:

UBA	UNIBUS adapter or Q22-bus interface
MBA	MASSBUS adapter
GENBI	Generic VAXBI adapter
DR	DR device
NULL	No actual device for driver

[flags=0]

Flags used in loading the driver. Drivers use the following flags:

DPT\$M_SVP	Indicates that the driver requires a permanently allocated system page. Disk drivers use this SPTE during ECC correction and when using the system routines IOC\$MOVFRUSER and IOC\$MOVTOUSER. When this flag is set, the driver-loading procedure allocates a permanent system page-table entry (SPTE) for the device. It stores an index to the virtual address of the SPTE in UCB\$L_SVPN when it creates the UCB. A driver can calculate the system virtual address of the page corresponding to this index by using the following formula:
------------	--

$$(index * 200_{16}) + 80000000_{16}$$

DPT\$M_NOUNLOAD	Indicates that the driver cannot be reloaded. When this bit is set, the driver can be unloaded only by rebooting the system.
-----------------	--

VMS Macros Invoked by Drivers

DPTAB

DPT\$M_SMPMOD Indicates that the driver has been designed to execute within a VMS multiprocessing environment. Use of any of the VMS multiprocessing synchronization macros (DEVICELOCK/DEVICEUNLOCK, FORKLOCK /FORKUNLOCK, or LOCK/UNLOCK) automatically sets this flag, as long as the code using the macro resides in the same module as the invocation of DPTAB.

ucbsize

Size in bytes of each UCB the driver-loading procedure creates for devices supported by the driver. This required argument allows drivers to extend the UCB to store device-dependent data describing an I/O operation. Figure 1-20 describes the VMS-defined extensions to the UCB and discusses the means by which a driver can define a device-specific extension.

[unload]

Address of the driver routine invoked by the SYSGEN RELOAD command before it unloads an old version of the driver to load a new version. The driver-loading procedure calls this routine before reinitializing all controllers and device units associated with the driver.

[maxunits=8]

Maximum number of units that this driver supports on a controller. This field affects the size of the IDB created by the driver-loading procedure. If you omit the **maxunits** argument, the default is eight units. You can override the value specified in the DPT by using the /MAXUNITS qualifier to the SYSGEN CONNECT command.

[defunits=1]

Maximum number of UCBs to be created by SYSGEN's AUTOCONFIGURE command (one for each device unit to be configured). The unit numbers assigned are zero through **defunits**-1.

If you do not specify the **deliver** argument, AUTOCONFIGURE creates the number of units specified by **defunits**. If you specify the address of a unit delivery routine in the **deliver** argument, AUTOCONFIGURE calls that routine to determine whether to create each UCB automatically.

[deliver]

Address of the driver unit delivery routine. The unit delivery routine determines which device units supported by this driver the SYSGEN AUTOCONFIGURE command should configure automatically. If you omit the **deliver** argument, the AUTOCONFIGURE command creates the number of units specified by the **defunits** argument.

[vector]

Address of a driver-specific transfer vector. A terminal port driver specifies the address of its vector table in this argument.

name

Name of the device driver. The driver-loading procedure will permit the loading of only one copy of the driver associated with this name. A driver name can be up to 11 alphabetic characters and, by convention, is formed by appending the string DRIVER to the 2-alphabetic-character generic device name, for example, QBDRIVER. (Digital reserves to customers driver names beginning with the letters J and Q.)

[psect=\$\$\$105_PROLOGUE]

Program section in which the DPT is created. The default value of this argument is required for all non-Digital-supplied device drivers.

[smp=NO]

Indication of whether the driver is suitably synchronized to execute in a VMS multiprocessing system. Note that use of any of the spin lock synchronization macros in a device driver causes the DPTAB macro to indicate multiprocessing synchronization.

[decode]

Offset to name used by workstation windowing software.

DESCRIPTION

The DPTAB macro, in conjunction with invocations of the DPT_STORE macro, creates a driver prologue table (DPT). The DPTAB macro places information in the DPT that allows the driver-loading procedure to identify the driver and the devices it supports. The DPTAB macro, in invoking the \$SPLCODDEF definition macro, also defines the spin lock indexes used in the DPT_STORE, FORKLOCK, and LOCK macros.

EXAMPLE

```

DPTAB      -                               ;DPT-creation macro
           END=XA_END,-                     ;End of driver label
           ADAPTER=UBA,-                     ;Adapter type
           FLAGS=<DPT$M_SVP!-                ;Allocate permanent SPTE
             DPT$M_SMPMOD>,-                 ;Multiprocessing driver
           UCBSIZE=UCB$K_SIZE,-              ;UCB size
           NAME=XADRIVER                     ;Driver name
DPT_STORE  INIT                             ;Start of load initialization table
DPT_STORE  UCB,UCB$B_FLCK,B,-                ;Fork lock index
           SPL$C_IOLOCK8
DPT_STORE  UCB,UCB$B_DIPL,B,22               ;Device interrupt IPL
DPT_STORE  UCB,UCB$L_DEVCHAR,L,<-            ;Device characteristics
           DEV$M_AVL!-                       ;Available
           DEV$M_RTM!-                       ;Real time device
           DEV$M_ELG!-                       ;Error logging enabled
           DEV$M_IDV!-                       ;Input device
           DEV$M_ODV>                       ;Output device
DPT_STORE  UCB,UCB$B_DEVCLASS,B,-            ;Device class
           DC$_REALTIME
DPT_STORE  UCB,UCB$B_DEVTYPE,B,-            ;Device type
           DT$_DR11W
DPT_STORE  UCB,UCB$W_DEVBUFSIZ,W,-          ;Default buffer size
           XA_DEF_BUFSIZ
DPT_STORE  REINIT                            ;Start of reload initialization table
DPT_STORE  DDB,DDB$L_DDT,D,XA$DDT           ;Address of DDT
DPT_STORE  CRB,CRB$L_INTD+VEC$L_ISR,D,-      ;Address of interrupt service routine
           XA_INTERRUPT
DPT_STORE  CRB,CRB$L_INTD+VEC$L_INITIAL,D,-  ;Address of controller initialization routine
           XA_CONTROL_INIT
DPT_STORE  END                               ;End of initialization

```

This excerpt from XADRIVER.MAR contains the DPTAB macro and the series of DPT_STORE macros that create its driver prologue table.

VMS Macros Invoked by Drivers

DPT_STORE

DPT_STORE

Instructs the VMS driver-loading procedure to store values in a table or data structure.

FORMAT **DPT_STORE** *str_type* , *str_off* , *oper* , *exp* [, *pos*] [, *size*]

PARAMETERS ***str_type***

Type of data structure (CRB, DDB, IDB, ORB, or UCB) into which the driver-loading procedure is to store the specified data, or a label denoting a table marker. Table marker labels indicate the start of a list of DPT_STORE macro invocations that store information for the driver-loading procedure in the driver initialization table and driver reinitialization table sections of the DPT. If this argument is a table marker label, no other argument is allowed. The following labels are used:

- INIT Indicates the start of fields to initialize when the driver is loaded
- REINIT Indicates the start of additional fields to initialize when the driver is loaded and reinitialized when the driver is reloaded
- END Indicates the end of the two lists

str_off

Unsigned offset into the data structure in which the data is to be stored. This value cannot be more than 65,535 bytes.

oper

Type of storage operation, one of the following:

Type	Meaning
B	Write a byte value.
W	Write a word value.
L	Write a longword value.
D	Write an address relative to the beginning of the driver.
V	Write a bit field. If you specify a V in the oper argument, the driver-loading procedure uses the exp , pos , and size arguments as operands to an INSV instruction.

If an at sign (@) precedes the **oper** argument, the **exp** argument indicates the address of the data that is to be stored and not the data itself.

exp

Expression indicating the value with which the driver-loading procedure is to initialize the indicated field. If an at-sign character (@) precedes the **oper** argument, the **exp** argument indicates the address of the data with which to initialize the field. For example, the following macro indicates that the contents of the location DEVICE_CHARS are to be written into the DEVCHAR field of the UCB.

DPT_STORE UCB,UCB\$L_DEVCHAR,@L,DEVICE_CHARS

[pos]

Starting bit position within the specified field; used only if **oper=V**.

[size]

Number of bits to be written; used only if **oper=V**.

DESCRIPTION

The DPT_STORE macro places information in the DPT that the driver-loading procedure uses to load specified values into specified fields. The DPT_STORE macro accepts two lists of fields:

- Fields to be initialized only when a driver is first loaded
- Fields to be initialized when a driver is first loaded and reinitialized if the driver is reloaded

The DPTAB macro stores the relative addresses of these two lists, called initialization and reinitialization tables, in the DPT. A driver constructs the initialization tables by following the DPTAB macro with one or more invocations of the DPT_STORE macro.

Drivers use the DPT_STORE macro with the **INIT** table marker label to begin a list of DPT_STORE invocations that supply initialization data for the following fields:

UCB\$B_FLCK Index of the fork lock under which the driver performs fork processing. Fork lock indexes are defined by the \$SPLCODDEF definition macro (invoked by DPTAB) as follows:

IPL	Fork Lock Index
8	SPL\$I_IOCTL8
9	SPL\$I_IOCTL9
10	SPL\$I_IOCTL10
11	SPL\$I_IOCTL11

UCB\$B_DIPL Device interrupt priority level.

Other commonly initialized fields are as follows:

UCB\$L_DEVCHAR Device characteristics.

UCB\$B_DEVCLASS Device class.

UCB\$B_DEVTYPE Device type.

UCB\$W_DEVBUFSIZ Default buffer size.

UCB\$Q_DEVDEPEND Device-dependent parameters.

Drivers use the DPT_STORE macro with the **REINIT** table marker label to begin a list of DPT_STORE invocations that supply initialization and reinitialization data for the following fields:

DDB\$L_DDT Driver dispatch table. Every driver must specify a value for this field.

VMS Macros Invoked by Drivers

DPT_STORE

CRB\$_INTD+ VEC\$_ISR	Interrupt service routine.
CRB\$_INTD2+ VEC\$_ISR	Interrupt service routine for second interrupt vector.
CRB\$_INTD+ VEC\$_INITIAL	Controller initialization routine.
CRB\$_INTD+ VEC\$_UNITINIT	Unit initialization routine (for UNIBUS, Q22 bus, and generic VAXBI device drivers). Note that MASSBUS drivers must specify the address of the unit initialization routine in an invocation of the DDTAB macro.

For an example of the use of the DPT_STORE macro, see the description of the DPTAB macro.

DSBINT

Blocks interrupts from occurring on the local processor at or below a specified IPL.

FORMAT	<p>DSBINT <i>[ipl=31] [,dst=-(SP)]</i> <i>[,environ=MULTIPROCESSOR]</i></p>
---------------	--

PARAMETERS	<p><i>[ipl=31]</i> IPL at which to block interrupts. If no ipl is specified, the default is IPL 31, which blocks all interrupts.</p> <p><i>[dst=-(SP)]</i> Location in which to save the current IPL. If no destination is specified, the current IPL is pushed onto the stack.</p> <p><i>[environ=MULTIPROCESSOR]</i> Processing environment in which the DSBINT synchronization macro is to be assembled. If you do not specify environ, or if you do specify environ=MULTIPROCESSOR, the DSBINT macro generates the following assembly-time warning message, where <i>xx</i> is an IPL above IPL 2:</p>
-------------------	--

%MACRO-W-GENWARN, Generated WARNING: Raising IPL to #xx provides no multiprocessing synchronization

If you are certain that the purpose of the macro invocation is to block only local processor events, you can disable the warning message by including **environ=UNIPROCESSOR** in the invocation.

DESCRIPTION	<p>The DSBINT macro first stores the current IPL of the local processor and then moves the specified IPL into the processor's IPL register (PR\$_IPL).</p> <p>Note that the DSBINT and ENBINT macros provide full synchronization only in a uniprocessing environment. In a multiprocessor configuration, DSBINT and ENBINT are suitable only for blocking events on the local processor. To provide synchronized access to system resources and devices in a multiprocessing environment, you <i>must</i> use the DEVICELOCK /DEVICEUNLOCK, FORKLOCK/FORKUNLOCK, and LOCK/UNLOCK macros.</p>
--------------------	---

VMS Macros Invoked by Drivers

ENBINT

ENBINT

Lowens the local processor's IPL to a specified value, thus permitting interrupts to occur at or beneath the current IPL.

FORMAT **ENBINT** *[src=(SP)+]*

PARAMETERS *[src=(SP)+]*
Location containing the IPL to be restored to the processor IPL register (PR\$_IPL) of the local processor. If you do not specify a value in **src**, ENBINT moves the value on the top of the stack into PR\$_IPL.

DESCRIPTION The ENBINT macro complements the actions of the DSBINT macro, restoring an IPL value to PR\$_IPL. Procedures invoke this macro to lower IPL to a previously saved level. If an interrupt is pending at the current IPL or at any IPL above the IPL specified by **src**, the current procedure is immediately interrupted.

Note that the DSBINT and ENBINT macros only provide full synchronization in a uniprocessor environment. In multiprocessor configurations, DSBINT and ENBINT are only suitable for blocking events on the local processor. To provide synchronized access to system resources and devices in a multiprocessing environment, you *must* use the DEVICELock/DEVICEUNLOCK, FORKLock/FORKUNLOCK, and LOCK/UNLOCK macros.

\$EQULST

Defines a list of symbols and assigns values to the symbols.

FORMAT **\$EQULST** *prefix* ,*[gbl=LOCAL]* ,*init* ,*[incr=1]* ,*list*

PARAMETERS

prefix

Prefix to be used in forming the names of the symbols.

[gbl=LOCAL]

Scope of the definition of the symbol, either **LOCAL**, the default, or **GLOBAL**.

init

Value to be assigned to the first symbol in the list.

[incr=1]

Increment by which to increase the value of each succeeding symbol in the list. The default is 1.

list

List of symbols to be defined. Each element in the list can have one of the following forms:

<**symbol**> — where **symbol** is the string appended to the prefix, forming the name of the symbol; the value of the symbol is assigned based on the values of **init** and **incr**.

<**symbol,value**> — where **symbol** is the string that is appended to the prefix, forming the name of the symbol, and **value** specifies the value of the symbol.

DESCRIPTION

See the descriptions of the \$DEFINI and _VIELD macros for additional information on defining symbols for data structure fields.

VMS Macros Invoked by Drivers

\$EQLST

EXAMPLE

```
$EQLST  XA_K_,,0,1,<-                ;Define CSR bit values
        <fnct1,2>-
        <fnct2,4>-
        <fnct3,8>-
        <statusa,2048>-
        <statusb,1024>-
        <statusc,512>-
        >
```

This code excerpt produces the following symbols:

XA_K_FNCT1	=	00000002
XA_K_FNCT2	=	00000004
XA_K_FNCT3	=	00000008
XA_K_STATUSA	=	00000800
XA_K_STATUSB	=	00000400
XA_K_STATUSC	=	00000200

FIND_CPU_DATA

Locates the start of the current process's per-CPU database area (CPU).

FORMAT **FIND_CPU_DATA** *reg* [*,amod=G^*] [*,istack=NO*]

PARAMETERS *reg*
Register to receive the base virtual address of the current processor's per-CPU database structure (CPU)).

[amod=G^]
Addressing mode.

[istack=NO]
Mechanism by which the base of the per-CPU database structure is calculated. Use **istack=YES** used only when it is certain that the processor is executing on the interrupt stack. The mechanism used when **istack=NO** is somewhat slower, but works whether the processor is executing on the interrupt stack or kernel stack.

DESCRIPTION The FIND_CPU_DATA macro loads the starting virtual address of the current processor's per-CPU database (CPU) into the specified register. A driver generally invokes the FIND_CPU_DATA macro in the process of determining the current process of the current CPU when executing in system context.

Such a driver must adhere to the following rules:

- It must invoke the FIND_CPU_DATA macro in kernel mode at or above IPL\$_RESCHED.
- It must ensure that it will not be rescheduled after issuing the macro while it is using the information returned by FIND_CPU_DATA. It typically does this by remaining at IPL\$_RESCHED or greater.

EXAMPLE

```
FIND_CPU_DATA R0
MOVL CPU$L_CURPCB (R0) , R1
```

The FIND_CPU_DATA macro returns the starting virtual address of the current processor's per-CPU database in R0. The subsequent MOVL instruction obtains the address of the process currently active on that processor and places it in R1.

VMS Macros Invoked by Drivers

FORK

FORK

Creates a fork process, in which context the code that follows the macro invocation executes.

FORMAT

FORK

DESCRIPTION

The FORK macro calls EXE\$FORK to create a fork process. When the FORK macro is invoked, the following registers must contain the values listed:

Register	Contents
R3	Contents to be placed in R3 of the fork process
R4	Contents to be placed in R4 of the fork process
R5	Address of fork block
00(SP)	Address of caller's caller

Unlike EXE\$IOfORK, EXE\$FORK does not disable device timeouts by clearing the UCB\$V_TIM bit in the field UCB\$L_STS.

VMS Macros Invoked by Drivers

FORKLOCK

- If offset FKB\$B_FLCK (FKB\$B_FIPL) contains a fork IPL, it sets IPL to that fork IPL.

In a *multiprocessing* environment, the FORKLOCK macro stores the fork lock index in R0 and calls SMP\$ACQUIRE. SMP\$ACQUIRE uses the value in R0 to locate the fork lock structure in the system spin lock database (a pointer to which is located at SMP\$AR_SPNLKVEC). Prior to securing the fork lock, SMP\$ACQUIRE raises IPL to its associated IPL (SPL\$B_IPL).

In both processing environments, the FORKLOCK macro performs the following tasks:

- Preserves R0 through the macro call (if **preserve=YES** is specified)
- Preserves the current IPL at the specified location (if **savipl** is specified)
- Sets the SMP-modified bit in the driver prologue table (DPT\$V_SMPMOD in DPT\$L_FLAGS)

EXAMPLE

```
FORKLOCK -
    LOCK=UCB$B_FLCK(R5),-      ;Lock fork database
    SAVIPL=-(SP),-           ;Save the current IPL
    PRESERVE=NO               ;Do not preserve R0
INCW  UCB$W_QLEN(R5)         ;Bump device queue length
BBSS  #UCB$V_BSY,UCB$W_STS(R5),-
    20$                       ;If set, device is busy
PUSHL R5                     ;Save UCB address
BSBW  IOC$INITIATE          ;Initiate I/O function
POPL  R5                     ;Restore UCB address
FORKUNLOCK -
    LOCK=UCB$B_FLCK(R5),-      ;Unlock fork database
    NEWIPL=(SP)+,-           ;Restore previous IPL
    PRESERVE=NO               ;Do not preserve R0
    RSB
.
.
.
20$:                               ;Place IRP in UCB pending-I/O queue
```

The VMS routine that determines whether a device can immediately service an I/O request synchronizes its access to the fork database by invoking the FORKLOCK macro. The FORKLOCK macro raises IPL to fork IPL and, in a multiprocessing environment, obtains the corresponding fork lock.

Thus synchronized, the VMS routine tests a bit in the UCB to determine whether the device is busy. If the device is not busy, VMS calls a routine that initiates driver processing of the I/O request, still at fork IPL and holding the fork lock. Later, possibly with an invocation of the WFIKPCH macro, the driver start-I/O routine returns control to this routine, which issues the FORKUNLOCK macro to relinquish fork level synchronization.

VMS Macros Invoked by Drivers

FORKUNLOCK

In either processing environment, the FORKUNLOCK macro sets the SMP-modified bit in the driver prologue table (DPT\$V_SMPMOD in DPT\$L_FLAGS).

For an example of the use of the FORKUNLOCK macro, see the description of the FORKLOCK macro.

FUNCTAB

Creates a driver's function decision table (FDT) and generates FDT entries.

FORMAT **FUNCTAB** *[action],codes*

PARAMETERS *[action]*

Address of an FDT routine that VMS calls when preprocessing an I/O request whose function code matches a function indicated in the **codes** argument. A plus sign (+) precedes the address of any specified FDT routine that is part of VMS. No plus sign precedes the address of an FDT routine that is contained within the driver module.

You cannot specify an **action** argument in a driver's first two invocations of the FUNCTAB macro.

codes

List of I/O function codes that VMS preprocessing services by calling the FDT routine specified in the **action** argument of the FUNCTAB macro invocation. The macro expansion prefixes each code with the string IO\$_; for example, READVBLK expands to IO\$_READVBLK.

DESCRIPTION

A device driver uses several invocations of the FUNCTAB macro to generate the three components of a function decision table:

- The list of valid I/O function codes
- The list of buffered I/O function codes
- One or more FDT entries

The first two invocations of the FUNCTAB macro in a driver generate the lists of valid I/O functions and buffered I/O functions, respectively. These invocations include the **codes** argument, but not the **action** argument. If no buffered I/O functions are defined for the device, the **codes** argument to the second invocation of the FUNCTAB macro specifies an empty list.

Each succeeding invocation of the FUNCTAB macro generates an FDT entry. Each FDT entry specifies all or a subset of the valid I/O function codes and the address of an FDT routine that performs I/O preprocessing for those function codes. You can specify any valid I/O function code in more than one of these FUNCTAB macro invocations, thus causing more than one FDT routine to be called for a single valid I/O function code.

VMS Macros Invoked by Drivers

FUNCTAB

EXAMPLE

```
XX_FUNCTABLE:                                ;Function decision table
FUNCTAB  ,-                                  ;Valid functions
<READLBLK,-                                  ;Read logical block
READPBLK,-                                  ;Read physical block
READVBLK,-                                  ;Read virtual block
SENSEMODE,-                                  ;Sense reader mode
SENSECHAR,-                                  ;Sense reader characteristics
SETMODE,-                                    ;Set reader mode
SETCHAR,-                                    ;Set reader characteristics
>
FUNCTAB  ,-                                  ;Buffered-I/O functions
<READLBLK,-                                  ;Read logical block
READPBLK,-                                  ;Read physical block
READVBLK,-                                  ;Read virtual block
SENSEMODE,-                                  ;Sense reader mode
SENSECHAR,-                                  ;Sense reader characteristics
SETMODE,-                                    ;Set reader mode
SETCHAR,-                                    ;Set reader characteristics
>
FUNCTAB  XX_READ,-                            ;Read function FDT routine
<READLBLK,-                                  ;Read logical block
READPBLK,-                                  ;Read physical block
READVBLK,-                                  ;Read virtual block
>
FUNCTAB  +EXE$SETMODE,-                       ;Set mode/characteristics FDT routine
<SETCHAR,-                                  ;Set reader characteristics
SETMODE,-                                    ;Set reader mode
>
FUNCTAB  +EXE$SENSEMODE,-                   ;Sense mode/characteristics FDT routine
<SENSECHAR,-                                ;Sense reader characteristics
SENSEMODE,-                                  ;Sense reader mode
>
```

This function decision table specifies that the routine `XX_READ` be called for all read functions that are valid for the device. `XX_READ` appears later in the driver module. VMS I/O preprocessing will call routines `EXE$SETMODE` and `EXE$SENSEMODE` for the device's set-characteristics and sense-mode functions. Because each of these routines is part of VMS, a plus sign (+) precedes its name in the `FUNCTAB` macro argument.

IFNORD, IFNOWRT, IFRD, IFWRT

Determines the read or write accessibility of a range of memory locations.

FORMAT { IFNORD
 IFNOWRT
 IFRD
 IFWRT } *siz,adr,dest [,mode=#0]*

PARAMETERS *siz*
Offset of the last byte to check from the first byte to check, a number less than or equal to 512.

adr
Address of first byte to check.

dest
Address to which the macro transfers control, according to the following conditions:

Macro	Condition
IFNORD	If either of the specified bytes cannot be read in the specified access mode
IFNOWRT	If either of the specified bytes cannot be written in the specified access mode
IFRD	If both bytes can be read in the specified access mode
IFWRT	If both bytes can be written in the specified access mode

[mode=#0]
Mode in which access is to be checked; zero, the default, causes the check to be performed in the mode contained in the previous-mode field of the current PSL.

DESCRIPTION The IFNORD and IFRD macros use the PROBER instruction to check the read accessibility of the specified range of memory by checking the accessibility of the first and last bytes in that range. The IFNORD macro passes control to the specified destination if either of the specified bytes cannot be read in the specified access mode. The IFRD macro transfers control if both bytes can be read in the specified access mode. Otherwise, the macros transfer to the next in-line instruction.

The IFNOWRT and IFWRT macros use the PROBEW instruction to check the write accessibility of the specified range of memory by checking the accessibility of the first and last bytes in that range. The IFNOWRT macro passes control to the specified destination if either of the specified

VMS Macros Invoked by Drivers

IFNORD, IFNOWRT, IFRD, IFWRT

bytes cannot be written in the specified access mode. The IFWRT macro transfers control to the specified destination if both bytes can be written in the specified access mode. Otherwise, the macros transfer to the next in-line instruction.

EXAMPLE

```
MOVZWL    $SS_ACCVIO,R0           ;Assume read access failure
MOVL      ENTRY_LIST(AP),R11      ;Get address of entry point list
IFRD      #4*4,(R11),50$          ;Branch forward if process
                                           ; has read access
BRW       ERROR                   ;Otherwise stop with error
.
.
.
```

The connect-to-interrupt driver uses the IFRD macro to verify that the process has read access to the four longwords that make up the entry point list. The address of the entry point list was specified in the **p2** argument of the \$QIO request to the driver.

INVALIDATE_TB

Allows a single page-table entry (PTE) to be modified while any translation buffer entry that maps it is invalidated, or invalidates the entire translation buffer.

FORMAT **INVALIDATE_TB** *[addr, inst1 [,inst2] [,inst3] [,inst4] [,inst5] [,inst6] [,save_r2=YES] [,checks=YES]]*

PARAMETERS **[addr]**
 Virtual address mapped by the PTE for which invalidation is required. If **addr** is blank, then the macro invalidates all PTEs in the translation buffer.

[inst1]
 First instruction that modifies the PTE.

[inst2]
 Second instruction that modifies the PTE.

[inst3]
 Third instruction that modifies the PTE.

[inst4]
 Fourth instruction that modifies the PTE.

[inst5]
 Fifth instruction that modifies the PTE.

[inst6]
 Sixth instruction that modifies the PTE.

[save_r2=YES]
 Indication that the value in R2 at the invocation of this macro should be preserved across the macro call. By default, INVALIDATE_TB preserves the value in R2; any value but **YES** supplied in this argument overrides this behavior.

[checks=YES]
 Argument enabling or disabling the generation of assembly-time warning messages that indicate misuse of the macro. When any value but **YES** is supplied in the **checks** argument, the INVALIDATE_TB macro does not generate these messages.

VMS Macros Invoked by Drivers

INVALIDATE_TB

DESCRIPTION

When privileged code alters page mapping information, modifying a valid PTE in an active page table, it must notify the operating system. The operating system then takes suitable steps to invalidate all translation buffer entries that reference this PTE.

The INVALIDATE_TB macro allows you modify a single PTE and invalidate a single translation buffer cache entry by supplying the virtual address mapped by the PTE in the **addr** argument and at least one instruction argument. INVALIDATE_TB executes up to six instructions that modify the PTE while preventing all other processors in the system from referencing the page it maps. Because the INVALIDATE_TB macro calls system routines that rely on the stack contents and use R2, none of the specified instruction arguments should reference the stack or use R2.

To invalidate the entire translation buffer (without modifying PTEs), invoke the INVALIDATE_TB macro with no **addr** and instruction arguments. Note that, if the **addr** argument is not present and any instruction arguments are specified, the INVALIDATE_TB macro invalidates the entire translation buffer but does not execute any of the instructions. In this case, if **checks=YES** is not overridden, the macro generates an assembly-time warning message if any instruction arguments are present.

To invoke INVALIDATE_TB, code must be executing at or below IPL\$_INVALIDATE, holding—in a VMS multiprocessing environment—no spin lock ranked higher than INVALIDATE. If you issue the INVALIDATE_TB macro from pageable code, you must ensure that the location of the code has been locked in memory.

EXAMPLE

```
MOVL      8(SP),R2                ;Load virtual address to invalidate
MOVL     12(SP),R3                ;Load address of PTE
INVALIDATE_TB  R2,-              ;Invalidate translation buffer
          INST1=<BICL2 #PTE$M_VALID, (R3)> ;Clear PTE valid bit
```

The INVALIDATE_TB macro causes the PTE corresponding to the virtual address supplied in R2 to be flushed from the system's translation buffers. The macro causes the specified BICL2 instruction to be executed while other processors in the system are prevented from referencing the stale PTE.

IOFORK

Disables timeouts from a target device and creates a fork process, in which context the code that follows the macro invocation executes.

FORMAT IOFORK

DESCRIPTION The IOFORK macro calls EXE\$IOFORK to disable timeouts from a target device (by clearing UCB\$V_TIM in UCB\$L_STS) and to create a fork process for a device driver.

When the IOFORK macro is invoked, the following registers must contain the values listed:

Register	Contents
R3	Contents to be placed in R3 of the fork process
R4	Contents to be placed in R4 of the fork process
R5	Address of a UCB that will be used as a fork block for the fork process to be created
00(SP)	Address of caller's caller

EXAMPLE

```
WFIKPCH XA_TIME_OUT,IRP$L_MEDIA(R3)            ;Wait for interrupt
IOFORK                                            ;Device has interrupted; fork
```

The start-I/O routine of a driver initiates an I/O request by invoking the WFIKPCH macro. The WFIKPCH macro sets UCB\$V_INT and UCB\$V_TIM in UCB\$L_STS to record an expected interrupt and enable timeouts from the device, saving the PC of the instruction following IOFORK at UCB\$L_FPC in the driver's fork block. When the device interrupts, the driver's interrupt service routine clears UCB\$V_INT and issues the instruction JSB @UCB\$L_FPC(R5), transferring control to the IOFORK macro invocation.

The IOFORK macro clears the UCB\$V_TIM bit, creates a fork block, inserts it in the appropriate fork queue, requests a software interrupt at that fork IPL from the local processor, and returns control to the driver's interrupt service routine at the instruction following the JSB. When the processor's IPL drops below the fork level, the fork dispatcher dequeues the fork block, obtains proper synchronization, and resumes execution at the instruction in the driver that follows the IOFORK invocation.

VMS Macros Invoked by Drivers

LOADALT

LOADALT

Lloads a set of Q22-bus alternate map registers.

FORMAT

LOADALT

DESCRIPTION

The LOADALT macro calls IOC\$LOADALTMAP to load a set of Q22-bus alternate map registers (registers 496 to 8191). Map registers must already be allocated before the LOADALT macro can be invoked.

When the LOADALT macro is invoked, register R5 must contain the address of the UCB. LOADALT destroys the contents of R0 through R2.

LOADMBA

Loads MASSBUS map registers.

FORMAT LOADMBA

DESCRIPTION

The LOADMBA macro calls IOC\$LOADMBAMAP to load MASSBUS map registers. The driver must own the MASSBUS adapter, and thus the map registers, before it can invoke LOADMBA.

When the LOADMBA macro is invoked, the following registers must contain the following values:

Register	Contents
R4	Address of the MBA's configuration register (MBA\$L_CSR)
R5	Address of UCB

LOADMBA destroys the contents of R0 through R2.

VMS Macros Invoked by Drivers

LOADUBA

LOADUBA

Lloads a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers.

FORMAT

LOADUBA

DESCRIPTION

The LOADUBA macro calls IOC\$LOADUBAMAP to load a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers. Map registers must already be allocated before the LOADUBA macro can be invoked.

When the LOADUBA macro is invoked, register R5 must contain the address of the UCB. LOADUBA destroys the contents of R0 through R2.

LOCK

Achieves synchronized access to a system resource as appropriate to the processing environment.

FORMAT **LOCK** *lockname* [,lockipl] [,savipl] [,condition]
 [,preserve=YES]

PARAMETERS **lockname**
Name of the resource to lock.

[lockipl]

Location containing the IPL at which the resource is synchronized. Although the value of this argument is ignored by the macro, Digital recommends that you specify a **lockipl** value to facilitate debugging.

[savipl]

Location at which to save the current IPL.

[condition]

Indication of a special use of the macro. The only defined **condition** is **NOSETIPL**, which causes the macro to omit setting IPL.

[preserve=YES]

Indication that the macro should preserve R0 across the invocation. If you do not need to retain the contents of R0, specifying **preserve=NO** can enhance system performance.

DESCRIPTION In a *uniprocessing* environment, the LOCK macro sets IPL to the IPL that corresponds to the constant **IPL\$lockname**.

In a *multiprocessing* environment, the LOCK macro performs the following actions:

- Preserves R0 through the macro call (if **preserve=YES** is specified).
- Generates a spin lock index of the form **SPL\$C_lockname** and stores it in R0.
- Calls **SMP\$ACQUIRE** to obtain the specified spin lock. **SMP\$ACQUIRE** indexes into the system spin lock database (a pointer to this database is located at **SMP\$AR_SPNLKVEC**) to obtain the spin lock. Prior to securing the spin lock, **SMP\$ACQUIRE** raises IPL to the IPL associated with the spin lock, determining the appropriate IPL from the spin lock structure (**SPL\$B_IPL**).

VMS Macros Invoked by Drivers

LOCK

In either processing environment, the LOCK macro performs the following tasks:

- Preserves the current IPL at the specified location (if **savipl** is specified)
- Sets the SMP-modified bit in the driver prologue table (DPT\$V_SMPMOD in DPT\$L_FLAGS)

LOCK_SYSTEM_PAGES

Locks a paged code segment in system memory.

FORMAT **LOCK_SYSTEM_PAGES** *[startva],endva [,ipl]*

PARAMETERS *[startva]*
System virtual address in the first page to be locked. If the **startva** argument is omitted, the starting virtual address defaults to the current PC.

endva
System virtual address in the last page to be locked.

[ipl]
IPL at which the locked code segment is to execute. If the **ipl** argument is omitted, the locked code segment executes at the current IPL.

DESCRIPTION

The LOCK_SYSTEM_PAGES macro calls a memory management routine to lock as many pages as necessary into the system working set. The macro accepts a virtual address that indicates the first page to be locked and a virtual address that indicates the last page to be locked. You can also supply the IPL at which the code in the locked pages is to execute.

The LOCK_SYSTEM_PAGES macro executes under the following conditions:

- The LOCK_SYSTEM_PAGES macro should be used only on system virtual addresses.
- All pages requested in a single LOCK_SYSTEM_PAGES macro call must be virtually contiguous. If you must lock discontinuous memory, you must invoke the LOCK_SYSTEM_PAGES macro once for each page or set of contiguous pages.
- You must invoke LOCK_SYSTEM_PAGES at IPL 2 or lower to allow page faulting to occur.
- When the locked code segment is finished, it must invoke the UNLOCK_SYSTEM_PAGES macro to release all previously locked pages. In other words, there must be exactly one UNLOCK_SYSTEM_PAGES macro call per LOCK_SYSTEM_PAGES macro call.
- When it invokes the UNLOCK_SYSTEM_PAGES macro, the code must ensure that the stack is exactly as it was when the LOCK_SYSTEM_PAGES macro was invoked. That is, if the code has pushed anything on the stack, it must remove it before invoking UNLOCK_SYSTEM_PAGES.

VMS Macros Invoked by Drivers

LOCK_SYSTEM_PAGES

- If the **i**pl argument is supplied to the LOCK_SYSTEM_PAGES macro, the locked code segment must invoke the appropriate system synchronization macros (LOCK, FORKLCK, or DEVICELOCK and UNLOCK, FORKUNLOCK or DEVICEUNLOCK) to obtain and release any spin locks required to protect the resources accessed at the elevated IPL.
- If it specified the **i**pl argument to the LOCK_SYSTEM_PAGES macro, the code segment must restore the previous IPL, either explicitly, through the use of the **i**pl argument to the UNLOCK_SYSTEM_PAGES macro, or through the use of one of the system synchronization macros.

EXAMPLE

```
30$:          TSTB      (R0)                ; Fault in page
              LOCK_SYSTEM_PAGES, -
              END=100$                    ; Lock down pages
              LOCK      LOCKNAME=MMG, -    ; Synch with MMG
              SAVIPL=-(SP)                ; Save current IPL
              MOVL      W^MMG$GL_SYSPHD, R3 ; Get system PHD
              .
              .
              UNLOCK   LOCKNAME=MMG, -    ; Unlock MMG
              NEWIPL=(SP)+                ; Restore IPL
              UNLOCK_SYSTEM_PAGES         ; Unlock pages
100$:
```

In this example, the LOCK_SYSTEM_PAGES macro locks all pages between labels 30\$ and 100\$ into the system working set. The UNLOCK_SYSTEM_PAGES macro does the coroutine return to unlock those pages locked by the LOCK_SYSTEM_PAGES macro call.

PURDPR

Purges a UNIBUS adapter buffered data path.

FORMAT PURDPR

DESCRIPTION

The PURDPR macro calls IOC\$PURGDATAP to purge a UNIBUS adapter buffered data path. A driver within an I/O subsystem configuration that does not provide buffered data paths may use the PURDPR macro because the purge operation detects memory parity errors that may have occurred during the transfer. When the PURDPR macro is invoked, R5 must contain the address of the UCB.

When PURDPR returns control to its caller, the following registers contain the following values:

Register	Contents
R0	Status of the purge (success or failure)
R1	Contents of data-path register, provided for the use of the driver's register dumping routine
R2	Address of first map register, provided for the use of the driver's register dumping routine
R3	Address of the CRB

VMS Macros Invoked by Drivers

READ_SYSTIME

READ_SYSTIME

Reads the current system time.

FORMAT **READ_SYSTIME** *dst*

PARAMETER *dst*
Quadword into which the macro inserts the system time.

DESCRIPTION The READ_SYSTIME macro generates the code required to obtain a consistent copy of the system time from EXE\$GQ_SYSTIME.

Use of the READ_SYSTIME macro is subject to the following restrictions:

- IPL must be less than 23.
- The processor must be executing in kernel mode.
- When using the macro within pageable program sections (or within code executing at IPL 2 and below), you must ensure that the pages involved are locked in memory.

EXAMPLE

```
READ_SYSTIME R0
```

The READ_SYSTIME macro inserts the current system time in R0 and R1.

RELALT

Releases a set of Q22-bus alternate map registers allocated to the driver.

FORMAT	RELALT
---------------	---------------

DESCRIPTION	The RELALT macro calls IOC\$RELALTMAP to release a set of Q22-bus alternate map registers (registers 496 to 8191) allocated to the driver. When the RELALT macro is invoked, R5 must contain the address of the UCB. RELALT destroys the contents of R0 through R2.
--------------------	---

VMS Macros Invoked by Drivers

RELCHAN

RELCHAN

Releases all controller data channels allocated to a device.

FORMAT

RELCHAN

DESCRIPTION

The RELCHAN macro calls IOC\$RELCHAN to release all controller data channels allocated to a device. When the RELCHAN macro is invoked, R5 must contain the address of the UCB. RELCHAN destroys the contents of R0 through R2.

RELDPR

Releases a UNIBUS adapter data path register allocated to the driver.

FORMAT	RELDPR
---------------	---------------

DESCRIPTION	The RELDPR macro calls IOC\$RELDATAP to release a UNIBUS adapter buffered data path allocated to the driver.
--------------------	--

When the RELDPR macro is invoked, R5 must contain the address of the UCB. RELDPR destroys the contents of R0 through R2.

VMS Macros Invoked by Drivers

RELMPR

RELMPR

Releases a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers allocated to the driver.

FORMAT

RELMPR

DESCRIPTION

The RELMPR macro calls IOC\$RELMAPREG to release a set of map registers allocated to the driver. When the RELMPR macro is invoked, R5 must contain the address of the UCB. RELMPR destroys the contents of R0 through R2.

RELSCHAN

Releases all secondary channels allocated to the driver.

FORMAT	RELSCHAN
---------------	-----------------

DESCRIPTION	<p>The RELSCHAN macro calls IOC\$RELSCHAN to release all secondary data channels (for example, the MASSBUS adapter's controller data channel) allocated to the driver.</p> <p>When the RELSCHAN macro is invoked, R5 must contain the address of the UCB. RELSCHAN destroys the contents of R0 through R2.</p>
--------------------	--

VMS Macros Invoked by Drivers

REQALT

REQALT

Obtains a set of Q22-bus alternate map registers.

FORMAT

REQALT

DESCRIPTION

The REQALT macro calls IOC\$REQALTMAP to obtain a set of Q22-bus alternate map registers (registers 496 to 8191). When the REQALT macro is invoked, the following registers must contain the following values:

Register	Contents
R5	Address of UCB
00(SP)	Address of caller's caller

The REQALT macro destroys the contents of R0 through R2.

REQCOM

Invokes VMS device-independent I/O postprocessing.

FORMAT REQCOM

DESCRIPTION

The REQCOM macro calls IOC\$REQCOM to complete the processing of an I/O request after the driver has finished its portion of the processing.

When the REQCOM macro is invoked, the following registers must contain the following values:

Register	Contents
R0	First longword of I/O status
R1	Second longword of I/O status
R5	Address of UCB

The REQCOM macro destroys the contents of R0 through R3. All other registers are also destroyed if the action of the macro initiates the processing of a waiting I/O request for the device.

VMS Macros Invoked by Drivers

REQDPR

REQDPR

Requests a UNIBUS adapter buffered data path.

FORMAT

REQDPR

DESCRIPTION

The REQDPR macro calls IOC\$REQDATAP to request a UNIBUS adapter buffered data path.

When the REQDPR macro is invoked, the following registers must contain the following values:

Register	Contents
R5	Address of UCB
00(SP)	Address of caller's caller

The REQDPR macro destroys the contents of R0 through R2.

REQMPR

Obtains a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers.

FORMAT REQMPR

DESCRIPTION The REQMPR macro calls IOC\$REQMAPREG to obtain a set of map registers. When the REQMPR macro is invoked, the following registers must contain the following values:

Register	Contents
R5	Address of UCB
00(SP)	Address of caller's caller

The REQMPR macro destroys the contents of R0 through R2.

VMS Macros Invoked by Drivers

REQPCHAN

REQPCHAN

Obtains a controller's data channel.

FORMAT **REQPCHAN** *[pri]*

PARAMETERS *[pri]*
Priority of request. If the priority is **HIGH**, REQPCHAN calls IOC\$REQPCHANH; otherwise it calls IOC\$REQPCHANL.

DESCRIPTION The REQPCHAN macro calls IOC\$REQPCHANH or IOC\$REQPCHANL, depending on the priority specified, to obtain a controller's data channel.

When the REQPCHAN macro is invoked, the following registers must contain the following values:

Register	Contents
R5	Address of UCB
00(SP)	Address of caller's caller

The REQPCHAN macro returns the address of the device's CSR in R4 and destroys the contents of R0 through R2.

REQSCHAN

Obtains a secondary MASSBUS data channel.

FORMAT **REQSCHAN** *[pri]*

PARAMETER *[pri]*
Priority of request. If the priority is **HIGH**, REQSCHAN calls IOC\$REQSCHANH; otherwise it calls IOC\$REQSCHANL.

DESCRIPTION The REQSCHAN macro calls IOC\$REQSCHANH or IOC\$REQSCHANL, depending on the priority specified, to obtain a secondary MASSBUS data channel.

When the REQSCHAN macro is invoked, the following registers must contain the following values:

Register	Contents
R5	Address of UCB
00(SP)	Address of caller's caller

The REQSCHAN macro returns the address of the device's CSR in R4 and destroys the contents of R0 through R2.

VMS Macros Invoked by Drivers

SAVIPL

SAVIPL

Saves the current IPL of the local processor.

FORMAT **SAVIPL** *[dst=-(SP)]*

PARAMETER *[dst=-(SP)]*
Address of longword in which to save the current IPL.

DESCRIPTION The SAVIPL macro stores the current IPL of the local processor, as recorded in the processor IPL register (PR\$_IPL), in the specified location.

SETIPL

Sets the current IPL of the local processor.

FORMAT **SETIPL** *[ipl=31]* *[environ=MULTIPROCESSOR]*

PARAMETERS ***[ipl=31]***
Level at which to set the current IPL. The default value sets IPL to 31, blocking all interrupts on the local processor.

[environ=MULTIPROCESSOR]
Processing environment in which the SETIPL synchronization macro is to be assembled. If you do not specify **environ**, or if you do specify **environ=MULTIPROCESSOR**, the SETIPL macro generates the following assembly-time warning message, where *xx* is an IPL above IPL 2:

%MACRO-W-GENWARN, Generated WARNING: Raising IPL to #xx provides no multiprocessing synchronization

If you are certain that the purpose of the macro invocation is to block only local processor events, you can disable the warning message by including **environ=UNIPROCESSOR** in the invocation.

DESCRIPTION The SETIPL macro sets the IPL of the local processor by moving the specified **ipl** or IPL 31 into its IPL register (PR\$_IPL).

Note that the SETIPL macro provides full synchronization only in a uniprocessing environment. In a multiprocessor configuration, SETIPL is suitable only for blocking events on the local processor. To provide synchronized access to system resources and devices in a multiprocessing environment, you *must* use the DEVICELOCK/DEVICEUNLOCK, FORKLOCK/FORKUNLOCK, and LOCK/UNLOCK macros.

VMS Macros Invoked by Drivers

SETIPL

EXAMPLE

```
DEVICELOCK - ;Secure device lock
  LOCKADDR=UCB$L_DLCK(R5),- ;(also raises IPL to device lock's IPL)
  SAVIPL=- (SP) ;Save current IPL on stack
SETIPL #IPL$POWER,- ;Raise IPL to 31
  ENVIRON=UNIPROCESSOR ;Avoid assembly-time warning
BBC #UCB$V_POWER, -
  UCB$W_STS(R5),30$ ;If clear, no power failure
;Service power failure
.
.
.
DEVICEUNLOCK - ;Release device lock
  LOCKADDR=UCB$L_DLCK(R5),-
  NEWIPL=(SP)+ ;Restore old IPL from stack
.
.
.
;Branch
30$: ;Start device
.
.
.
WFIKPCH ;Wait for interrupt
```

Here, the `DEVICELOCK` macro achieves synchronized systemwide access to the device registers. The `SETIPL` macro then synchronizes the local processor against its own powerful interrupt event. The code does not need to synchronize systemwide against powerful events, because its interest is truly limited to the local processor.

Note that the `WFIKPCH` macro conditionally releases the device lock and restores the old IPL prior to returning control to the caller's caller.

SOFTINT

Requests a software interrupt from the local processor at a specified IPL.

FORMAT **SOFTINT** *ipl*

PARAMETER *ipl*
IPL at which the software interrupt is being requested.

DESCRIPTION The SOFTINT macro moves the specified **ipl** into the local processor's Software Interrupt Request Register (PR\$_SIRR), thus requesting a software interrupt at that IPL on the processor.

The processor may take either of the following actions:

- If the local processor is executing at an IPL below the level of the requested interrupt, it immediately transfers control to a software interrupt service routine for the appropriate IPL.
- If the local processor is executing at an IPL equal or above the level of the requested interrupt, it does not transfer control to the software interrupt service routine until its IPL drops below the specified **ipl**.

The SOFTINT macro does not provide the capability of requesting a software interrupt from another processor in a VMS multiprocessing environment.

VMS Macros Invoked by Drivers

SPI\$ABORT_COMMAND

SPI\$ABORT_COMMAND

Aborts execution of the outstanding SCSI command on a given connection.

FORMAT SPI\$ABORT_COMMAND

DESCRIPTION The SPI\$ABORT_COMMAND macro aborts the outstanding SCSI command on the connection specified in SCDRP\$L_CDT. The SCSI port driver's abort routine sends the SCSI ABORT command to the target device.

Note: VAXstation 3520/3540 systems do not implement the abort-SCSI-command function.

Inputs to the SPI\$ABORT_COMMAND macro include the following:

Location	Contents
R4	Address of the SPDT
R5	Address of the SCDRP
SCDRP\$L_CDT	Address of the SCDT

The port driver returns SS\$_NORMAL status in R0, and preserves the contents of R3, R4, and R5. The original SPI\$SEND_COMMAND call completes with SS\$_ABORT status.

SPI\$ALLOCATE_COMMAND_BUFFER

Allocates a port command buffer for a SCSI command descriptor block.

FORMAT SPI\$ALLOCATE_COMMAND_BUFFER

DESCRIPTION The SPI\$ALLOCATE_COMMAND_BUFFER macro allocates a port command buffer for a SCSI command descriptor block.

Typically a SCSI class driver requests two additional longwords when specifying the size of the requested buffer, the first for the SCSI status byte and the second for the length of the SCSI command. The port command buffer allows the SCSI port driver to access both the SCSI command descriptor block and the SCSI status byte during the SCSI COMMAND and STATUS phases.

Inputs to the SPI\$ALLOCATE_COMMAND_BUFFER macro include the following:

Location	Contents
R1	Size of requested buffer. This value should include the size of the SCSI command, plus 4 bytes reserved for the SCSI status byte and 4 bytes in which the SCSI class driver places the size of the SCSI command.
R4	Address of the SPDT.
R5	Address of the SCDRP.
SCDRP\$L_CDT	Address of the SCDT.
SCDRP\$W_CMD_MAPREG	Page number of the first port DMA buffer page allocated for the port command buffer.
SCDRP\$W_CMD_NUMREG	Number of port DMA buffer pages allocated for the port DMA buffer.

The port driver returns the following values to the class driver, preserving the contents of R3, R4, and R5:

Location	Contents
R0	SS\$_NORMAL
R1	Size of port command buffer
R2	Address of port command buffer

VMS Macros Invoked by Drivers

SPI\$CONNECT

SPI\$CONNECT

Creates a connection from a class driver to a SCSI device.

FORMAT **SPI\$CONNECT** [*select_callback* [,*select_context*]]

PARAMETERS ***select_callback***

Address of a routine in the class driver that executes in response to asynchronous event notification from the target device. The port driver invokes the selection callback routine at this address, holding the fork lock and no other locks at IPL 8; it passes to the routine the address of the SPDT in R4 and any optional selection context in R5.

If the SCSI class driver does not provide a callback address, no selections are allowed on the connection that is established.

select_context

Longword context value to be passed to selection callback routine. When the port driver invokes the selection callback routine, it passes this value to it in R5. For instance, some class drivers may specify the address of the UCB in this argument (**select_context=R5**) if the selection callback routine needs access to the device unit's UCB. The **select_context** value can help a class driver that supports multiple device units to identify which unit is generating the asynchronous event.

DESCRIPTION

The SPI\$CONNECT macro establishes a connection between the class driver and a SCSI device. It also links a SCSI class driver to the port driver. Before a SCSI class driver can exchange commands and data with a SCSI device, it must invoke SPI\$CONNECT.

In response to the call to SPI\$CONNECT, the port driver allocates and links an SCDT for the connection. It marks the connection state open and initializes default connection information. If the connection already exists, it returns SS\$_DEVALLOC status to the class driver.

Inputs to the SPI\$CONNECT macro include the following:

Location	Contents
R1	SCSI device ID (bits <31:16>) and SCSI port ID (bits <15:0>). Valid SCSI device IDs are integers from 0 to 7; valid SCSI port IDs are integers 0 and 1, corresponding to controller IDs <i>A</i> and <i>B</i> .
R2	SCSI logical unit number (bits <31:16>). Bits <15:0> are reserved. Valid SCSI logical unit numbers are integers from 0 to 7.

VMS Macros Invoked by Drivers

SPI\$CONNECT

The port driver returns the following values to the class driver:

Location	Contents
R0	<p>Port status. The port driver returns one of the following values:</p> <p>SS\$_DEVALLOC Connection already open for this target.</p> <p>SS\$_DEVOFFLINE Port is off line and allows no connections.</p> <p>SS\$_INSFMEM Insufficient memory to allocate SCDT.</p> <p>SS\$_NORMAL Connection formed.</p> <p>SS\$_NOSUCHDEV Port not found.</p>
R2	Address of the SCDT.
R3	<p>Port capability mask. The following bits are defined by the \$SPDTDEF macro (in SYS\$LIBRARY:LIB.MLB):</p> <p>SPDT\$_M_SYNCH Supports synchronous mode.</p> <p>SPDT\$_M_ASYNCH Supports asynchronous mode.</p> <p>SPDT\$_M_MAPPING_REG Supports map registers.</p> <p>SPDT\$_M_BUF_DMA Supports buffered DMA.</p> <p>SPDT\$_M_DIR_DMA Supports direct DMA.</p> <p>SPDT\$_M_AEN Supports asynchronous event notification.</p> <p>SPDT\$_M_LUNS Supports LUNs (logical unit numbers).</p>
R4	Address of the SPDT.

VMS Macros Invoked by Drivers

SPI\$DEALLOCATE_COMMAND_BUFFER

SPI\$DEALLOCATE_COMMAND_BUFFER

Deallocates a port command buffer.

FORMAT SPI\$DEALLOCATE_COMMAND_BUFFER

DESCRIPTION

The SPI\$DEALLOCATE_COMMAND_BUFFER macro deallocates a port command buffer.

Inputs to the SPI\$DEALLOCATE_COMMAND_BUFFER macro include the following:

Location	Contents
R4	Address of the SPDT.
R5	Address of the SCDRP.
SCDRP\$L_CDT	Address of the SCDT.
SCDRP\$W_CMD_MAPREG	Page number of the first port DMA buffer page allocated for the port command buffer.
SCDRP\$W_CMD_NUMREG	Number of the port DMA buffer pages allocated for the port DMA buffer.

The port driver returns SS\$_NORMAL status in R0, and preserves the contents of R3, R4, and R5.

SPI\$DISCONNECT

Breaks a connection between a class driver and a SCSI port.

FORMAT SPI\$DISCONNECT

DESCRIPTION

The SPI\$DISCONNECT macro breaks a connection between a class driver and a SCSI device unit and deallocates the associated SCDT. The connection must not be busy when it is being disconnected.

Normally a connection between a class driver and a SCSI device unit lasts throughout the runtime life of a system. A SCSI class driver should never need to invoke this macro.

Inputs to the SPI\$DISCONNECT macro include the following:

Location	Contents
R1	SCSI device ID (bits <31:16>) and SCSI port ID (bits <15:0>). Valid SCSI device IDs are integers from 0 to 7; valid SCSI port IDs are integers 0 and 1, corresponding to controller IDs <i>A</i> and <i>B</i> .
R2	SCSI logical unit number (bits <15:0>). Valid SCSI logical unit numbers are integers from 0 to 7.
R4	Address of the SPDT.
R5	Address of the SCDT.

The port driver returns SS\$_NORMAL status in R0, and preserves the contents of R3, R4, and R5.

VMS Macros Invoked by Drivers

SPI\$FINISH_COMMAND

SPI\$FINISH_COMMAND

Completes an I/O operation initiated with asynchronous event notification.

FORMAT SPI\$FINISH_COMMAND

DESCRIPTION The SPI\$FINISH_COMMAND macro allows the host acting as a target to send a status byte, return the COMMAND COMPLETE message, and drive the SCSI bus to BUS FREE. The class driver's callback routine should invoke SPI\$FINISH_COMMAND or SPI\$RELEASE_BUS, but not both, before exiting.

The SPI\$FINISH_COMMAND function is a higher-level function that class drivers can use to finish an I/O operation that is executing with asynchronous event notification.

Inputs to the SPI\$FINISH_COMMAND macro include the following:

Location	Contents
R1	Address of the system buffer containing the SCSI status byte
R4	Address of the SPDT

The port driver returns SS\$_NORMAL status in R0, destroys R2, and preserves all other registers.

SPI\$GET_CONNECTION_CHAR

Returns characteristics of an existing connection to a specified buffer.

FORMAT SPI\$GET_CONNECTION_CHAR

DESCRIPTION The SPI\$GET_CONNECTION_CHAR macro returns characteristics of an existing connection to a specified buffer.

The connection characteristics buffer has the following format:

Longword	Contents						
1	Number of longwords in the buffer, <i>not</i> including this longword. The value of this field must be 10.						
2	Connection flags. Bits in this longword are defined as follows: <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>ENA_DISCON. When set, this bit indicates that disconnect and reselection are enabled on this connection.</td> </tr> <tr> <td>1</td> <td>DIS_RETRY. When set, this bit indicates that command retry is disabled on this connection.</td> </tr> </tbody> </table>	Bit	Description	0	ENA_DISCON. When set, this bit indicates that disconnect and reselection are enabled on this connection.	1	DIS_RETRY. When set, this bit indicates that command retry is disabled on this connection.
Bit	Description						
0	ENA_DISCON. When set, this bit indicates that disconnect and reselection are enabled on this connection.						
1	DIS_RETRY. When set, this bit indicates that command retry is disabled on this connection.						
3	Synchronous. When this longword contains 0, the connection supports asynchronous data transfers; when it contains a nonzero value, the connection supports synchronous data transfers.						
4	Transfer period. If the synchronous parameter is nonzero, this field contains the number of 4-nanosecond ticks between a REQ and an ACK. The default is 64 ₁₀ .						
5	REQ-ACK offset. If the synchronous parameter is nonzero, this field contains the maximum number of REQs outstanding before there must be an ACK.						
6	Busy retry count. Maximum number of retries allowed on this connection while waiting for the bus to become free.						
7	Select retry count. Maximum number of retries allowed on this connection while waiting for the port to be selected by the target device.						
8	Arbitration retry count. Maximum number of retries allowed on this connection while waiting for the port to win arbitration of the bus.						

VMS Macros Invoked by Drivers

SPI\$GET_CONNECTION_CHAR

Longword	Contents
9	Command retry count. Maximum number of retries allowed on this connection to successfully send a command to the target device.
10	Phase change timeout. Default timeout value (in seconds) for a target to change the SCSI bus phase or complete a data transfer. This value is also known as the DMA timeout. Upon sending the last command byte, the port driver waits this many seconds for the target to change the bus phase lines and assert REQ (indicating a new phase). Or, if the target enters the DATA IN or DATA OUT phase, the transfer must be completed within this interval. If this value is not specified, the default value is 4 seconds.
11	Disconnect timeout. Default timeout value (in seconds) for a target to reselect the initiator to proceed with a disconnected I/O transfer. If this value is not specified, the default value is 4 seconds.

Inputs to the SPI\$GET_CONNECTION_CHAR macro include the following:

Location	Contents
R2	Address of the connection characteristics buffer.
R4	Address of the SPDT.
R5	Address of the SCDRP.
SCDRP\$L_CDT	Address of the SCDT.

The port driver returns the following values to the class driver, preserving R3, R4, and R5:

Location	Contents
R0	Port status. The port driver returns one of the following values: SS\$_NORMAL Normal, successful completion SS\$_NOSUCHID No connection for this SCSI connection ID
R2	Address of the connection characteristics buffer in which device characteristics are returned.

SPI\$MAP_BUFFER

Makes the process buffer involved in a data transfer available to the port driver.

FORMAT	SPI\$MAP_BUFFER
---------------	------------------------

DESCRIPTION

The SPI\$MAP_BUFFER macro makes the process buffer involved in a data transfer accessible to the port driver. Typically, the I/O buffer is specified in the \$QIO call, is in process space (P0 space), and is mapped by process page-table entries. Because a port driver executes in system context, it cannot access a process's page table.

The means by which the SPI\$MAP_BUFFER macro makes a process buffer available to the port driver depends upon the port hardware. For certain implementations, it allocates a segment of the port's DMA buffer and a set of system page-table entries that double-map the process buffer. In others, it obtains a set of port map registers and loads them with the page-frame numbers of the process buffer pages.

VMS Macros Invoked by Drivers

SPI\$MAP_BUFFER

Inputs to the SPI\$MAP_BUFFER macro include the following:

Location	Contents
R4	Address of the SPDT.
R5	Address of the SCDRP. The class driver must provide values in the following fields:
	SCDRP\$L_BCNT Size in bytes of the buffer to be mapped. The largest single transfer that can be mapped is determined by the port driver in the call to SPI\$CONNECT. The SPI\$CONNECT macro returns this value to the class driver in R1. If the class driver must accomplish transfers larger than this value, it must segment them.
	SCDRP\$W_BOFF Byte offset into the first page of the buffer.
	SCDRP\$L_SVA_USER For direct DMA buffering, system virtual address of the process buffer to map in system space (S0 space)
	SCDRP\$L_SVAPTE System virtual address of the page-table entry that maps the first byte of the user buffer.
	SCDRP\$L_SCSI_FLAGS SCSI mapping flags. If SCDRP\$V_S0BUF is set, SPI\$MAP_BUFFER does not double-map the buffer into system space.
	SCDRP\$W_STS Transfer direction flags. IRP\$V_FUNC must be set for read I/O functions and clear for write I/O functions.

The port driver returns the following values to the class driver, preserving R3, R4, and R5:

Location	Contents
R0	Port status. The port driver returns one of the following values:
	SS\$_NORMAL Normal, successful completion
	SS\$_BADPARAM Bad parameter provided by class driver

VMS Macros Invoked by Drivers

SPI\$MAP_BUFFER

Location	Contents
R5	<p>Address of the SCDRP. The port driver initializes the following fields:</p> <p>SCDRP\$L_SVA_USER System virtual address of the process buffer as mapped in system space (S0 space)</p> <p>SCDRP\$L_SVA_SPTE System virtual address of the system page-table entry that maps the first page of the process buffer in S0 space</p> <p>SCDRP\$W_NUMREG Number of port DMA buffer pages allocated</p> <p>SCDRP\$W_MAPREG Page number of the first port DMA buffer page allocated</p>

VMS Macros Invoked by Drivers

SPI\$RECEIVE_BYTES

SPI\$RECEIVE_BYTES

Receives command, message, and data bytes from a device acting as an initiator on the SCSI bus.

FORMAT SPI\$RECEIVE_BYTES

DESCRIPTION The SPI\$RECEIVE_BYTES macro allows the host to receive information from the device acting as an initiator. A class driver uses SPI\$RECEIVE_BYTES to receive command, message, and data bytes. This macro uses DMA operations for the transfer of large segments of data where appropriate.

Inputs to the SPI\$RECEIVE_BYTES macro include the following:

Location	Contents
R0	Size of the system buffer into which the target returns the requested bytes
R1	Address of the system buffer into which the target device returns the requested bytes
R4	Address of the SPDT

The port driver returns the following values to the class driver, destroying R2, and preserving all other registers:

Location	Contents
R0	Port status. The port driver returns one of the following values: SS\$_NORMAL Normal, successful completion. SS\$_CTRLERR Timeout occurred during the operation.
R1	Actual number of bytes received.

SPI\$RELEASE_BUS

Releases the SCSI bus.

FORMAT SPI\$RELEASE_BUS

DESCRIPTION

The SPI\$RELEASE_BUS macro allows the host acting as a target to release the SCSI bus. The class driver's callback routine should invoke either SPI\$RELEASE_BUS or SPI\$FINISH_COMMAND, but not both, before exiting.

The class driver should use SPI\$RELEASE_BUS instead of SPI\$FINISH_COMMAND if it must explicitly send the SCSI status byte and COMMAND COMPLETE message using SPI\$SEND_BYTES, or if it simply wants to drop off the bus and terminate the thread in certain error conditions.

Inputs to the SPI\$RELEASE_BUS macro include the following:

Location	Contents
R4	Address of the SPDT

The port driver returns SS\$_NORMAL status in R0, destroys R2, and preserves all other registers.

VMS Macros Invoked by Drivers

SPI\$RESET

SPI\$RESET

Resets the SCSI bus and SCSI port hardware.

FORMAT

SPI\$RESET

DESCRIPTION

The SPI\$RESET macro first resets the SCSI bus and then resets the port hardware. A SCSI class driver should rarely invoke this macro; those class drivers that do use it should be aware of the impact of a reset operation on other devices on the same bus. The VMS SCSI port driver logs an error when a class driver invokes the SPI\$RESET macro.

Inputs to the SPI\$RESET macro include the following:

Location	Contents
R4	Address of the SPDT.
R5	Address of the SCDRP.
SCDRP\$L_CDT	Address of the SCDT.

The port driver returns the following value to the class driver, preserving R3, R4, and R5:

Location	Contents
R0	Port status. The port driver returns one of the following values: SS\$_NORMAL Normal, successful completion. SS\$_ABORT Reset aborted before completion.

SPI\$SEND_BYTES

Sends command, message, and data bytes to a device acting as an initiator on the SCSI bus.

FORMAT SPI\$SEND_BYTES

DESCRIPTION

The SPI\$SEND_BYTES macro allows the host to send information to the device acting as an initiator. A class driver uses SPI\$SEND_BYTES to send command, message, and data bytes. This macro uses DMA operations for the transfer of large segments of data where appropriate.

Inputs to the SPI\$SEND_BYTES macro include the following:

Location	Contents
R0	Size of the system buffer that contains the bytes to be sent
R1	Address of the system buffer that contains the bytes to be sent
R4	Address of the SPDT

The port driver returns the following values to the class driver, destroying R2, and preserving all other registers:

Location	Contents
R0	Port status. The port driver returns one of the following values: SS\$_NORMAL Normal, successful completion. SS\$_CTRLERR Timeout occurred during the operation.
R1	Actual number of bytes sent.

VMS Macros Invoked by Drivers

SPI\$SEND_COMMAND

SPI\$SEND_COMMAND

Sends a command to a SCSI device.

FORMAT

SPI\$SEND_COMMAND

DESCRIPTION

The SPI\$SEND_COMMAND macro sends a command to a SCSI device. A class driver invokes this macro, after calling SPI\$ALLOCATE_COMMAND_BUFFER to allocate a port command buffer and formatting a SCSI command descriptor block in it.

The port driver responds to the SPI\$SEND_COMMAND macro call by arbitrating for ownership of the SCSI bus, selecting the target device, sending the SCSI command descriptor block to the target, and waiting for a response. Prior to returning to the class driver, the port driver sends data to or receives data from the target device, obtains command status, processes SCSI message bytes, and transfers the data. When it returns from the SPI\$SEND_COMMAND call, the port driver returns port status and SCSI status to the class driver.

VMS Macros Invoked by Drivers

SPI\$SEND_COMMAND

Inputs to the SPI\$SEND_COMMAND macro include the following:

Location	Contents												
R4	Address of the SPDT.												
R5	Address of the SCDRP. The class driver must provide values in the following fields:												
	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 30%; vertical-align: top;">SCDRP\$L_CMD_PTR</td> <td style="vertical-align: top;">Address of the port command buffer. The first longword of the port command buffer contains the number of bytes in the buffer (not including the count longword). Subsequent bytes contain the SCSI command descriptor block.</td> </tr> <tr> <td style="vertical-align: top;">SCDRP\$L_BCNT</td> <td style="vertical-align: top;">Size in bytes of the mapped process buffer.</td> </tr> <tr> <td style="vertical-align: top;">SCDRP\$W_PAD_BCNT</td> <td style="vertical-align: top;">Number of bytes to make the size of the buffer equal to the data length value required in the command.</td> </tr> <tr> <td style="vertical-align: top;">SCDRP\$L_SVA_USER</td> <td style="vertical-align: top;">System virtual address of the process buffer as mapped in system space (S0 space).</td> </tr> <tr> <td style="vertical-align: top;">SCDRP\$L_STS_PTR</td> <td style="vertical-align: top;">Address of the status longword. The port driver copies the SCSI status byte it receives in the bus STATUS phase into the low-order byte of this buffer.</td> </tr> <tr> <td style="vertical-align: top;">SCDRP\$W_FUNC</td> <td style="vertical-align: top;">Read or write operation.</td> </tr> </table>	SCDRP\$L_CMD_PTR	Address of the port command buffer. The first longword of the port command buffer contains the number of bytes in the buffer (not including the count longword). Subsequent bytes contain the SCSI command descriptor block.	SCDRP\$L_BCNT	Size in bytes of the mapped process buffer.	SCDRP\$W_PAD_BCNT	Number of bytes to make the size of the buffer equal to the data length value required in the command.	SCDRP\$L_SVA_USER	System virtual address of the process buffer as mapped in system space (S0 space).	SCDRP\$L_STS_PTR	Address of the status longword. The port driver copies the SCSI status byte it receives in the bus STATUS phase into the low-order byte of this buffer.	SCDRP\$W_FUNC	Read or write operation.
SCDRP\$L_CMD_PTR	Address of the port command buffer. The first longword of the port command buffer contains the number of bytes in the buffer (not including the count longword). Subsequent bytes contain the SCSI command descriptor block.												
SCDRP\$L_BCNT	Size in bytes of the mapped process buffer.												
SCDRP\$W_PAD_BCNT	Number of bytes to make the size of the buffer equal to the data length value required in the command.												
SCDRP\$L_SVA_USER	System virtual address of the process buffer as mapped in system space (S0 space).												
SCDRP\$L_STS_PTR	Address of the status longword. The port driver copies the SCSI status byte it receives in the bus STATUS phase into the low-order byte of this buffer.												
SCDRP\$W_FUNC	Read or write operation.												
SCDRP\$L_CDT	Address of the SCDT.												

VMS Macros Invoked by Drivers

SPI\$SEND_COMMAND

The port driver returns the following values to the class driver, preserving R3, R4, and R5:

Location	Contents
R0	Port status. The port driver returns one of the following status values:
	SS\$_BADPARAM Bad parameter specified by the class driver.
	SS\$_CTRLERR Controller error or port hardware failure.
	SS\$_DEVACTIVE Command outstanding on this connection.
	SS\$_LINKABORT Connection no longer exists.
	SS\$_NORMAL Normal, successful completion.
	SS\$_TIMEOUT Failed during selection or arbitration.
R5	Address of the SCDRP. The port driver provides information in the following fields:
	SCDRP\$L_STS_PTR Address of the status longword. The port driver copies the SCSI status byte it receives in the bus STATUS phase into the low-order byte of this buffer.
	SCDRP\$L_TRANS_CNT Actual number of bytes sent or received by the port driver during the Data phase.

SPI\$SENSE_PHASE

Returns the current phase of the SCSI bus.

FORMAT SPI\$SENSE_PHASE


DESCRIPTION The SPI\$SENSE_PHASE macro allows the host to read the current SCSI bus phase, and the state of the ATN signal, while using the asynchronous event notification feature.

A class driver must supply the address of the SPDT in R4 as input to the SPI\$SENSE_PHASE macro.

The port driver returns the following values to the class driver, destroying R2, and preserving all other registers:

Location	Contents
R0	SS\$_NORMAL.
R1	SCSI bus phase (and ATN signal). This SCSI-defined longword has the format illustrated in Figure 2-1.

Figure 2-1 SCSI Bus Phase Longword Returned to SPI\$SENSE_PHASE



ZK-1377A-GE

VMS Macros Invoked by Drivers

SPI\$SET_CONNECTION_CHAR

SPI\$SET_CONNECTION_CHAR

Sets characteristics of an existing connection.

FORMAT SPI\$SET_CONNECTION_CHAR

DESCRIPTION The SPI\$SET_CONNECTION_CHAR macro sets characteristics of an existing SCSI connection. Prior to altering the characteristics of a connection, a SCSI class driver should read and examine the current connection characteristics using the SPI\$GET_CONNECTION_CHAR macro.

The class driver specifies the characteristics to be set for the connection in a connection characteristics buffer. The buffer has the following format:

Longword	Contents						
1	Number of longwords in the buffer, <i>not</i> including this longword. The value of this field must be 10.						
2	Connection flags. Bits in this longword are defined as follows: <table border="1"><thead><tr><th>Bit</th><th>Description</th></tr></thead><tbody><tr><td>0</td><td>ENA_DISCON. When set, this bit enables disconnect and reselection on the connection.</td></tr><tr><td>1</td><td>DIS_RETRY. When set, this bit disables command retry on the connection.</td></tr></tbody></table>	Bit	Description	0	ENA_DISCON. When set, this bit enables disconnect and reselection on the connection.	1	DIS_RETRY. When set, this bit disables command retry on the connection.
Bit	Description						
0	ENA_DISCON. When set, this bit enables disconnect and reselection on the connection.						
1	DIS_RETRY. When set, this bit disables command retry on the connection.						
3	Synchronous. When this longword contains 0, the connection uses asynchronous data transfer mode; when it contains a nonzero value, the connection uses synchronous data transfer mode.						
4	Transfer period. If the synchronous parameter is nonzero, this field controls the number of 4-nanosecond ticks between a REQ and an ACK. The default is 64 ₁₀ .						
5	REQ-ACK offset. If the synchronous parameter is nonzero, this field controls the maximum number of REQs outstanding before there must be an ACK.						
6	Busy retry count. Maximum number of retries allowed on this connection while waiting for the port to become free.						
7	Select retry count. Maximum number of retries allowed on this connection while waiting for the port to be selected by the target device.						

VMS Macros Invoked by Drivers

SPI\$SET_CONNECTION_CHAR

Longword	Contents
8	Arbitration retry count. Maximum number of retries allowed on this connection while waiting for the port to win arbitration of the bus.
9	Command retry count. Maximum number of retries allowed on this connection to successfully send a command to the target device.
10	Phase change timeout. Default timeout value (in seconds) for a target to change the SCSI bus phase or complete a data transfer. This value is also known as the DMA timeout. Upon sending the last command byte, the port driver waits this many seconds for the target to change the bus phase lines and assert REQ (indicating a new phase). Or, if the target enters the DATA IN or DATA OUT phase, the transfer must be completed within this interval. If this value is not specified, the default value is 4 seconds.
11	Disconnect timeout. Default timeout value (in seconds) for a target to reselect the initiator to proceed with a disconnected I/O transfer. If this value is not specified, the default value is 4 seconds.

Inputs to the SPI\$SET_CONNECTION_CHAR macro include the following:

Location	Contents
R2	Address of the connection characteristics buffer.
R4	Address of the SPDT.
R5	Address of the SCDRP.
SCDRP\$L_CDT	Address of the SCDT.

The port driver returns the following values to the class driver, preserving R3, R4, and R5:

Location	Contents				
R0	Port status. The port driver returns one of the following values: <table style="margin-left: 2em; border: none;"> <tr> <td>SS\$_NORMAL</td> <td>Normal, successful completion</td> </tr> <tr> <td>SS\$_NOSUCHID</td> <td>No connection for this SCSI connection ID</td> </tr> </table>	SS\$_NORMAL	Normal, successful completion	SS\$_NOSUCHID	No connection for this SCSI connection ID
SS\$_NORMAL	Normal, successful completion				
SS\$_NOSUCHID	No connection for this SCSI connection ID				

VMS Macros Invoked by Drivers

SPI\$SET_PHASE

SPI\$SET_PHASE

Sets the bus to a new phase.

FORMAT SPI\$SET_PHASE


DESCRIPTION

The SPI\$SET_PHASE macro allows the host to set the SCSI bus to a new phase. A class driver uses this macro to drive the phase transitions of the SCSI bus while using the asynchronous event notification feature.

Inputs to the SPI\$SET_PHASE macro include the following:

Location	Contents
R0	New SCSI bus phase. This SCSI-defined longword has the format shown in Figure 2-2.
R4	Address of the SPDT.

Figure 2-2 SCSI Bus Phase Longword Supplied to SPI\$SET_PHASE



ZK-1376A-GE

The port driver returns SS\$_NORMAL status in R0, destroys R2, and preserves all other registers.

SPI\$UNMAP_BUFFER

Releases port mapping resources and deallocates port DMA buffer space, as required to unmap a process buffer.

FORMAT SPI\$UNMAP_BUFFER

DESCRIPTION The SPI\$UNMAP_BUFFER macro releases mapping resources and deallocates port DMA buffer space, as required to unmap a process buffer. Inputs to the SPI\$UNMAP_BUFFER macro include the following:

Location	Contents
R4	Address of the SPDT.
R5	Address of the SCDRP. The class driver must provide values in the following fields:
	SCDRP\$W_NUMREG Number of port DMA buffer pages allocated
	SCDRP\$W_MAPREG Page number of the first port DMA buffer page

The port driver returns the following values to the class driver, preserving R3, R4, and R5:

Location	Contents
R0	SS\$_NORMAL.
R5	Address of the SCDRP. The port driver clears SCDRP\$W_NUMREG and SCDRP\$W_MAPREG.

VMS Macros Invoked by Drivers

TIMEDWAIT

[ublbl]

Label placed at the instruction that performs the processor-specific delay after each execution of the loop of embedded instructions; embedded instructions can pass control here in order to skip the execution of the rest of the embedded instructions in a given execution of the embedded loop.

DESCRIPTION

The TIMEDWAIT macro waits for a period of time for an event or condition to occur. You can specify up to six instructions for this macro to execute in a loop to determine whether the event has occurred.

The TIMEDWAIT macro does not read the processor's clock. The interval it waits is approximate and depends upon the processor and the set of instructions you choose for testing to see if the condition exists.

TIMEDWAIT returns a status code (success or failure) in R0, destroys the contents of R1, and preserves all other registers.

EXAMPLE

```
TIMEDWAIT TIME=#600*1000,-           ;6-second wait loop
INS1=<TSTB RL_CS(R4)>,-               ;Is controller ready?
INS2=<BLSS 15$>,-                     ;If LSS - yes
DONELBL=15$                           ;Label to exit wait loop
BLBC R0,25$                             ;Time expired - exit
```

The unit initialization routine of DLDRIVER issues the TIMEDWAIT macro to wait a maximum of six seconds if another unit is busy on the controller's channel.

EXAMPLE

```
MOVQ      R0,-(SP)           ;Save R0,R1
TIMEWAIT  #3,#RL_CS_M_CRDY,-
          RL_CS(R4),W
MOVQ      (SP)+,R0          ;Restore R0,R1
```

DLDRIVER's unit initialization routine uses the TIMEWAIT macro to wait 30 microseconds for the RL11 controller to be ready before proceeding.

UNLOCK_SYSTEM_PAGES

Terminates a request to lock down a series of system pages.

FORMAT **UNLOCK_SYSTEM_PAGES** *[ipl]*

PARAMETERS *[ipl]*
IPL at which to continue execution.

DESCRIPTION The UNLOCK_SYSTEM_PAGES macro terminates a request to lock down a series of contiguous system pages. In a code segment that uses this locking technique, there must be exactly one UNLOCK_SYSTEM_PAGES macro call per LOCK_SYSTEM_PAGES macro call. When the locked code segment completes, it must invoke the UNLOCK_SYSTEM_PAGES macro to release all previously locked pages.

The UNLOCK_SYSTEM_PAGES macro executes under the following conditions:

- When it invokes the UNLOCK_SYSTEM_PAGES macro, the code must ensure that the stack is exactly as it was when the LOCK_SYSTEM_PAGES macro was invoked. That is, if the code has pushed anything on the stack, it must remove it before invoking UNLOCK_SYSTEM_PAGES.
- If it specified the **ipl** argument to the LOCK_SYSTEM_PAGES macro, the code segment must restore the previous IPL, either explicitly, through the use of the **ipl** argument to the UNLOCK_SYSTEM_PAGES macro, or through the use of one of the system synchronization macros (UNLOCK, FORKUNLOCK or DEVICEUNLOCK). If it lowers IPL, the locked code segment must invoke the appropriate system synchronization macro to release any spin locks that were required to protect the resources accessed at the elevated IPL.

VMS Macros Invoked by Drivers

\$VEC

\$VEC

Defines an entry in a port driver vector table within the context of a \$VECINI macro.

FORMAT **\$VEC** *entry, routine*

PARAMETERS ***entry***

Name of the vector table entry, specified without the PORT_ prefix.

routine

Name of the service routine within the driver that corresponds to the entry point.

DESCRIPTION

A terminal port driver uses the \$VEC macro to validate and generate a vector table entry. A driver need not invoke the \$VEC macro to associate a routine with each entry in the vector table. The \$VECINI macro initializes all unspecified entry points with the address of the driver's null entry point.

To use the \$VEC macro, the driver must include an invocation of the \$TTYMACS definition macro (from SYS\$LIBRARY:LIB.MLB). See the description of the \$VECINI macro for an example of creating a port driver vector table.

\$VECEND

Ends the scope of the \$VECINI macro, thereby completing the definition of a port driver vector table.

FORMAT **\$VECEND** *[end]*

PARAMETER *[end]*
Flag controlling the generation of the end of the vector table. This argument is generally omitted so that the \$VECEND macro can generate the end of the vector table. Otherwise, the \$VECEND macro does not generate the end of the table.

DESCRIPTION A terminal port driver uses the \$VECEND macro to generate the longword of zeros that terminates a port driver vector table initialized by the \$VECINI and \$VEC macros. It also positions the location counter at label **drivername\$VECEND**, as defined by the \$VECINI macro.

To use the \$VECEND macro, the driver must include an invocation of the \$TTYMACS definition macro (from SYS\$LIBRARY:LIB.MLB). See the descriptions of the \$VECINI and \$VEC macros for additional information on creating a port driver vector table.

VMS Macros Invoked by Drivers

\$VECINI

\$VECINI

Begins the definition of a port vector table.

FORMAT **\$VECINI** *drivername*, *nullRoutine* [, *prefix=PORT_*]
 [*,size=_LENGTH*]

PARAMETERS ***drivername***
Prefix (usually two letters) of the driver name (for example, DZ).

nullRoutine
Address of the driver's null entry point, usually specified in the format ***drivername\$NULL***. This address contains an RSB instruction.

 [***prefix=PORT_***]
Prefix to be added to the symbols defined in subsequent invocations of the \$VEC macro.

 [***size***]
Number of bytes allocated for the vector table.

DESCRIPTION A terminal port driver uses the \$VECINI macro to begin the definition of a port vector table and initialize each table entry to point to the driver's null entry point. The \$VECINI macro generates the label ***drivername\$VEC*** at the beginning of the table and ***drivername\$VECEND*** at the end of the table.

 The \$VEC macro defines valid entries within the port driver vector table specified by the invocation of the \$VECINI macro, and the \$VECEND macro ends the table's definition.

 To use the \$VECINI macro, the driver must include an invocation of the \$TTYMACS definition macro (from SYS\$LIBRARY:LIB.MLB).

EXAMPLE

```

$VECINI DZ32,DZ$NULL
$VEC  STARTIO,DZ32$STARTIO      ;Start new output
$VEC  SET_LINE,DZ32$SET_LINE    ;Set new parity/speed
$VEC  XON,DZ32$XON              ;Send XON
$VEC  XOFF,DZ32$XOFF            ;Send XOFF
$VEC  STOP,DZ32$STOP           ;Stop current output
$VEC  ABORT,DZ32$ABORT         ;Abort current output
$VEC  RESUME,DZ32$RESUME       ;Resume stopped output
$VEC  MAINT,DZ32$MAINT        ;Invoke maintenance functions
$VECEND

```

In this example, the \$VECINI macro creates a port driver vector table. The table entries defined by the eight subsequent invocations of the \$VEC

VMS Macros Invoked by Drivers

\$VECINI

macro (PORT_STARTIO, PORT_SET_LINE, and so on) are set up to point to the specified routines in the port driver. The \$VECINI macro initializes any entry point not defined by a \$VEC macro (for instance, PORT_SET_MODEM) with the address of the null entry point, DZ\$NULL.

The \$VECEND macro concludes the definition of the port driver vector table.

VMS Macros Invoked by Drivers

\$VIELD, _VIELD

\$VIELD, _VIELD

Defines symbolic offsets and masks for bit fields.

FORMAT $\left\{ \begin{array}{l} \$VIELD \\ _VIELD \end{array} \right\} \text{ mod ,inibit ,fields}$

PARAMETERS

mod

Module in which this bit field is defined; the prefix portion of the name of the symbol to be defined.

inibit

Bit within the field on which the positions of the bits to be defined are based.

fields

One or more fields of the form `<sym,[size=1],[mask]>`, where these arguments are defined as follows:

Argument	Meaning
sym	String appended to the string "mod\$" to form the name of this bit field.
[size=1]	Size in bits of this bit field. If you specify a value greater than 1, the VIELD macro generates a symbol for the size of the bit field.
[mask]	Character "M" if the VIELD macro is to generate a symbol for the mask of the bit field, blank otherwise.

DESCRIPTION

The \$VIELD and _VIELD macros define bit fields whose names have the form `mod$x_sym` and `mod_x_sym` (where *x* can be V, S, or M and **sym** is a value supplied in the **fields** argument). Because the dollar-sign character (\$) is reserved for use in VMS-defined symbols, use of the _VIELD macro is recommended for non-Digital-supplied device drivers.

See the descriptions of the \$DEFINI and \$EQLST macros for additional information on defining symbols for data structure fields.

EXAMPLE

```

$EQU_LST  XA_K_ , , 0, 1, <-          ;Define CSR bit values
          <fnct1, 2>-
          <fnct2, 4>-
          <fnct3, 8>-
_VIELD    XX_CSR, 0, <-              ;Control/status register
          <GO, , M>, -                ;Start device
          <FNCT, 3, M>, -             ;Function bits
          <XBA, 2, M>, -             ;Extended address bits
          <IE, , M>, -                ;Enable interrupts
          <MAINT>, -                 ;Maintenance bit
          <ATTN>, -                   ;Status from other processors
          >
    
```

This code excerpt produces the following symbols:

```

.
.
.
XX_CSR_M_FNCT          = 0000000E
XX_CSR_M_GO           = 00000001
XX_CSR_M_IE           = 00000040
XX_CSR_M_XBA          = 00000030
XX_CSR_S_FNCT         = 00000003
XX_CSR_S_XBA          = 00000002
XX_CSR_V_FNCT         = 00000001
XX_CSR_V_GO           = 00000000
XX_CSR_V_IE           = 00000006
XX_CSR_V_MAINT        = 00000007
XX_CSR_V_XBA          = 00000004
    
```

VMS Macros Invoked by Drivers

WFIKPCH, WFIRLCH

WFIKPCH, WFIRLCH

Suspends a driver fork thread and folds its context into a fork block in anticipation of a device interrupt or timeout. When WFIKPCH is invoked, the fork thread keeps ownership of the controller channel while waiting; when WFIRLCH is invoked, the fork thread releases ownership of the controller channel.

FORMAT { WFIKPCH }
 { WFIRLCH } *except* [, *time=65536*]

PARAMETERS *except*
Name of a device timeout handling routine; the address of this routine must be within 65,536 bytes of the address at which the WFIKPCH macro is invoked.

[*time=65536*]
Timeout interval, expressed as the number of seconds to wait for an interrupt before a device timeout is considered to exist. A value equal to or greater than 2 is required because the timeout detection mechanism is accurate only to within one second.

DESCRIPTION The WFIKPCH and WFIRLCH macros push **time** on the stack and call IOC\$WFIKPCH and IOC\$WFIRLCH, respectively. After the JSB instruction that makes the routine call, either of these macros constructs a word that contains the relative offset to the timeout handling routine specified in **except**. Because these routines compute and store the address of the following instruction in the fork block at UCB\$L_FPC, the software timer interrupt service routine can determine the routine's location and call it if the device times out before it can deliver an interrupt.

IOC\$WFIKPCH and IOC\$WFIRLCH assume that, prior to the invocation of the macro, a DEVICELock macro has been issued—both to synchronize with other device activity and to leave the IPL of the previous code thread on the top of the stack. Upon storing the context of and suspending the current code thread, IOC\$WFIKPCH and IOC\$WFIRLCH return control to their caller's caller at the stored IPL.

VMS Macros Invoked by Drivers

WFIKPCH, WFIRLCH

When the WFIKPCH or WFIRLCH macro is invoked, the following locations must contain the values listed:

Location	Contents
R5	Address of UCB
00(SP)	IPL at which control is passed to the caller's caller
04(SP)	Address (in the caller's caller) at which to return control

The suspended code thread is resumed by the occurrence of an interrupt signaling the successful completion of a device operation. When an interrupt occurs, control returns to the instruction following the macro. If a device timeout occurs before an interrupt can be posted, the timeout handling routine specified in **except** is called. In both instances, subsequent code can assume that only R3 and R4 have been preserved across the suspension.

See the descriptions of the DEVICELock, IOFORK, and SETIPL macros for examples of the use of the WFIKPCH macro.

3

Operating System Routines

This chapter describes the VMS operating system routines that are used by device drivers and employs the following conventions:

- Most routines reside in modules within the [SYS] facility of VMS. A routine description provides a facility name (in brackets) only if the module is not located in the [SYS] facility.
- Many routines are not directly called by device drivers. Rather, VMS supplies macros that drivers invoke to accomplish the routine call. The description of a routine that has such a macro interface lists the name of the associated macro. Chapter 2 describes how a driver can use these macros.
- System routines generally return a status value in R0 (for instance, SS\$_NORMAL). The low-order bit of this value indicates successful (1) or unsuccessful (0) completion of the routine. Additional information on returned status values appears in the *VMS System Services Reference Manual* and the *VMS System Messages and Recovery Procedures Reference Manual*.
- If a register is not used to transfer output or is not explicitly indicated as destroyed, a driver can assume that its contents are preserved.

Operating System Routines

COM\$DELATTNAST

COM\$DELATTNAST

Delivers all attention ASTs linked in the specified list.

module

COMDRVSUB

input

Location	Contents
R4	Address of listhead of AST control blocks
R5	Address of UCB

output

Location	Contents
Specified listhead	Empty
R0 through R11	Preserved

synchronization

COM\$DELATTNAST executes and exits at the caller's IPL, and acquires no spin locks.

DESCRIPTION

COM\$DELATTNAST removes all AST control blocks (ACBs) from the specified list. Using each ACB as a fork block, it schedules a fork process at IPL\$_QUEUEAST to queue the AST to its target process. COM\$DELATTNAST dequeues each ACB from the head of the list, thus removing them in the reverse order of their declaration by COM\$SETATTNAST. Note that in certain circumstances attention ASTs can be delivered to a user process before the delivery of I/O completion ASTs previously posted by the driver.

COM\$DRVDEALMEM

Deallocates system dynamic memory.

module

COMDRVSUB

input

Location	Contents
R0	Address of block to be deallocated
IRP\$W_SIZE	Size of block in bytes (must be at least 24 bytes long)

output

Location	Contents
R0 through R11	Preserved

synchronization

Drivers can call COM\$DRVDEALMEM from any IPL. COM\$DRVDEALMEM executes at the caller's IPL and returns control at that IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

COM\$DRVDEALMEM calls EXE\$DEANONPAGED to deallocate the buffer specified by R0. If COM\$DRVDEALMEM cannot deallocate memory at the caller's IPL, it transforms the block being deallocated into a fork block and queues the block in the fork queue. The code that executes in the fork process then jumps to EXE\$DEANONPAGED.

If the buffer to be deallocated is less than FKB\$C_LENGTH in size, or its address is not aligned on a 16-byte boundary, COM\$DRVDEALMEM issues a BADDALRQSZ bugcheck.

Operating System Routines

COM\$FLUSHATTNS

COM\$FLUSHATTNS

Flushes an attention AST list.

module COMDRVSUB

input

Location	Contents
R4	Address of PCB
R5	Address of UCB
R6	Number of the assigned I/O channel
R7	Address of listhead of AST control blocks
UCB\$L_DLCK	Address of device lock
PCB\$L_PID	Process ID
PCB\$W_ASTCNT	ASTs remaining in quota

output

Location	Contents
R0	SS\$_NORMAL
R1, R2, R7	Destroyed
PCB\$W_ASTCNT	Incremented by the number of AST control blocks that are flushed
Specified listhead	Updated

synchronization

COM\$FLUSHATTNS raises IPL to device IPL, acquiring the corresponding device lock. Before returning control to its caller at the caller's IPL, COM\$FLUSHATTNS releases the device lock. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

A driver's cancel-I/O routine calls COM\$FLUSHATTNS to flush an attention AST list. A driver FDT routine calls COM\$FLUSHATTNS to service a \$QIO request that specifies a set-attention-AST function and a value of 0 in the **p1** argument.

COM\$FLUSHATTNS locates all AST control blocks whose channel number and PID match those supplied as input to the routine. It removes them from the specified list, deallocates them, and returns control to its caller.

COM\$POST, COM\$POST_NOCNT

Initiates device-independent postprocessing of an I/O request independent of the status of the device unit.

module

COMDRVSUB

input

Location	Contents
R3	Address of IRP
R5	Address of UCB (COM\$POST only)
IRP\$L_MEDIA	Data to be copied to the I/O status block
IRP\$L_MEDIA+4	Data to be copied to the I/O status block

output

Location	Contents
R0	Destroyed
UCB\$L_OPCNT	Incremented (COM\$POST only)

synchronization

Drivers call COM\$POST and COM\$POST_NOCNT at or above fork IPL. These routines execute at their callers' IPL and return control at that IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

A driver fork process calls COM\$POST or COM\$POST_NOCNT after it has completed device-dependent I/O processing for an I/O request initiated by EXE\$ALTQUEPKT. Because COM\$POST_NOCNT, unlike COM\$POST, does not increment the unit's operations count (UCB\$L_OPCNT), a driver uses COM\$POST_NOCNT to initiate completion processing for an I/O request when the associated UCB is not available.

COM\$POST and COM\$POST_NOCNT insert the IRP into the systemwide I/O postprocessing queue, request an IPL_IOPOST software interrupt, and return control to the caller. Unlike IOC\$REQCOM, these routines do not attempt to dequeue any IRP waiting for the device or change the busy status of the device.

Operating System Routines

COM\$SETATTNAST

COM\$SETATTNAST

Enables or disables attention ASTs.

module

COMDRVSUB

input

Location	Contents
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
R7	Address of listhead of AST control blocks
AP	Address of \$QIO system service argument list
IRP\$W_CHAN	I/O request channel index number
UCB\$L_DLCK	Address of device lock
PCB\$W_ASTCNT	Number of ASTs remaining in process quota
PCB\$L_PID	Process ID
00(AP)	Address of process's AST routine
04(AP)	AST parameter
08(AP)	Access mode for AST

output

Location	Contents
R0	SS\$_NORMAL, SS\$_EXQUOTA, or SS\$_INSFMEM
R1 and R2	Destroyed
R3	Address of IRP
R5	Address of UCB
R6, R7, R8	Destroyed
PCB\$W_ASTCNT	Decrementing
Specified listhead	Updated

synchronization

COM\$SETATTNAST raises IPL to device IPL, acquiring the corresponding device lock. It returns control to its caller at the caller's IPL.

DESCRIPTION

A driver FDT routine calls COM\$SETATTNAST to service a \$QIO request that specifies a set-attention-AST function.

If the **p1** argument of the request contains a zero, COM\$SETATTNAST transfers control to COM\$FLUSHATTNS, which disables all ASTs indicated by the PID and I/O channel number (IRP\$W_CHAN). COM\$FLUSHATTNS searches through the AST control block (ACB)

Operating System Routines

COM\$SETATTNAST

list, extracts each identified ACB, deallocates, and returns to the caller of COM\$SETATTNAST.

If the **p1** argument of the request contains the address of an AST routine, COM\$SETATTNAST decrements PCB\$W_ASTCNT and allocates an expanded AST control block (ACB) that contains the following information:

- Spin lock index SPL\$C_QUEUEAST
- Address of the AST routine (as specified in **p1**)
- AST parameter (as specified in **p2**)
- Access mode (as specified in **p3** and maximized against the current process's access mode and bit ACB\$V_QUOTA set to indicate a process-requested AST)
- Number of the assigned I/O channel
- PID of the requesting process

COM\$SETATTNAST links the ACB to the start of the specified linked list of ACBs located in a UCB extension area. (See Section 1.17 for information on defining an extension to a UCB.) COM\$DELATTNAST can later use the expanded ACB to fork to IPL\$_QUEUEAST, at which IPL it reformats the block into a standard ACB.

If the process exceeds buffered I/O or AST quotas, or if there is no memory available to allocate the expanded ACB, COM\$SETATTNAST restores PCB\$W_ASTCNT to its original value and transfers control to EXE\$ABORTIO with error status.

Operating System Routines

ERL\$DEVICERR, ERL\$DEVICTMO, ERL\$DEVICEATTN

ERL\$DEVICERR, ERL\$DEVICTMO, ERL\$DEVICEATTN

Allocate an error message buffer and record in it information concerning the error.

module

ERRORLOG

input

Location	Contents
R5	Address of UCB
DDT\$W_ERRORBUF	Size of error message buffer in bytes
UCB\$L_DEVCHAR	Bit DEV\$V_ELG set
UCB\$W_FUNC	Bit IO\$V_INHERLOG clear
UCB\$L_IRP	Address of IRP currently being processed (ERL\$DEVICERR and ERL\$DEVICTMO only)
UCB\$L_ORB	ORB address

output

Location	Contents
UCB\$W_ERRCNT	Incremented
UCB\$L_EMB	Address of error message buffer
UCB\$L_STS	UCB\$V_ERLOGIP set
R0 through R11	Preserved

synchronization

A driver calls ERL\$DEVICERR, ERL\$DEVICTMO, or ERL\$DEVICEATTN, at or above fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment. These routines return control to the caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

ERL\$DEVICERR and ERL\$DEVICTMO log an error associated with a particular I/O request. ERL\$DEVICEATTN logs an error that is not associated with an I/O request. Each of these routines performs the following steps:

- Increments UCB\$W_ERRCNT to record a device error. If the error-log-in-progress bit (UCB\$V_ERLOGIP in UCB\$L_STS) is set, the routine returns control to its caller.
- Allocates from the current error log allocation buffer an error message buffer of the length specified in the device's DDT (in argument **erlgbf** to the DDTAB macro). This allocation is performed at IPL\$EMB holding the EMB spin lock.

Operating System Routines

ERL\$DEVICERR, ERL\$DEVICTMO, ERL\$DEVICEATTN

- Initializes the buffer with the current system time, error log sequence number, and error type code. These routines use the following error type codes:

ERL\$DEVICERR	Device error (EMB\$C_DE)
ERL\$DEVICTMO	Device timeout (EMB\$C_DT)
ERL\$DEVICEATTN	Device attention (EMB\$C_DA)

- Places the address of the error message buffer in UCB\$L_EMB.
- Sets UCB\$V_ERLOGIP in UCB\$L_STS.
- Loads fields from the UCB, the IRP, and the DDB into the buffer, including the following:

UCB\$B_DEVCLASS	Device class
UCB\$B_DEVTYPE	Device type
IRP\$L_PID	Process ID of the process originating the I/O request (ERL\$_DEVICERR and ERL\$_DEVICTMO)
IRP\$W_BOFF	Transfer parameter (ERL\$DEVICERR and ERL\$DEVICTMO)
IRP\$W_BCNT	Transfer parameter (ERL\$DEVICERR and ERL\$DEVICTMO)
UCB\$L_MEDIA	Disk size
UCB\$W_UNIT	Unit number
UCB\$W_ERRCNT	Count of device errors
UCB\$L_OPCNT	Count of completed operations
ORB\$L_OWNER	UIC of volume owner
UCB\$L_DEVCHAR	Device characteristics
UCB\$B_SLAVE	Slave unit number
IRP\$W_FUNC	I/O function value (ERL\$DEVICERR and ERL\$DEVICTMO)
DDB\$T_NAME	Device name (concatenated with cluster node name if appropriate)

- Loads into R0 the address of the location in the buffer in which the contents of the device registers are to be stored.
- Calls the driver's register dumping routine, the address of which is specified in the **regdmp** argument to the DDTAB macro.

Note that a driver must define the local disk UCB extension or local tape UCB extension, as described in Section 1.17, to use these error logging routines.

Operating System Routines

EXE\$ABORTIO

EXE\$ABORTIO

Completes the servicing of an I/O request without returning status to the I/O status block specified in the request.

module

SYSQIOREQ

input

Location	Contents
R0	First longword of status for the I/O status block
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
IRP\$L_IOSB	Address of I/O status block
IRP\$B_RMOD	ACB\$V_QUOTA set indicates process-specified AST pending
PCB\$W_ASTCNT	Count of available AST queue entries

output

Location	Contents
IRP\$L_IOSB	Zero
IRP\$B_RMOD	ACB\$V_QUOTA clear
PCB\$W_ASTCNT	Incremented if ACB\$V_QUOTA was set

synchronization

EXE\$ABORTIO executes at its caller's IPL and raises to fork IPL, acquiring the associated fork lock in a VMS multiprocessing environment. As a result, its caller cannot be executing above fork IPL. A driver usually transfers control to EXE\$ABORTIO at IPL\$_ASTDEL.

EXE\$ABORTIO exits at normal process IPL (IPL 0).

DESCRIPTION

EXE\$ABORTIO performs the following actions:

- 1 Clears IRP\$L_IOSB so that no status is returned by I/O postprocessing
- 2 Clears ACB\$V_QUOTA in IRP\$B_RMOD to prevent the delivery of any AST to the process specified in the I/O request
- 3 Updates the count of available AST entries at PCB\$W_ASTCNT, if necessary
- 4 Inserts the IRP in the local processor's I/O postprocessing queue
- 5 If the queue is empty, requests a software interrupt from the local processor at IPL\$_IOPOST

This interrupt causes I/O postprocessing to occur before the remaining instructions in EXE\$ABORTIO are executed.

Operating System Routines

EXE\$ABORTIO

When all I/O postprocessing has been completed, EXE\$ABORTIO regains control and completes the I/O operation as follows:

- Lowers IPL to zero
- Issues the RET instruction that restores the original access mode of the caller of the \$QIO system service and returns control to the system service dispatcher

EXE\$ABORTIO returns in R0 the final status code saved when the exit routine was called. Any ASTs specified when the I/O request was issued will not be delivered, and any event flags requested will not be set.

Operating System Routines
EXE\$ALLOCBUF, EXE\$ALLOCIRP

EXE\$ALLOCBUF, EXE\$ALLOCIRP

Allocates a buffer from nonpaged pool for a buffered-I/O operation.

module MEMORYALC

input

Location	Contents
R1	Size of requested buffer in bytes (EXE\$ALLOCBUF only). This value should include the 12 bytes required to store header information.
PCB\$L_STS	PCB\$V_SSRWAIT clear if the process should wait if no memory is available for requested buffer; set if resource wait mode is disabled.

output

Location	Contents
R0	SS\$_NORMAL or SS\$_INSFMEM.
R1	Size of requested buffer in bytes (IRP\$_LENGTH for EXE\$ALLOCIRP).
R2	Address of allocated buffer.
R4	See the following discussion.
IRP\$_SIZE (in allocated buffer)	Size of requested buffer in bytes (for EXE\$ALLOCBUF), IRP\$_LENGTH (for EXE\$ALLOCIRP).
IRP\$_B_TYPE (in allocated buffer)	DYN\$_BUFIO (for EXE\$ALLOCBUF), DYN\$_IRP (for EXE\$ALLOCIRP).

synchronization EXE\$ALLOCBUF and EXE\$ALLOCIRP set IPL to IPL\$_ASTDEL. As a result they cannot be called by code executing above IPL\$_ASTDEL. They return control to their callers at the caller's IPL.

DESCRIPTION

EXE\$ALLOCBUF attempts to allocate a buffer of the requested size from nonpaged pool; EXE\$ALLOCIRP attempts to allocate an IRP from nonpaged pool.

If sufficient memory is not available, EXE\$ALLOCBUF and EXE\$ALLOCIRP move the current PCB (CTL\$GL_PCB) into R4 to determine whether the process has resource wait mode enabled. If PCB\$V_SSRWAIT in PCB\$L_STS is clear, these routines place the process in a resource wait state until memory is released.

The caller must check and adjust process quotas (JIB\$L_BYTCNT or JIB\$L_BYTLM, or both) by calling EXE\$DEBIT_BYTCNT or EXE\$DEBIT_BYTCNT_BYTLM. (Note that you can perform this task *and* allocate a buffer of the requested size by using the routines EXE\$DEBIT_

Operating System Routines EXE\$ALLOCBUF, EXE\$ALLOCIRP

BYTCNT_ALO and EXE\$DEBIT_BYTCNT_BYTLM_ALO. These routines invoke EXE\$ALLOCBUF.)

The normal buffered I/O postprocessing routine (IOC\$REQCOM), initiated by the REQCOM macro, readjusts quotas and also deallocates the buffer.

Note that the value returned in R1 and placed at IRP\$W_SIZE in the allocated buffer is the size of the requested buffer. The actual size of the allocated buffer is determined according to the algorithms used by EXE\$ALONONPAGED and the size of the lookaside list packets. The nonpaged pool deallocation routine (EXE\$DEANONPAGED), called in buffered I/O postprocessing, uses similar algorithms when returning memory to nonpaged pool.

Operating System Routines

EXE\$ALONONPAGED

EXE\$ALONONPAGED

Allocates a block of memory from nonpaged pool.

module

MEMORYALC

input

Location	Contents
R1	Size of requested block in bytes

output

Location	Contents
R0	SS\$_NORMAL or SS\$_INSFMEM
R1	Size of the allocated block, which may be larger than the requested size
R2	Address of allocated block

synchronization

EXE\$ALONONPAGED executes at its caller's IPL and at IPL\$_POOL, obtaining the POOL spin lock in a VMS multiprocessing environment. For this reason, it cannot be called by code executing above IPL\$_POOL.

EXE\$ALONONPAGED returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

Depending upon the size of the requested block, EXE\$ALONONPAGED allocates nonpaged pool either from one of the lookaside lists (SRP, IRP, or LRP) or from the variable region of nonpaged dynamic memory.

EXE\$ALONONPAGED does not initialize the header of the allocated block of memory.

EXE\$ALONPAGVAR

Allocates a block of memory from the variable region of nonpaged pool.

module

MEMORYALC

input

Location	Contents
R1	Size of requested block in bytes

output

Location	Contents
R0	SS\$_NORMAL or SS\$_INSFMEM
R1	Size of requested buffer, rounded up to a 16-byte multiple
R2	Address of allocated block

synchronization

EXE\$ALONPAGVAR executes at its caller's IPL and at IPL\$_POOL, holding the POOL spin lock in a VMS multiprocessing environment. For this reason, its caller cannot be executing at an IPL above IPL\$_POOL.

EXE\$ALONPAGVAR returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

EXE\$ALONPAGVAR allocates a block of memory of the requested size from the variable region of nonpaged dynamic memory. Because EXE\$ALONPAGVAR does not attempt to service the request from the lookaside lists, it is suitable for driver fork processes that may afterwards return the allocated block to nonpaged pool in pieces.

EXE\$ALONPAGVAR does not initialize the header of the allocated block of memory.

Operating System Routines

EXE\$ALOPHYCNTG

EXE\$ALOPHYCNTG

Allocates a physically contiguous block of memory.

module

MEMORYALC

input

Location	Contents
R1	Number of physically contiguous pages to allocate

output

Location	Contents
R0	SS\$_NORMAL, SS\$_INSFMEM, or SS\$_INSFSPTS
R2	System virtual address of allocated block, if the allocation succeeds

synchronization

EXE\$ALOPHYCNTG raises IPL to IPL\$_SYNCH and obtains the MMG spin lock. As a result, its caller cannot be executing above IPL\$_SYNCH or hold any spin lock ranked higher than MMG. (For instance, a driver fork process executing at IPL\$_SYNCH holding the IOLOCK8 fork lock can call EXE\$ALOPHYCNTG.)

EXE\$ALOPHYCNTG returns control to its caller at IPL\$_SYNCH. The caller retains any spin lock it held at the time of the call.

DESCRIPTION

EXE\$ALOPHYCNTG allocates a physically contiguous block of memory. You cannot deallocate memory allocated by EXE\$ALOPHYCNTG.

Note that the number of SPT slots available depends on the value of the SPTREQ system parameter.

EXE\$ALTQUEPKT

Delivers an IRP to a driver's alternate start-I/O routine without regard for the status of the device.

module

SYSQIOREQ

input

Location	Contents
R3	Address of IRP
R5	Address of UCB
DDT\$L_ALTSTART	Address of alternate start-I/O routine
UCB\$B_FLCK	Fork lock index
UCB\$L_DDB	Address of unit's DDB
DDB\$L_DDT	Address of DDT

output

Location	Contents
R0 through R5	Destroyed

synchronization

A driver FDT routine calls EXE\$ALTQUEPKT at IPL\$_ASTDEL. EXE\$ALTQUEPKT raises to fork IPL (acquiring any required fork lock) before calling the driver's alternate start-I/O routine. When the alternate start-I/O routine returns control to it, EXE\$ALTQUEPKT returns control to its caller at the caller's IPL (having released its acquisition of the fork lock).

DESCRIPTION

EXE\$ALTQUEPKT calls the driver's alternate start-I/O routine. It does not test whether the unit is busy before making the call.

Operating System Routines

EXE\$CREDIT_BYTCNT, EXE\$CREDIT_BYTCNT_BYTLM

EXE\$CREDIT_BYTCNT, EXE\$CREDIT_BYTCNT_BYTLM

Return credit to a job's buffered-I/O byte count quota and byte limit.

module

EXSUBROUT

input

Location	Contents
R0	Number of bytes to return to the byte count quota (and byte limit)
R4	Address of current PCB
JIB\$B_FLAGS	JIB\$V_BYTCNT_WAITERS set if there are processes waiting for byte count quota from this JIB
JIB\$L_BYTCNT	Job's byte count usage quota
JIB\$L_BYTLM	Job's byte limit (used by EXE\$CREDIT_BYTCNT_BYTLM)

output

Location	Contents
R0	Destroyed
JIB\$L_BYTCNT	Updated
JIB\$L_BYTLM	Updated (by EXE\$CREDIT_BYTCNT_BYTLM)

synchronization

EXE\$CREDIT_BYTCNT and EXE\$CREDIT_BYTCNT_BYTLM raise IPL to IPL\$_SYNCH and obtain the JIB spin lock and the SCHED spin lock (if JIB\$V_BYTCNT_WAITERS is set) in a VMS multiprocessing environment. As a result, their callers cannot be executing above IPL\$_SYNCH or hold any spin lock ranked higher than JIB. (For instance, a driver fork process executing at IPL\$_SYNCH holding the IOLOCK8 fork lock can call these routines. It cannot, however, hold the SCHED spin lock.)

EXE\$CREDIT_BYTCNT and EXE\$CREDIT_BYTCNT_BYTLM return control to their callers at the caller's IPL. Their caller retains any spin locks it held at the time of the call.

DESCRIPTION

EXE\$CREDIT_BYTCNT provides a synchronized method of crediting a job's byte count quota to JIB\$L_BYTCNT. EXE\$CREDIT_BYTCNT_BYTLM also credits a job's byte limit to JIB\$L_BYTLM.

Both routines round the value specified in R0 up to the nearest 16-byte boundary before applying it to the JIB. Both check JIB\$V_BYTCNT_WAITERS to determine if any process is waiting for the return of nonpaged pool quota for this JIB. If a process is waiting, EXE\$CREDIT_BYTCNT calls a system routine that attempts to fill any pending requests.

EXE\$DEANONPAGED, EXE\$DEANONPGDSIZ

Deallocates a block of memory and returns it to nonpaged pool.

module

MEMORYALC

input

Location	Contents
R0	Address of block to be deallocated
R1	Size of block in bytes, if from variable pool (EXE\$DEANONPGDSIZ only)
IRP\$W_SIZE	Size of block in bytes (EXE\$DEANONPAGED only)
IRP\$B_TYPE	Type of block to be deallocated (EXE\$DEANONPAGED only)

output

Location	Contents
R1 and R2	Destroyed

synchronization

EXE\$DEANONPAGED and EXE\$DEANONPGDSIZ execute at the caller's IPL, at IPL\$_SYNCH holding the SCHED spin lock, and at IPL\$_POOL holding the POOL spin lock. As a result, the caller cannot be executing above IPL\$_SYNCH. EXE\$DEANONPAGED and EXE\$DEANONPGDSIZ return control to the caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

EXE\$DEANONPAGED and EXE\$DEANONPGDSIZ deallocate the specified block of memory to nonpaged dynamic memory, returning it to a lookaside list or the variable region of nonpaged pool as appropriate. These routines also report to the scheduler the availability of the deallocated pool.

EXE\$DEANONPAGED issues a BADDALRQSZ bugcheck if the address of the pool to be deallocated is not aligned on a 16-byte boundary.

If enabled by the SYSGEN parameter POOLCHECK, these routines overwrite portions of the deallocated pool with a checksum and a one-byte pattern. This action is helpful when tracking pool corruption problems.

Operating System Routines

EXE\$DEBIT_BYTCNT(_NW), EXE\$DEBIT_BYTCNT_BYTLM(_NW)

EXE\$DEBIT_BYTCNT(_NW), EXE\$DEBIT_BYTCNT_BYTLM(_NW)

Determine whether a job's buffered I/O byte count quota usage permits the process to be granted additional buffered I/O and, if so, adjust the job's byte count quota and byte limit.

module

EXSUBROUT

input

Location	Contents
R1	Number of bytes to be deducted; bit 31, when set, disables the routine's check against IOC\$GW_MAXBUF
R4	Address of current PCB
PCB\$V_STS	PCB\$V_SSRWAIT clear if the process should wait for buffered-I/O byte quota; set if resource wait mode is disabled
IOC\$GW_MAXBUF	Maximum number of buffered I/O bytes the system allows to a single request
JIB\$L_BYTCNT	Job's byte count usage quota
JIB\$L_BYTLM	Job's byte limit (used by EXE\$DEBIT_BYTCNT_BYTLM and EXE\$DEBIT_BYTCNT_BYTLM_NW)

output

Location	Contents
R0	SS\$_NORMAL or SS\$_EXQUOTA
R1	Number of bytes deducted; bit 31 cleared
JIB\$L_BYTCNT	Updated if successful
JIB\$L_BYTLM	Updated if successful (by EXE\$DEBIT_BYTCNT_BYTLM and EXE\$DEBIT_BYTCNT_BYTLM_NW)

synchronization

EXE\$DEBIT_BYTCNT, EXE\$DEBIT_BYTCNT_NW, EXE\$DEBIT_BYTCNT_BYTLM, and EXE\$DEBIT_BYTCNT_BYTLM_NW raise IPL to IPL\$_SYNCH and obtain the JIB spin lock in a VMS multiprocessing environment. As a result, their callers cannot be executing above IPL\$_SYNCH or hold any spin lock ranked higher than JIB. (For instance, a driver fork process executing at IPL\$_SYNCH holding the IOLOCK8 fork lock can call these routines. It cannot, however, hold the SCHED spin lock.)

EXE\$DEBIT_BYTCNT, EXE\$DEBIT_BYTCNT_NW, EXE\$DEBIT_BYTCNT_BYTLM, and EXE\$DEBIT_BYTCNT_BYTLM_NW return control to their callers at the caller's IPL. The caller retains any spin locks it held at the time of the call.

Operating System Routines

EXE\$DEBIT_BYTCNT(_NW), EXE\$DEBIT_BYTCNT_BYTLM(_NW)

DESCRIPTION

EXE\$DEBIT_BYTCNT and EXE\$DEBIT_BYTCNT_NW check whether a process has sufficient quota for a buffer of the specified size and, if so, deduct the corresponding number of bytes from the job's byte count quota. EXE\$DEBIT_BYTCNT_BYTLM and EXE\$DEBIT_BYTCNT_BYTLM_NW also adjust the job's byte limit. All routines round the value specified in R1 up to the nearest 16-byte boundary before applying it to the JIB.

If the process's quota usage is too large, EXE\$DEBIT_BYTCNT and EXE\$DEBIT_BYTCNT_BYTLM place the process into a resource wait state, based on the setting of PCB\$V_SSRWAIT, until sufficient quota is returned to the job. EXE\$DEBIT_BYTCNT_NW and EXE\$DEBIT_BYTCNT_BYTLM_NW do not refer to PCB\$V_SSRWAIT and return an error if the process has exceeded its job's quota. These latter routines never wait for sufficient quota.

If bit 31 in R1 is clear, all routines compare the byte count in R1 against IOC\$GW_MAXBUF, returning an error if the system's maximum buffer allotment to a process is exceeded.

Operating System Routines

EXE\$DEBIT_BYTCNT_ALO, EXE\$DEBIT_BYTCNT_BYTLM_ALO

EXE\$DEBIT_BYTCNT_ALO, EXE\$DEBIT_BYTCNT_BYTLM_ALO

Determine whether a job's buffered I/O byte count quota usage permits the process to be granted additional buffered I/O and, if so, allocates the requested amount of nonpaged pool and adjust the job's byte count quota and byte limit.

module

EXSUBROUT

input

Location	Contents
R1	Number of bytes to be allocated (including the 12 bytes required for the buffer's header) and deducted; bit 31, when set, disables the routine's check against IOC\$GW_MAXBUF
R4	Address of current PCB
PCB\$L_STS	PCB\$V_SSRWAIT clear if the process should wait for buffered-I/O byte quota; set if resource wait mode is disabled
IOC\$GW_MAXBUF	Maximum number of buffered I/O bytes the system allows to a single request
JIB\$L_BYTCNT	Job's byte count usage quota
JIB\$L_BYTLM	Job's byte limit (used by EXE\$DEBIT_BYTCNT_BYTLM_ALO)

output

Location	Contents
R0	SS\$_NORMAL, SS\$_EXQUOTA, or SS\$_INSFMEM
R1	Number of bytes deducted; bit 31 cleared
R2	Address of requested buffer
R3	Destroyed
JIB\$L_BYTCNT	Updated if successful
JIB\$L_BYTLM	Updated if successful (by EXE\$DEBIT_BYTCNT_BYTLM_ALO)
IRP\$W_SIZE (in allocated buffer)	Size of requested buffer in bytes
IRP\$B_TYPE (in allocated buffer)	DYN\$C_BUFIO

Operating System Routines

EXE\$DEBIT_BYTCNT_ALO, EXE\$DEBIT_BYTCNT_BYTLM_ALO

synchronization

EXE\$DEBIT_BYTCNT_ALO and EXE\$DEBIT_BYTCNT_BYTLM_ALO raise IPL to IPL\$_SYNCH and obtain the JIB spin lock in a VMS multiprocessing environment. Their callers cannot be executing above IPL\$_SYNCH or hold any spin lock.

EXE\$DEBIT_BYTCNT_ALO and EXE\$DEBIT_BYTCNT_BYTLM_ALO return control to their callers at IPL\$_ASTDEL.

DESCRIPTION

EXE\$DEBIT_BYTCNT_ALO checks whether a process has sufficient quota for a buffer of the specified size and, if so, allocates the buffer from nonpaged pool and deducts the corresponding number of bytes from the job's byte count quota. EXE\$DEBIT_BYTCNT_BYTLM_ALO also adjusts the job's byte limit. Both routines round the value specified in R1 up to the nearest 16-byte boundary before applying it to the JIB.

If there is insufficient nonpaged pool available for the buffer, these routines return SS\$_INSFMEM status to the caller.

If the process's quota usage is too large, EXE\$DEBIT_BYTCNT_ALO and EXE\$DEBIT_BYTCNT_BYTLM_ALO place the process into a resource wait state, based on the setting of PCB\$V_SSRWAIT, until sufficient quota is returned to the job.

If bit 31 in R1 is clear, these routines compare the byte count in R1 against IOC\$GW_MAXBUF, returning an error if the system's maximum buffer allotment to a process is exceeded.

Operating System Routines

EXE\$FINISHIO, EXE\$FINISHIOC

EXE\$FINISHIO, EXE\$FINISHIOC

Complete the servicing of an I/O request and return status to the I/O status block specified in the request.

module

SYSQIOREQ

input

Location	Contents
R0	First longword of status for the I/O status block
R1	Second longword of status for the I/O status block (EXE\$FINISHIO only)
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB

output

Location	Contents
R0	SS\$_NORMAL
IRP\$L_IOST1	First longword of I/O status
IRP\$L_IOST2	Second longword of I/O status (cleared by EXE\$FINISHIOC)
UCB\$L_OPCNT	Incremented

synchronization

EXE\$FINISHIO and EXE\$FINISHIOC execute at their caller's IPL and raise to fork IPL, acquiring the associated fork lock in a VMS multiprocessing environment. As a result, their callers cannot be executing above fork IPL. A driver usually transfers control to these routines at IPL\$_ASTDEL.

EXE\$FINISHIO and EXE\$FINISHIOC exit at IPL 0 (normal process IPL).

DESCRIPTION

EXE\$FINISHIOC clears the contents of R1. Then, EXE\$FINISHIO or EXE\$FINISHIOC takes the following steps to complete the processing of the I/O request:

- Increases the number of I/O operations completed on the current device in the operation count field of the UCB (UCB\$L_OPCNT). This task is performed at fork IPL, holding the associated fork lock in a VMS multiprocessing environment.
- Stores the contents of R0 and R1 in the IRP.
- Inserts the IRP in the local processor's I/O postprocessing queue.

Operating System Routines

EXE\$FINISHIO, EXE\$FINISHIOC

- If the queue is empty, requests a software interrupt from the local processor at IPL\$_IOPOST.

This interrupt causes postprocessing to occur before the remaining instructions in EXE\$FINISHIO or EXE\$FINISHIOC are executed.

When all I/O postprocessing has been completed, EXE\$FINISHIO or EXE\$FINISHIOC regains control and completes the I/O operation as follows:

- Places status SS\$_NORMAL in R0
- Lowers IPL to zero
- Issues the RET instruction that restores the original access mode of the caller of the \$QIO system service and returns control to the system service dispatcher

The image that issued the \$QIO receives SS\$_NORMAL status in R0, indicating that the I/O request has completed without device-independent error.

Operating System Routines

EXE\$FORK

EXE\$FORK

Creates a fork process on the local processor.

module

FORKCNTRL

macro

FORK

input

Location	Contents
R5	Address of fork block
00(SP)	Return PC of caller
04(SP)	Return PC of caller's caller
FKB\$B_FLCK	Fork lock index or fork IPL

output

Location	Contents
R3	Destroyed
R4	Fork IPL
FKB\$L_FR3 (UCB\$L_FR3)	R3 of caller
FKB\$L_FR4 (UCB\$L_FR4)	R4 of caller
FKB\$L_FPC (UCB\$L_FPC)	00(SP)

synchronization

EXE\$FORK acquires no spin locks and leaves IPL unchanged. It returns control to its caller's caller.

DESCRIPTION

EXE\$FORK saves the contents of R3 and R4 (in FKB\$L_FR3 and FKB\$L_FR4, respectively) in the fork block specified by R5, and pops the return PC value from the top of the stack into FKB\$L_FPC.

If FKB\$B_FLCK contains a fork lock index, EXE\$FORK determines the fork IPL by using this value as an index into the spin lock IPL vector (SMP\$AR_IPLVEC). EXE\$FORK inserts the fork block into the fork queue on the local processor (headed by CPU\$Q_SWIQFL) corresponding to this IPL. If the queue is empty, EXE\$FORK issues a SOFTINT macro, requesting a software interrupt from the local processor at that fork IPL.

Unlike EXE\$IIOFORK, EXE\$FORK does *not* disable timeouts by clearing UCB\$V_TIM in the UCB\$L_STS field.

EXE\$INSERTIRP

Inserts an IRP into the specified queue of IRPs according to the base priority of the process that issued the I/O request.

module

SYSQIOREQ

input

Location	Contents
R2	Address of I/O queue listhead for the device
R3	Address of IRP
IRP\$B_PRI	Base priority of process requesting the I/O

output

Location	Contents
R1	Destroyed
PSL<2> (Z bit)	Set if the entry is first in the queue, cleared if at least one entry is already in the queue
Pending-I/O queue	IRP inserted

synchronization

EXE\$INSERTIRP must be called at fork IPL or higher. In a VMS multiprocessing environment, the caller must also hold the associated fork lock. EXE\$INSERTIRP does not alter IPL or acquire any spin locks. It returns to its caller.

DESCRIPTION

EXE\$INSERTIRP determines the position of the specified IRP in the pending-I/O queue according to two factors:

- Priority of the IRP, which is derived from the requesting process's base priority as stored in the IRP\$B_PRI
- Time that the entry is queued; for each priority, the queue is ordered on a first-in/first-out basis

EXE\$INSERTIRP inserts the IRP into the queue at that position, adjusts the queue links, and sets the Z bit in the PSL to indicate the status of the queue.

Operating System Routines

EXE\$INSIOQ, EXE\$INSIOQC

EXE\$INSIOQ, EXE\$INSIOQC

Insert an IRP in a device's pending-I/O queue and call the driver's start-I/O routine if the device is not busy.

module

SYSQIOREQ

input

Location	Contents
R3	Address of IRP
R5	Address of UCB
UCB\$_FLCK	Fork lock index
UCB\$_STS	UCB\$_BSY set indicates device is busy, clear indicates device is idle
UCB\$_IOQFL	Address of pending-I/O queue listhead
UCB\$_QLEN	Length of pending-I/O queue

output

Location	Contents
R0, R1, R2	Destroyed. Other registers (used by the driver's start-I/O routine) are destroyed if the start-I/O routine is called.
UCB\$_STS	UCB\$_BSY set.
UCB\$_QLEN	Incremented.

synchronization

EXE\$INSIOQ and EXE\$INSIOQC immediately raise to fork IPL and, in a VMS multiprocessing environment, obtain the corresponding fork lock. As a result, their callers must not be executing at an IPL higher than fork IPL or hold a spin lock ranked higher than the fork lock.

EXE\$INSIOQ unconditionally releases ownership of the fork lock before returning control to the caller without possession of the fork lock. If a fork process must retain possession of the fork lock, it should call EXE\$INSIOQC instead.

DESCRIPTION

EXE\$INSIOQ and EXE\$INSIOQC increment UCB\$_QLEN and proceed according to the status of the device (as indicated by UCB\$_BSY in UCB\$_STS) as follows:

- If the device is busy, call EXE\$INSERTIRP to place the IRP on the device's pending-I/O queue.
- If the device is idle, call IOC\$INITIATE to begin device processing of the I/O request immediately. IOC\$INITIATE transfers control to the driver's start-I/O routine.

EXE\$INSTIMQ

Inserts a timer queue element (TQE) into the timer queue.

module

EXSUBROUT

input

Location	Contents
R0, R1	Quadword expiration time for TQE
R5	Address of TQE to be inserted
EXE\$GQ_1ST_TIME	Expiration time of first TQE in timer queue

output

Location	Contents
R2, R3	Destroyed
TQE\$Q_TIME	Quadword expiration time for TQE
EXE\$GQ_1ST_TIME	Updated if TQE is inserted at the head of the timer queue

synchronization

EXE\$INSTIMQ immediately raises to IPL\$_TIMER (IPL\$_SYNCH), obtaining the TIMER spin lock in a VMS multiprocessing environment. As a result, its caller must not be executing above IPL\$_SYNCH or hold any spin locks of a higher rank. (For instance, a driver fork process executing at IPL\$_SYNCH holding the IOLOCK8 fork lock can call EXE\$INSTIMQ.)

EXE\$INSTIMQ returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

EXE\$INSTIMQ inserts the specified TQE into the timer queue according to its expiration time. If the expiration time of the new TQE is sooner than that of the first TQE in the queue, EXE\$INSTIMQ raises IPL to interval clock IPL (obtaining the HWCLK spin lock in a VMS multiprocessing environment), inserts it on the head of the queue, and updates EXE\$GQ_1ST_TIME.

Operating System Routines

EXE\$IOFORK

EXE\$IOFORK

Creates a fork process on the local processor for a device driver, disabling timeouts from the associated device.

module

FORKCNTRL

macro

IOFORK

input

Location	Contents
R5	Address of fork block (usually the UCB)
00(SP)	Return PC of caller
04(SP)	Return PC of caller's caller
FKB\$B_FLCK (UCB\$B_FLCK)	Fork lock index or fork IPL

output

Location	Contents
R3	Destroyed
R4	Fork IPL
UCB\$L_STS	UCB\$V_TIM cleared, disabling device timeouts
FKB\$L_FR3 (UCB\$L_FR3)	R3 of caller
FKB\$L_FR4 (UCB\$L_FR4)	R4 of caller
FKB\$L_FPC (UCB\$L_FPC)	00(SP)

synchronization

EXE\$IOFORK acquires no spin locks and leaves IPL unchanged. It returns control to its caller's caller.

DESCRIPTION

EXE\$IOFORK first disables timeouts from the target device by clearing UCB\$V_TIM in UCB\$L_STS.

It saves the contents of R3 and R4 (in FKB\$L_FR3 and FKB\$L_FR4, respectively) in the fork block specified by R5, and pops the return PC value from the top of the stack into FKB\$L_FPC.

If FKB\$B_FLCK contains a fork lock index, EXE\$IOFORK determines the fork IPL by using this value as an index into the spin lock IPL vector (SMP\$AR_IPLVEC). EXE\$IOFORK inserts the fork block into the fork queue on the local processor (headed by CPU\$Q_SWIQFL) corresponding to this IPL. If the queue is empty, EXE\$IOFORK issues a SOFTINT macro, requesting a software interrupt from the local processor at that fork IPL.

EXE\$MODIFY

Translates a logical read or write function into a physical read or write function, transfers \$QIO system service parameters to the IRP, validates and prepares a user buffer, and proceeds with or aborts a direct-I/O, DMA read/write operation.

module

SYSQIOFDT

input

Location	Contents
R3	Address of IRP.
R4	Address of current PCB.
R5	Address of UCB.
R6	Address of CCB.
R7	Bit number of the I/O function code.
R8	Address of FDT entry for this routine.
00(AP)	Virtual address of buffer (p1).
04(AP)	Number of bytes in transfer (p2). The maximum number of bytes that EXE\$MODIFY can transfer is 65,535 (128 pages minus one byte).
12(AP)	Carriage control byte (p4).
IRP\$W_FUNC	I/O function code.

output

Location	Contents
R0, R1, R2	Destroyed
IRP\$L_IOST2	p4
IRP\$W_STS	IRP\$V_FUNC set, indicating a read function
IRP\$W_FUNC	Logical read or write function code converted to physical function
IRP\$L_SVAPTE	System virtual address of the process page-table entry (PTE) that maps the first page of the buffer
IRP\$W_BOFF	Byte offset to start of transfer in page
IRP\$L_BCNT	Size of transfer in bytes

synchronization

EXE\$MODIFY is called as a driver FDT routine at IPL\$_ASTDEL.

Operating System Routines

EXE\$MODIFY

DESCRIPTION

A driver uses EXE\$MODIFY as an FDT routine when the driver must both read from and write to the user-specified buffer. Because EXE\$MODIFY transfers control to EXE\$QIODRVPKT if its operations are successful or EXE\$ABORTIO if they are not, it must be the last FDT routine called to perform the preprocessing of I/O read/write requests. A driver cannot use EXE\$MODIFY for buffered I/O operations.

EXE\$MODIFY performs the following functions:

- Sets IRP\$V_FUNC in IRP\$W_STS to indicate a read function.
- Writes the **p4** argument of the \$QIO request into IRP\$L_IOST2 (IRP\$B_CARCON).
- Translates logical read and write functions to physical read and write functions.
- Examines the size of the transfer, as specified in the **p2** argument of the \$QIO request, and takes one of the following actions:
 - If the transfer byte count is zero, EXE\$MODIFY transfers control to EXE\$QIODRVPKT to deliver the IRP to the driver's start-I/O routine. The driver start-I/O routine should check for zero-length buffers to avoid mapping them to UNIBUS, Q22-bus, MASSBUS, or VAXBI node space. An attempted mapping can cause a system failure.
 - If the byte count is not zero, EXE\$MODIFY loads the byte count and the starting address of the transfer into R1 and R0, respectively, and calls EXE\$MODIFYLOCK.

EXE\$MODIFYLOCK calls EXE\$MODIFYLOCKR. EXE\$MODIFYLOCKR calls EXE\$READCHKR, which performs the following tasks:

- Moves the transfer byte count into IRP\$L_BCNT. If the byte count is negative, it returns SS\$_BADPARAM status to EXE\$MODIFYLOCKR.
- Determines if the specified buffer is write accessible for a read I/O function, with one of the following results:
 - If the buffer allows write access, EXE\$READCHKR sets IRP\$V_FUNC in IRP\$W_STS and returns SS\$_NORMAL to EXE\$MODIFYLOCKR.
 - If the buffer does not allow write access, EXE\$READCHKR returns SS\$_ACCVIO status to EXE\$MODIFYLOCKR.

Operating System Routines

EXE\$MODIFY

If EXE\$READCHKR succeeds, EXE\$MODIFYLOCKR moves into IRP\$W_BOFF the byte offset to the start of the buffer and calls MMG\$IOLOCK. MMG\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:¹

- If MMG\$IOLOCK succeeds, EXE\$MODIFYLOCKR stores in IRP\$L_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns control to EXE\$MODIFY. EXE\$MODIFY calls EXE\$QIODRVPKT to deliver the IRP to the driver's start-I/O routine.
- If MMG\$IOLOCK fails, it returns SS\$_ACCVIO, SS\$_INSFWSL, or page fault status to EXE\$MODIFYLOCKR.

If either EXE\$READCHKR or MMG\$IOLOCK returns an error status other than a page fault condition, EXE\$MODIFYLOCKR calls EXE\$ABORTIO. In the event of a page fault, EXE\$MODIFYLOCKR adjusts direct I/O count and AST count to the values they held before the I/O request, deallocates the IRP, and restarts the I/O request at the \$QIO system service. This procedure is carried out so that the user process can receive ASTs while it waits for the page fault to complete. Once the page is faulted into memory, the \$QIO system service will resubmit the I/O request.

¹ For read requests, MMG\$IOLOCK performs an optimization for any nonvalid page contained within the buffer. It creates a demand-zero page rather than fault into memory the requested page. However, if the buffer extends to more than one page, this optimization is not possible.

Operating System Routines

EXE\$MODIFYLOCK, EXE\$MODIFYLOCKR

EXE\$MODIFYLOCK, EXE\$MODIFYLOCKR

Validate and prepare a user buffer for a direct-I/O, DMA read/write operation.

module

SYSQIOFDT

input

Location	Contents
R0	Virtual address of buffer
R1	Number of bytes in transfer
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
R6	Address of CCB
R7	Bit number of the I/O function code

output

Location	Contents
R0	SS\$_NORMAL
R1	System virtual address of the process page-table entry (PTE) that maps the first page of the buffer
R2	1, indicating a read function
IRP\$W_STS	IRP\$V_FUNC set, indicating a read function
IRP\$L_SVAPTE	System virtual address of the PTE that maps the first page of the buffer
IRP\$W_BOFF	Byte offset to start of transfer in page
IRP\$L_BCNT	Size of transfer in bytes

synchronization

EXE\$MODIFYLOCK and EXE\$MODIFYLOCKR are called by a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

A driver typically calls EXE\$MODIFYLOCKR instead of EXE\$MODIFYLOCK when it must lock multiple areas into memory for a single I/O request and must regain control, if the request is to be aborted, to unlock these areas. A driver uses either of these routines when it must both read and write to the user-specified buffer and it is not desirable to automatically deliver the IRP to the device unit after the buffer has been successfully locked. A driver cannot use EXE\$MODIFYLOCK or EXE\$MODIFYLOCKR for buffered I/O operations.

EXE\$MODIFYLOCK calls EXE\$MODIFYLOCKR.

Operating System Routines

EXE\$MODIFYLOCK, EXE\$MODIFYLOCKR

EXE\$MODIFYLOCKR calls EXE\$READCHKR, which performs the following tasks:

- Moves the transfer byte count into IRP\$L_BCNT. If the byte count is negative, it returns SS\$_BADPARAM status to EXE\$MODIFYLOCKR.
- Determines if the specified buffer is write accessible for a read I/O function, with one of the following results:
 - If the buffer allows write access, EXE\$READCHKR sets IRP\$V_FUNC in IRP\$W_STS and returns SS\$_NORMAL to EXE\$MODIFYLOCKR.
 - If the buffer does not allow write access, EXE\$READCHKR returns SS\$_ACCVIO status to EXE\$MODIFYLOCKR.

If EXE\$READCHKR succeeds, EXE\$MODIFYLOCKR moves into IRP\$W_BOFF the byte offset to the start of the buffer and calls MMG\$IOLOCK, disabling a paging mechanism used in write-only operations. MMG\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:²

- If MMG\$IOLOCK succeeds, EXE\$MODIFYLOCKR stores in IRP\$L_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns success status to its caller.
- If MMG\$IOLOCK fails, it returns SS\$_ACCVIO, SS\$_INSFWSL, or page fault status to EXE\$MODIFYLOCKR.

If the initial call was to EXE\$MODIFYLOCK and either EXE\$READCHKR or MMG\$IOLOCK returns an error status other than a page fault condition, EXE\$MODIFYLOCKR calls EXE\$ABORTIO. In the event of a page fault, EXE\$MODIFYLOCKR adjusts direct I/O count and AST count to the values they held before the I/O request, deallocates the IRP, and restarts the I/O request at the \$QIO system service. This procedure is carried out so that the user process can receive ASTs while it waits for the page fault to complete. Once the page is faulted into memory, the \$QIO system service will resubmit the I/O request.

If the initial call was to EXE\$MODIFYLOCKR and an error occurs, EXE\$MODIFYLOCKR, by means of a coroutine call, returns control to the driver's FDT routine with status in R0. The driver performs whatever device-specific actions are required to abort the request, preserving the contents of R0 and R1. When the driver issues the RSB instruction, control is returned to EXE\$MODIFYLOCKR. EXE\$MODIFYLOCKR proceeds to abort or resubmit the I/O request.

Otherwise, these routines return success status to their callers.

² For read requests, MMG\$IOLOCK performs an optimization for any nonvalid page contained within the buffer. It creates a demand-zero page rather than fault into memory the requested page. However, if the buffer extends to more than one page, this optimization is not possible.

Operating System Routines

EXE\$MODIFYLOCK, EXE\$MODIFYLOCKR

A driver FDT routine that calls EXE\$MODIFYLOCKR must distinguish between successful and unsuccessful status when it resumes, as shown in the following example:

```
        JSB      G^EXE$MODIFYLOCKR
        BLBS     BUF_LOCK_OK
BUF_LOCK_FAIL:
;
; clean up this $QIO bookkeeping
;
        RSB
BUF_LOCK_OK:
        .
        .
        .
;
;continue processing this I/O request
;
```

EXE\$ONEPARM

Copies a single \$QIO parameter into the IRP and delivers the IRP to a driver's start-I/O routine.

module

SYSQIOFDT

input

Location	Contents
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
R6	Address of CCB
R7	Bit number of the I/O function code
R8	Address of FDT entry for this routine
00(AP)	Address of first function-dependent parameter of the \$QIO request (p1)

output

Location	Contents
IRP\$L_MEDIA	p1

synchronization

EXE\$ONEPARM is called as a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

EXE\$ONEPARM processes an I/O function code that requires only one parameter. This parameter should need no checking: for instance, for read or write accessibility. EXE\$ONEPARM stores the parameter, found at 00(AP), in IRP\$L_MEDIA and transfers control to EXE\$QIODRVPKT to deliver the IRP to the driver.

Operating System Routines

EXE\$QIODRVPKT

EXE\$QIODRVPKT

Delivers an IRP to the driver's start-I/O routine or pending-I/O queue, returns success status in R0, lowers IPL to 0, and returns to the system service dispatcher.

module

SYSQIOREQ

input

Location	Contents
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
UCB\$B_FLCK	Fork lock index or fork IPL
UCB\$L_STS	UCB\$V_BSY set if device is busy, clear if device is idle
UCB\$L_IOQFL	Address of pending-I/O queue listhead
UCB\$W_QLEN	Length of pending-I/O queue

output

UCB\$L_STS	UCB\$V_BSY set
UCB\$W_QLEN	Incremented

synchronization

EXE\$QIODRVPKT is called by a driver's FDT routine at IPL\$_ASTDEL. It exits at IPL 0 (normal process IPL).

DESCRIPTION

EXE\$QIODRVPKT calls EXE\$INSIOQ. EXE\$INSIOQ checks the status of the device and calls either EXE\$INSERTIRP or IOC\$INITIATE to place the IRP in the device's pending-I/O queue or deliver it to the driver's start-I/O routine, respectively.

When EXE\$INSIOQ returns to EXE\$QIODRVPKT at IPL\$_ASTDEL, EXE\$QIODRVPKT returns control to the system service dispatcher in the following steps:

- 1 Loads SS\$_NORMAL into R0
- 2 Lowers IPL to zero
- 3 Issues the RET instruction that restores the original access mode of the caller of the \$QIO system service and returns control to the system service dispatcher

The image that requested the I/O operation receives status SS\$_NORMAL in R0, indicating that the I/O request has completed without device-independent error.

EXE\$QIORETURN

Sets a success status code in R0, lowers IPL to 0, and returns to the system service dispatcher.

module

SYSQIOREQ

input

Location	Contents
R5	Address of UCB
UCB\$B_FLCK	Fork lock index or fork IPL

output

Location	Contents
R0	SS\$_NORMAL

synchronization

EXE\$QIORETURN is typically called by a driver FDT routine at IPL\$_ASTDEL. Its caller cannot be executing above fork IPL or hold any spin locks other than the appropriate fork lock.

EXE\$QIORETURN releases any fork lock held by its caller before it issues the RET instruction.

DESCRIPTION

EXE\$QIORETURN performs the following actions:

- Loads SS\$_NORMAL into R0
- Lowers IPL to zero
- Issues the RET instruction that restores the original access mode of the caller of the \$QIO system service and returns control to the system service dispatcher

The image that requested the I/O operation receives status SS\$_NORMAL in R0, indicating that the I/O request has completed without device-independent error.

Operating System Routines

EXE\$READ

EXE\$READ

Translates a logical read function into a physical read function, transfers \$QIO system service parameters to the IRP, validates and prepares a user buffer, and proceeds with or aborts a direct-I/O, DMA read/write operation.

module

SYSQIOFDT

input

Location	Contents
R3	Address of IRP.
R4	Address of current PCB.
R5	Address of UCB.
R6	Address of CCB.
R7	Bit number of the I/O function code.
R8	Address of FDT entry for this routine.
00(AP)	Virtual address of buffer (p1).
04(AP)	Number of bytes in transfer (p2). The maximum number of bytes that EXE\$READ can transfer is 65,535 (128 pages minus one byte).
12(AP)	Carriage control byte (p4).
IRP\$W_FUNC	I/O function code.

output

Location	Contents
R0, R1, R2	Destroyed
IRP\$B_IOST2	p4
IRP\$W_STS	IRP\$V_FUNC set, indicating a read function
IRP\$W_FUNC	Logical read function code converted to physical
IRP\$L_SVAPTE	System virtual address of the process page-table entry (PTE) that maps the first page of the buffer
IRP\$W_BOFF	Byte offset to start of transfer in page
IRP\$L_BCNT	Size of transfer in bytes

synchronization

EXE\$READ is called as a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

A driver uses EXE\$READ as an FDT routine when the driver must write to the user-specified buffer. Because EXE\$READ transfers control to EXE\$QIODRVPKT if its operations are successful or EXE\$ABORTIO if they are not, it must be the last FDT routine called to perform the preprocessing of read I/O requests. A driver cannot use EXE\$READ for buffered-I/O operations.

EXE\$READ performs the following functions:

- Sets IRP\$V_FUNC in IRP\$W_STS to indicate a read function
- Writes the **p4** argument of the \$QIO request into IRP\$L_IOST2 (IRP\$B_CARCON).
- Translates a logical read function to a physical read function.
- Examines the size of the transfer, as specified in the **p2** argument of the \$QIO request, and takes one of the following actions:
 - If the transfer byte count is zero, EXE\$READ transfers control to EXE\$QIODRVPKT to deliver the IRP to the driver's start-I/O routine. The driver start-I/O routine should check for zero-length buffers to avoid mapping them to UNIBUS, Q22-bus, MASSBUS, or VAXBI node space. An attempted mapping can cause a system failure.
 - If the byte count is not zero, EXE\$READ loads the byte count and the starting address of the transfer into R1 and R0, respectively, and calls EXE\$READLOCK.

EXE\$READLOCK calls EXE\$READLOCKR.

EXE\$READLOCKR calls EXE\$READCHKR, which performs the following tasks:

- Moves the transfer byte count into IRP\$L_BCNT. If the byte count is negative, it returns SS\$_BADPARAM status to EXE\$READLOCKR.
- Determines whether the specified buffer is write accessible for a read I/O function, with one of the following results:
 - If the buffer allows write access, EXE\$READCHKR sets IRP\$V_FUNC in IRP\$W_STS, and returns SS\$_NORMAL to EXE\$READLOCKR.
 - If the buffer does not allow write access, EXE\$READCHKR returns SS\$_ACCVIO status to EXE\$READLOCKR.

If EXE\$READCHKR succeeds, EXE\$READLOCKR moves into IRP\$W_BOFF the byte offset to the start of the buffer and calls MMG\$IOLOCK. MMG\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:³

- If MMG\$IOLOCK succeeds, EXE\$READLOCKR stores in IRP\$L_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns control to EXE\$READ. EXE\$READ transfers control to EXE\$QIODRVPKT to deliver the IRP to the driver's start-I/O routine.
- If MMG\$IOLOCK fails, it returns SS\$_ACCVIO, SS\$_INSFWSL, or page fault status to EXE\$READLOCKR.

³ For read requests, MMG\$IOLOCK performs an optimization for any nonvalid page contained within the buffer. It creates a demand-zero page rather than fault into memory the requested page. However, if the buffer extends to more than one page, this optimization is not possible.

Operating System Routines

EXE\$READ

If either EXE\$READCHKR or MMG\$IOLOCK returns an error status other than a page fault condition, EXE\$READLOCKR transfers control to EXE\$ABORTIO. In the event of a page fault, EXE\$READLOCKR adjusts direct I/O count and AST count to the values they held before the I/O request, deallocates the IRP, and restarts the I/O request at the \$QIO system service. This procedure is carried out so that the user process can receive ASTs while it waits for the page fault to complete. Once the page is faulted into memory, the \$QIO system service will resubmit the I/O request.

EXE\$READCHK, EXE\$READCHKR

Verify that a process has write access to the pages in the buffer specified in a \$QIO request.

module

SYSQIOFDT

input

Location	Contents
R0	Virtual address of buffer
R1	Size of transfer in bytes
R3	Address of IRP

output

Location	Contents
R0	Virtual address of buffer (EXE\$READCHK), SS\$_NORMAL (EXE\$READCHKR), or error status
R1	Size of transfer in bytes
R2	1, indicating a read function
R3	Address of IRP
IRP\$W_STS	IRP\$V_FUNC set, indicating a read function
IRP\$L_BCNT	Size of transfer in bytes

synchronization

EXE\$READCHK and EXE\$READCHKR are called by a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

A driver uses either of these routines to check the write accessibility of a user-specified buffer. A driver typically calls EXE\$READCHKR instead of EXE\$READCHK when it must regain control before the request is aborted in the event the buffer is inaccessible.

EXE\$READCHK calls EXE\$READCHKR.

EXE\$READCHKR performs the following tasks:

- Moves the transfer byte count into IRP\$L_BCNT. If the byte count is negative, it returns SS\$_BADPARAM status to its caller.
- Determines whether the specified buffer is write accessible for a read I/O function, with one of the following results:
 - If the buffer allows write access, EXE\$READCHKR sets IRP\$V_FUNC in IRP\$W_STS and returns SS\$_NORMAL to its caller.
 - If the buffer does not allow write access, EXE\$READCHKR returns SS\$_ACCVIO status to its caller.

Operating System Routines

EXE\$READCHK, EXE\$READCHKR

If the initial call was to EXE\$READCHK, and EXE\$READCHKR returns error status, EXE\$READCHK transfers control to EXE\$ABORTIO to terminate the I/O request. If the initial call was to EXE\$READCHKR, and an error occurs, EXE\$READCHKR returns control to the driver. Otherwise, these routines return success status to their callers.

A driver FDT routine that calls EXE\$READCHKR must distinguish between successful and unsuccessful status when it resumes, as shown in the following example:

```
        JSB      G^EXE$READCHKR
        BLBS    R0,BUF_ACCESS_OK
BUF_ACCESS_FAIL:
;
; clean up this $QIO bookkeeping
;
        JSB      G^EXE$ABORTIO
BUF_ACCESS_OK:
        .
        .
        .
;
;continue processing this I/O request
;
```

EXE\$READLOCK, EXE\$READLOCKR

Validate and prepare a user buffer for a direct-I/O, DMA read operation.

module

SYSQIOFDT

input

Location	Contents
R0	Virtual address of buffer
R1	Number of bytes in transfer
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
R6	Address of CCB
R7	Bit number of the I/O function code

output

Location	Contents
R0	SS\$_NORMAL
R1	System virtual address of the process page-table entry (PTE) that maps the first page of the buffer
R2	1, indicating a read function
IRP\$W_STS	IRP\$V_FUNC set, indicating a read function
IRP\$L_SVAPTE	System virtual address of the PTE that maps the first page of the buffer
IRP\$W_BOFF	Byte offset to start of transfer in page
IRP\$L_BCNT	Size of transfer in bytes

synchronization

EXE\$READLOCK and EXE\$READLOCKR are called by a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

A driver typically calls EXE\$READLOCKR instead of EXE\$READLOCK when it must lock multiple areas into memory for a single I/O request and must regain control, if the request is to be aborted, to unlock these areas. A driver uses either of these routines when it must write to the user-specified buffer and it is not desirable to automatically deliver the IRP to the device unit after the buffer has been successfully locked. A driver cannot use EXE\$READLOCK or EXE\$READLOCKR for buffered I/O operations.

EXE\$READLOCK calls EXE\$READLOCKR.

Operating System Routines

EXE\$READLOCK, EXE\$READLOCKR

EXE\$READLOCKR calls EXE\$READCHKR, which performs the following tasks:

- Moves the transfer byte count into IRP\$L_BCNT. If the byte count is negative, it returns SS\$_BADPARAM status to EXE\$READLOCKR.
- Determines whether the specified buffer is write accessible for a read I/O function, with one of the following results:
 - If the buffer allows write access, EXE\$READCHKR sets IRP\$V_FUNC in IRP\$W_STS and returns SS\$_NORMAL to EXE\$READLOCKR.
 - If the buffer does not allow write access, EXE\$READCHKR returns SS\$_ACCVIO status to EXE\$READLOCKR.

If EXE\$READCHKR succeeds, EXE\$READLOCKR moves into IRP\$W_BOFF the byte offset to the start of the buffer and calls MMG\$IOLOCK. MMG\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:⁴

- If MMG\$IOLOCK succeeds, EXE\$READLOCKR stores in IRP\$L_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns success status to its caller.
- If MMG\$IOLOCK fails, it returns SS\$_ACCVIO, SS\$_INSFWSL, or page fault status to EXE\$READLOCKR.

If the initial call was to EXE\$READLOCK and either EXE\$READCHKR or MMG\$IOLOCK returns an error status other than a page fault condition, EXE\$READLOCKR transfers control to EXE\$ABORTIO. In the event of a page fault, EXE\$READLOCKR adjusts direct I/O count and AST count to the values they held before the I/O request, deallocates the IRP, and restarts the I/O request at the \$QIO system service. This procedure is carried out so that the user process can receive ASTs while it waits for the page fault to complete. Once the page is faulted into memory, the \$QIO system service will resubmit the I/O request.

If the initial call was to EXE\$READLOCKR and an error occurs, EXE\$READLOCKR, by means of a coroutine call, returns control to the driver's FDT routine with status in R0. The driver performs whatever device-specific actions are required to abort the request, preserving the contents of R0 and R1. When the driver issues the RSB instruction, control is returned to EXE\$READLOCKR. EXE\$READLOCKR proceeds to abort or resubmit the I/O request.

Otherwise, these routines return success status to their callers.

⁴ For read requests, MMG\$IOLOCK performs an optimization for any nonvalid page contained within the buffer. It creates a demand-zero page rather than fault into memory the requested page. However, if the buffer extends to more than one page, this optimization is not possible.

Operating System Routines EXE\$READLOCK, EXE\$READLOCKR

A driver FDT routine that calls EXE\$READLOCKR must distinguish between successful and unsuccessful status when it resumes, as shown in the following example:

```
        JSB      G^EXE$READLOCKR
        BLBS     BUF_LOCK_OK
BUF_LOCK_FAIL:
;
; clean up this $QIO bookkeeping
;
        RSB
BUF_LOCK_OK:
        .
        .
        .
;
;continue processing this I/O request
;
```

Operating System Routines

EXE\$RMVTIMQ

EXE\$RMVTIMQ

Removes timer queue elements (TQEs) from the timer queue.

module

EXSUBROUT

input

Location	Contents
R2	Access mode (unused by system subroutine)
R3	Request identification (unused by system subroutine)
R4	Type of TQE entry (TQE\$B_RQTYPE) to remove from queue (TQE\$C_SSNGL) if bit 31 is zero. If bit 31 is set, then R4 contains the address of the TQE.
R5	Process ID (TQE\$L_PID)

output

Location	Contents
R0	If R0=1, then at least one TQE was removed. If R0=0, then no TQE was removed.
R1	Destroyed

synchronization

EXE\$RMVTIMQ immediately raises to IPL\$_TIMER (IPL\$_SYNCH), obtaining the TIMER spin lock in a VMS multiprocessing environment. As a result, its caller must not be executing above IPL\$_SYNCH or hold any spin locks of a higher rank. (For instance, a driver fork process executing at IPL\$_SYNCH holding the IOLOCK8 fork lock can call EXE\$RMVTIMQ and might need the SCHED and HWCLK spin locks, but these impose no additional restrictions on the caller.)

EXE\$RMVTIMQ returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

EXE\$RMVTIMQ removes the specified TQEs from the timer queue. Entries are removed by address, type, access mode, request identification, and process ID. Any entries which meet matching criteria are removed from queue.

If a system subroutine or a wake request TQE is being removed, access mode and request identification need not be supplied. If the TQE address is supplied in R4, no other input need be supplied.

EXE\$SENSEMODE

Copies device-dependent characteristics from the device's UCB into R1, writes a success code into R0, and completes the I/O operation.

module SYSQIOFDT

input

Location	Contents
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
R6	Address of CCB
R7	Bit number of the I/O function code
R8	Address of FDT entry for this routine
00(AP)	Address of first function-dependent parameter of the \$QIO request
UCB\$Q_DEVDEPEND	Device-dependent status

output

Location	Contents
R0	SS\$_NORMAL
R1	Device-dependent status

synchronization EXE\$SENSEMODE is called as a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

A driver uses EXE\$SENSEMODE as an FDT routine to process the sense-device-mode (IO\$_SENSEMODE) and sense-device-characteristics (IO\$_SENSECHAR) I/O functions.

EXE\$SENSEMODE loads the contents of UCB\$Q_DEVDEPEND into R1, places SS\$_NORMAL status into R0, and transfers control to EXE\$FINISHIO to insert the IRP in the systemwide I/O postprocessing queue.

Operating System Routines

EXE\$SETCHAR, EXE\$SETMODE

EXE\$SETCHAR, EXE\$SETMODE

Write device-specific status and control information into the device's UCB and complete the I/O request (EXE\$SETCHAR); or write the information into the IRP and deliver the IRP to the driver's start-I/O routine (EXE\$SETMODE).

module

SYSQIOFDT

input

Location	Contents
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
R6	Address of CCB
R7	Bit number of the I/O function code
R8	Address of FDT entry for this routine
00(AP)	Address of location containing device characteristics quadword (p1)
UCB\$B_DEVCLASS	Device class

output

Location	Contents
R0	SS\$_NORMAL, SS\$_ACCVIO, or SS\$_ILLIOFUNC
UCB\$B_DEVCLASS	Byte 0 of quadword (EXE\$SETCHAR, IO\$_SETCHAR function only)
UCB\$B_DEVTYPE	Byte 1 of quadword (EXE\$SETCHAR, IO\$_SETCHAR function only)
UCB\$W_DEVBUFSIZ	Bytes 2 and 3 of quadword (EXE\$SETCHAR)
UCB\$Q_DEVDEPEND	Bytes 4 through 7 of quadword (EXE\$SETCHAR)
IRP\$L_MEDIA	First longword of device characteristics (EXE\$SETMODE)
IRP\$L_MEDIA+4	Second longword of device characteristics (EXE\$SETMODE)

synchronization

EXE\$SETCHAR or EXE\$SETMODE is called as a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

A driver uses EXE\$SETCHAR or EXE\$SETMODE as an FDT routine to process the set-device-mode (IO\$_SETMODE) and set-device-characteristics (IO\$_SETCHAR) functions. If setting device characteristics requires device activity or synchronization with fork processing, the driver's FDT entry *must* specify EXE\$SETMODE. Otherwise, it can specify EXE\$SETCHAR.

Operating System Routines

EXE\$SETCHAR, EXE\$SETMODE

EXE\$SETCHAR and EXE\$SETMODE examine the current value of UCB\$B_DEVCLASS to determine whether the device permits the specified function. If the device class is disk (DC\$_DISK), the routines place SS\$_ILLIOFUNC status in R0 and transfer control to EXE\$ABORTIO to terminate the request.

EXE\$SETCHAR and EXE\$SETMODE then ensure that the process has read access to the quadword containing the new device characteristics. If it does not, the routines place SS\$_ACCVIO status in R0 and transfer control to EXE\$ABORTIO to terminate the request.

If the request passes these checks, EXE\$SETCHAR and EXE\$SETMODE proceed as follows:

- EXE\$SETCHAR stores the specified characteristics in the UCB. For an IO\$_SETCHAR function, the device type and class fields (UCB\$B_DEVCLASS and UCB\$B_DEVTYPE, respectively) receive the first word of data. For both IO\$_SETCHAR and IO\$_SETMODE functions, EXE\$SETCHAR writes the second word into the default-buffer-size field (UCB\$W_DEVBUSIZ) and the third and fourth words into the device-dependent-characteristics field (UCB\$Q_DEVDEPEND).

Finally, EXE\$SETCHAR stores normal completion status (SS\$_NORMAL) in R0 and transfers control to EXE\$FINISHIO to insert the IRP in the systemwide I/O postprocessing queue.

- EXE\$SETMODE stores the specified quadword of characteristics in IRP\$L_MEDIA, places normal completion status (SS\$_NORMAL) in R0, and transfers control to EXE\$QIODRVPKT to deliver the IRP to the driver's start-I/O routine.

The driver's start-I/O routine copies data from IRP\$L_MEDIA and the following longword into UCB\$W_DEVBUSIZ, UCB\$Q_DEVDEPEND, and, if the I/O function is IO\$_SETCHAR, UCB\$B_DEVCLASS and UCB\$B_DEVTYPE as well.

Operating System Routines

EXE\$SNDEVMSG

EXE\$SNDEVMSG

Builds and sends a device-specific message to the mailbox of a system process, such as the job controller or OPCOM.

module

MBDRIVER

input

Location	Contents
R3	Address of mailbox UCB. (SYS\$AR_JOBCTLMB contains the address of the job controller's mailbox; SYS\$AR_OPRMBX contains the address of OPCOM's mailbox.)
R4	Message type
R5	Address of device UCB
UCB\$W_UNIT	Device unit number
UCB\$L_DDB	Address of device DDB
DDB\$T_NAME and mailbox UCB fields	Device controller name

output

Location	Contents
R0	SS\$_NORMAL, SS\$_MBTOOSML, SS\$_MBFULL, SS\$_INSFMEM, or SS\$_NOPRIV
R1 through R4	Destroyed

synchronization

Because EXE\$SNDEVMSG raises IPL to IPL\$_MAILBOX and obtains the MAILBOX spin lock in a VMS multiprocessing environment, its caller cannot be executing above IPL\$_MAILBOX. EXE\$SNDEVMSG returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

EXE\$SNDEVMSG builds a 32-byte message on the stack that includes the following information:

Bytes	Contents
0 and 1	Low word of R4 (message type)
2 and 3	Device unit number (UCB\$W_UNIT)
4 through 31	Counted string of device controller name, formatted as <i>node\$controller</i> for clusterwide devices

EXE\$SNDEVMSG then calls EXE\$WRTMAILBOX to send the message to a mailbox.

Operating System Routines

EXE\$SNDEVMSG

EXE\$SNDEVMSG can fail for any of the following reasons:

- The message is too large for the mailbox (SS\$_MBTOOSML).
- The message mailbox is full of messages (SS\$_MBFULL).
- The system is unable to allocate memory for the message (SS\$_INSFMEM).
- The caller lacks privilege to write to the mailbox (SS\$_NOPRIV).

Operating System Routines

EXE\$WRITE

EXE\$WRITE

Translates a logical write function into a physical write function, transfers \$QIO system service parameters to the IRP, validates and prepares a user buffer, and proceeds with or aborts a direct-I/O, DMA read/write operation.

module

SYSQIOFDT

input

Location	Contents
R3	Address of IRP.
R4	Address of current PCB.
R5	Address of UCB.
R6	Address of CCB.
R7	Bit number of the I/O function code.
R8	Address of FDT entry for this routine.
00(AP)	Virtual address of buffer (p1).
04(AP)	Number of bytes in transfer (p2). The maximum number of bytes that EXE\$WRITE can transfer is 65,535 (128 pages minus one byte).
12(AP)	Carriage control byte (p4).
IRP\$W_FUNC	I/O function code.

output

Location	Contents
R0, R1, R2	Destroyed
IRP\$L_IOST2	p4
IRP\$W_FUNC	Logical read function code converted to physical
IRP\$W_STS	IRP\$V_FUNC clear, indicating a write function
IRP\$L_SVAPTE	System virtual address of the process page-table entry (PTE) that maps the first page of the buffer
IRP\$W_BOFF	Byte offset to start of transfer in page
IRP\$L_BCNT	Size of transfer in bytes

synchronization

EXE\$WRITE is called as a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

A driver uses EXE\$WRITE as an FDT routine when the driver must read from the user-specified buffer. Because EXE\$WRITE transfers control to EXE\$QIODRVPKT if its operations are successful or EXE\$ABORTIO if they are not, it must be the last FDT routine called to perform the preprocessing of write I/O requests. A driver cannot use EXE\$WRITE for buffered I/O operations.

Operating System Routines

EXE\$WRITE

EXE\$WRITE performs the following functions:

- Writes the **p4** argument of the \$QIO request into IRP\$L_IOST2 (IRP\$B_CARCON).
- Translates a logical write function to a physical write function.
- Examines the size of the transfer, as specified in the **p2** argument of the \$QIO request, and takes one of the following actions:
 - If the transfer byte count is zero, EXE\$WRITE transfers control to EXE\$QIODRVPKT to deliver the IRP to the driver's start-I/O routine. The driver start-I/O routine should check for zero-length buffers to avoid mapping them to UNIBUS, Q22-bus, MASSBUS, or VAXBI node space. An attempted mapping can cause a system failure.
 - If the byte count is not zero, EXE\$READ loads the byte count and the starting address of the transfer into R1 and R0, respectively, and calls EXE\$WRITELOCK.

EXE\$WRITELOCK calls EXE\$WRITELOCKR.

EXE\$WRITELOCKR calls EXE\$WRITECHKR, which performs the following tasks:

- Moves the transfer byte count into IRP\$L_BCNT. If the byte count is negative, it returns SS\$_BADPARAM status to EXE\$WRITELOCKR.
- Determines whether the specified buffer is read accessible for a write I/O function, with one of the following results:
 - If the buffer allows read access, EXE\$WRITECHKR returns SS\$_NORMAL to EXE\$WRITELOCKR.
 - If the buffer does not allow read access, EXE\$WRITECHKR returns SS\$_ACCVIO status to EXE\$WRITELOCKR.

If EXE\$WRITECHKR succeeds, EXE\$WRITELOCKR moves into IRP\$W_BOFF the byte offset to the start of the buffer and calls MMG\$IOLOCK. MMG\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:

- If MMG\$IOLOCK succeeds, EXE\$WRITELOCKR stores in IRP\$L_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns control to EXE\$WRITE. EXE\$WRITE transfers control to EXE\$QIODRVPKT to deliver the IRP to the driver's start-I/O routine.
- If MMG\$IOLOCK fails, it returns SS\$_ACCVIO, SS\$_INSFWSL, or page fault status to EXE\$WRITELOCKR.

If either EXE\$WRITECHKR or MMG\$IOLOCK returns an error status, EXE\$WRITELOCKR transfers control to EXE\$ABORTIO.

Operating System Routines

EXE\$WRITECHK, EXE\$WRITECHKR

EXE\$WRITECHK, EXE\$WRITECHKR

Verify that a process has read access to the pages in the buffer specified in a \$QIO request.

module

SYSQIOFDT

input

Location	Contents
R0	Virtual address of buffer
R1	Size of transfer in bytes
R3	Address of IRP

output

Location	Contents
R0	Virtual address of buffer (EXE\$WRITECHK), SS\$_NORMAL (EXE\$WRITECHKR), or error status
R1	Size of transfer in bytes
R2	0, indicating a write function
IRP\$W_STS	IRP\$V_FUNC clear, indicating a write function
IRP\$L_BCNT	Size of transfer in bytes

synchronization

EXE\$WRITECHK and EXE\$WRITECHKR are called by a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

A driver uses either of these routines to check the read accessibility of a user-specified buffer. A driver typically calls EXE\$WRITECHKR instead of EXE\$WRITECHK when it must regain control before the request is aborted in the event the buffer is inaccessible.

EXE\$WRITECHK calls EXE\$WRITECHKR.

EXE\$WRITECHKR performs the following tasks:

- Moves the transfer byte count into IRP\$L_BCNT. If the byte count is negative, it returns SS\$_BADPARAM status to its caller.
- Determines if the specified buffer is read accessible for a write I/O function, with one of the following results:
 - If the buffer allows read access, EXE\$WRITECHKR returns SS\$_NORMAL to its caller.
 - If the buffer does not allow read access, EXE\$WRITECHKR returns SS\$_ACCVIO status to its caller.

Operating System Routines

EXE\$WRITECHK, EXE\$WRITECHKR

If the initial call was to EXE\$WRITECHK, and EXE\$WRITECHKR returns error status, EXE\$WRITECHK transfers control to EXE\$ABORTIO to terminate the I/O request. If the initial call was to EXE\$WRITECHKR, and an error occurs, EXE\$WRITECHKR returns control to the driver. Otherwise, these routines return success status to their callers.

A driver FDT routine that calls EXE\$WRITECHKR must distinguish between successful and unsuccessful status when it resumes, as shown in the following example:

```
        JSB      G^EXE$WRITECHKR
        BLBS    R0,BUF_ACCESS_OK
BUF_ACCESS_FAIL:
;
; clean up this $QIO bookkeeping
;
        JSB      G^EXE$ABORTIO
BUF_ACCESS_OK:
.
.
.
;
;continue processing this I/O request
;
```


Operating System Routines

EXE\$WRITELOCK, EXE\$WRITELOCKR

EXE\$WRITELOCK, EXE\$WRITELOCKR

Validate and prepare a user buffer for a direct-I/O, DMA write operation.

module

SYSQIOFDT

input

Location	Contents
R0	Virtual address of buffer
R1	Number of bytes in transfer
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
R6	Address of CCB
R7	Bit number of the I/O function code

output

Location	Contents
R0	SS\$_NORMAL
R1	System virtual address of the process page-table entry (PTE) that maps the first page of the buffer
R2	0, indicating a write function
IRP\$_STS	IRP\$_FUNC clear, indicating a write function
IRP\$_SVAPTE	System virtual address of the PTE that maps the first page of the buffer
IRP\$_BOFF	Byte offset to start of transfer in page
IRP\$_BCNT	Size of transfer in bytes

synchronization

EXE\$WRITELOCK and EXE\$WRITELOCKR are called by a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

A driver typically calls EXE\$WRITELOCKR instead of EXE\$WRITELOCK when it must lock multiple areas into memory for a single I/O request and must regain control, if the request is to be aborted, to unlock these areas. A driver uses either of these routines when it must read from the user-specified buffer and it is not desirable to automatically deliver the IRP to the device unit after the buffer has been successfully locked. A driver cannot use EXE\$WRITELOCK or EXE\$WRITELOCKR for buffered I/O operations.

EXE\$WRITELOCK calls EXE\$WRITELOCKR.

Operating System Routines

EXE\$WRITELOCK, EXE\$WRITELOCKR

EXE\$WRITELOCKR calls EXE\$WRITECHKR, which performs the following tasks:

- Moves the transfer byte count into IRP\$L_BCNT. If the byte count is negative, it returns SS\$_BADPARAM status to EXE\$WRITELOCKR.
- Determines if the specified buffer is write accessible for a write I/O function, with one of the following results:
 - If the buffer allows read access, EXE\$WRITECHKR returns SS\$_NORMAL to EXE\$WRITELOCKR.
 - If the buffer does not allow read access, EXE\$WRITECHKR returns SS\$_ACCVIO status to EXE\$WRITELOCKR.

If EXE\$WRITECHKR succeeds, EXE\$WRITELOCKR moves into IRP\$W_BOFF the byte offset to the start of the buffer and calls MMG\$IOLOCK. MMG\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:

- If MMG\$IOLOCK succeeds, EXE\$WRITELOCKR stores in IRP\$L_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns success status to its caller.
- If MMG\$IOLOCK fails, it returns SS\$_ACCVIO, SS\$_INSFWSL, or page fault status to EXE\$WRITELOCKR.

If the initial call was to EXE\$WRITELOCK and either EXE\$WRITECHKR or MMG\$IOLOCK returns an error status other than a page fault condition, EXE\$WRITELOCKR transfers control to EXE\$ABORTIO.

In the event of a page fault, EXE\$WRITELOCKR adjusts direct I/O count and AST count to the values they held before the I/O request, deallocates the IRP, and restarts the I/O request at the \$QIO system service. This procedure is carried out so that the user process can receive ASTs while it waits for the page fault to complete. Once the page is faulted into memory, the \$QIO system service will resubmit the I/O request.

If the initial call was to EXE\$WRITELOCKR and an error occurs, EXE\$WRITELOCKR, by means of a coroutine call, returns control to the driver's FDT routine with status in R0. The driver performs whatever device-specific actions are required to abort the request, preserving the contents of R0 and R1. When the driver issues the RSB instruction, control is returned to EXE\$WRITELOCKR. EXE\$WRITELOCKR proceeds to abort the I/O request.

Otherwise, these routines return success status to their callers.

Operating System Routines

EXE\$WRITELOCK, EXE\$WRITELOCKR

A driver FDT routine that calls EXE\$WRITELOCKR must distinguish between successful and unsuccessful status when it resumes, as shown in the following example:

```
        JSB    G^EXE$WRITELOCKR
        BLBS   BUF_LOCK_OK
BUF_LOCK_FAIL:
;
; clean up this $QIO bookkeeping
;
        RSB
BUF_LOCK_OK:
        .
        .
;
; continue processing this I/O request
;
```

EXE\$WRTMAILBOX

Sends a message to a mailbox.

module

MBDRIVER

input

Location	Contents
R3	Message size
R4	Message address
R5	Address of mailbox UCB
Mailbox UCB fields	

output

Location	Contents
R0	SS\$_NORMAL, SS\$_MBTOOSML, SS\$_MBFULL, SS\$_INSFMEM, or SS\$_NOPRIV
R1 and R2	Destroyed

synchronization

Because EXE\$WRTMAILBOX raises IPL to IPL\$_MAILBOX and obtains the MAILBOX spin lock in a VMS multiprocessing environment, its caller cannot be executing above IPL\$_MAILBOX. EXE\$WRTMAILBOX returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

EXE\$WRTMAILBOX checks fields in the mailbox UCB (UCB\$W_BUFQUO, UCB\$W_DEVBUFFSIZ) to determine whether it can deliver a message of the specified size to the mailbox. It also checks fields in the associated ORB to determine whether the caller is sufficiently privileged to write to the mailbox. Finally, it calls EXE\$ALONONPAGED to allocate a block of nonpaged pool to contain the message. If it fails any of these operations, EXE\$WRTMAILBOX returns error status to its caller.

If it is successful thus far, EXE\$WRTMAILBOX creates a message and delivers it to the mailbox's message queue, adjusts its UCB fields accordingly, and returns success status to its caller.

Operating System Routines

EXE\$ZEROPARM

EXE\$ZEROPARM

Processes an I/O function code that requires no parameters.

module SYSQIOFDT

input

Location	Contents
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
R6	Address of CCB
R7	Bit number of the I/O function code
R8	Address of FDT entry for this routine

output

Location	Contents
IRP\$L_MEDIA	0

synchronization EXE\$ZEROPARM is called as a driver FDT routine at IPL\$_ASTDEL.

DESCRIPTION

EXE\$ZEROPARM processes an I/O function code that describes an I/O operation completely without any additional function-specific arguments. It clears IRP\$L_MEDIA and transfers control to EXE\$QIODRVPKT to deliver the IRP to the driver.

Operating System Routines
IOC\$ALOALTMAP, IOC\$ALOALTMAPN, IOC\$ALOALTMAPSP

IOC\$ALOALTMAP, IOC\$ALOALTMAPN, IOC\$ALOALTMAPSP

Allocate a set of Q22-bus alternate map registers.

module

[SYSLOA]MAPSUBxxx

input

Location	Contents
R3	Number of alternate map registers to allocate (IOC\$ALOALTMAPN and IOC\$ALOALTMAPSP only). The value should account for one extra register needed to prevent a transfer overrun.
R4	Number of first alternate map register to allocate (IOC\$ALOALTMAPSP only).
R5	Address of UCB.
UCB\$W_BCNT	Transfer byte count (IOC\$ALOALTMAP only).
UCB\$W_BOFF	Byte offset in page (IOC\$ALOALTMAP only).
UCB\$L_CRB	Address of CRB.
CRB\$L_INTD+ VEC\$L_ADP	Address of ADP.
CRB\$L_INTD+ VEC\$W_MAPALT	VEC\$V_ALTLOCK set indicates that alternate map registers have been permanently allocated to this controller.
ADP\$W_MR2NREGAR, ADP\$W_MR2FREGAR, ADP\$L_MR2ACTMDR	Alternate map register descriptor arrays.

output

Location	Contents
R0	SS\$_NORMAL, SS\$_INSFMAPREG, or SS\$_SSFAIL
R1	Destroyed
R2	Address of ADP
CRB\$L_INTD+ VEC\$W_NUMALT	Number of alternate map registers allocated
CRB\$L_INTD+ VEC\$W_MAPALT	Starting alternate map register number
ADP\$W_MR2NREGAR, ADP\$W_MR2FREGAR, ADP\$L_MR2ACTMDR	Updated

Operating System Routines

IOC\$ALOALTMAP, IOC\$ALOALTMAPN, IOC\$ALOALTMAPSP

synchronization

Callers of IOC\$ALOALTMAP, IOC\$ALOALTMAPN, or IOC\$ALOALTMAPSP may be executing at fork IPL or above and must hold the corresponding fork lock in a VMS multiprocessing environment. Each routine returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

IOC\$ALOALTMAP, IOC\$ALOALTMAPN, and IOC\$ALOALTMAPSP allocate a contiguous set of Q22-bus alternate map registers (registers 496 to 8191) and record the allocation in the ADP and CRB. These routines differ in the way in which they determine the number and location of the alternate map registers they allocate:

- IOC\$ALOALTMAP calculates the number of needed map registers using the values contained in UCB\$W_BCNT and UCB\$W_BOFF. It automatically allocates one extra map register. When it is later called by the driver, IOC\$LOADALTMAP marks this register invalid to prevent a transfer overrun.
- IOC\$ALOALTMAPN uses the value in R3 as the number of required registers.
- IOC\$ALOALTMAPSP uses the value in R3 as the number of required registers and attempts to allocate these registers starting at the one indicated by R4.

If an odd number of map registers is required, these routines round this value up to an even multiple.

If alternate map registers have been permanently allocated to the controller, IOC\$ALOALTMAP, IOC\$ALOALTMAPN, or IOC\$ALOALTMAPSP returns successfully to its caller without allocating the requested map registers. Otherwise, it searches the alternate map register descriptor arrays for the required number of map registers. If there are not enough contiguous map registers available, the routine returns SS\$_INSFMAPREG status.

If the VAX system does not support alternate map registers, the routine exits with SS\$_SSFAIL status.

IOC\$ALOUBAMAP, IOC\$ALOUBAMAPN

Allocate a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers.

module

IOSUBNPAG

input

Location	Contents
R3	Number of map registers to allocate (IOC\$ALOUBAMAPN only). The value should account for one extra register needed to prevent a transfer overrun.
R5	Address of UCB.
UCB\$W_BCNT	Transfer byte count (IOC\$ALOUBAMAP only).
UCB\$W_BOFF	Byte offset in page (IOC\$ALOUBAMAP only).
UCB\$L_CRB	Address of CRB.
CRB\$L_INTD+ VEC\$L_ADP	Address of ADP.
CRB\$L_INTD+ VEC\$W_MAPREG	VEC\$V_MAPLOCK set indicates that map registers have been permanently allocated to this controller.
ADP\$W_MRNREGARY, ADP\$W_MRFREGARY, ADP\$L_MRACTMDRS	Map register descriptor arrays.

output

Location	Contents
R0	SS\$_NORMAL or 0
R1	Destroyed
R2	Address of ADP
CRB\$L_INTD+VEC\$B_NUMREG	Number of map registers allocated
CRB\$L_INTD+VEC\$W_MAPREG	Starting map register number
ADP\$W_MRNREGARY, ADP\$W_MRFREGARY, ADP\$L_MRACTMDRS	Updated

synchronization

The caller of IOC\$ALOUBAMAP or IOC\$ALOUBAMAPN may be executing at fork IPL or above and must hold the corresponding fork lock in a VMS multiprocessing environment. Either routine returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

Operating System Routines

IOC\$ALOUBAMAP, IOC\$ALOUBAMAPN

DESCRIPTION

IOC\$ALOUBAMAP and IOC\$ALOUBAMAPN allocate a contiguous set of UNIBUS map registers or a set of the first 496 Q22-bus map registers and record the allocation in the ADP and CRB. These routines differ in the way in which they determine the number of the map registers they allocate:

- IOC\$ALOUBAMAP calculates the number of needed map registers using the values contained in UCB\$W_BCNT and UCB\$W_BOFF. It automatically allocates one extra map register. When it is later called by the driver, IOC\$LOADUBAMAP marks this register invalid to prevent a transfer overrun.
- IOC\$ALOUBAMAPN uses the value in R3 as the number of required registers.

If an odd number of map registers is required, both routines round this value up to an even multiple.

If map registers have been permanently allocated to the controller, IOC\$ALOUBAMAP or IOC\$ALOUBAMAPN returns successfully to its caller without allocating the requested map registers. Otherwise, it searches the map register descriptor arrays for the required number of map registers. If there are not enough contiguous map registers available, the routine returns an error status of zero to its caller.

IOC\$APPLYECC

Applies an ECC correction to data transferred from a disk device into memory.

module

IOSUBRAMS

input

Location	Contents
R0	Number of bytes of data that have been transferred, not including the block to be corrected; this must be a multiple of 512 bytes
R5	Address of UCB
UCB\$W_BCNT	Length of transfer in bytes
UCB\$W_EC1	Starting bit number of the error burst
UCB\$W_EC2	Exclusive OR correction pattern
UCB\$L_SVPN	Address of system PTE for a page that is available for use by driver
UCB\$L_SVAPTE	System virtual address of PTE that maps the transfer

output

Location	Contents
R0, R1, R2	Destroyed
UCB\$W_DEVSTS	UCB\$V_ECC set to indicate that an ECC correction was made

synchronization

IOC\$APPLYECC executes at the caller's IPL, obtains no spin locks, and returns control to its caller at its caller's IPL.

DESCRIPTION

IOC\$APPLYECC corrects data transferred from a disk device to memory by performing an exclusive-OR operation on the data and applying a correction pattern from the UCB. IOC\$APPLYECC also sets a UCB bit (UCB\$V_ECC in UCB\$W_DEVSTS) to indicate that it has made an ECC correction.

Note that, to use this routine, the driver must define the local UCB disk extension, as described in Section 1.17.

Operating System Routines

IOC\$CANCELIO

IOC\$CANCELIO

Conditionally marks a UCB so that its current I/O request will be canceled.

module

IOSUBNPAG

input

Location	Contents
R2	Channel index number
R3	Address of IRP
R4	Address of current PCB
R5	Address of UCB
IRP\$L_PID	Process identification of the process that queued the I/O request
IRP\$W_CHAN	I/O request channel index number
PCB\$L_PID	Process identification of the process that requested cancellation
UCB\$L_STS	UCB\$V_BSY set if device is busy, clear if device is idle

output

Location	Contents
UCB\$L_STS	UCB\$V_CANCEL set if the I/O request should be canceled

synchronization

IOC\$CANCELIO executes at its caller's IPL, obtains no spin locks, and returns control to its caller at the caller's IPL. It is usually called by EXE\$CANCEL (if specified in the DDT as the driver's cancel-I/O routine) at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment.

DESCRIPTION

IOC\$CANCELIO cancels I/O to a device in the following device-independent manner:

- 1 It confirms that the device is busy by examining the device-busy bit in the UCB status longword (UCB\$V_BSY in UCB\$L_STS).
- 2 It confirms that the IRP in progress on the device originates from the current process (that is, the contents of IRP\$L_PID and PCB\$L_PID are identical).
- 3 It confirms that the specified channel-index number is the same as the value stored in the IRP's channel-index field (IRP\$W_CHAN).
- 4 It sets the cancel-I/O bit in the UCB status longword (UCB\$V_CANCEL in UCB\$L_STS).

IOC\$DIAGBUFILL

Fills a diagnostic buffer if the original \$QIO request specified such a buffer.

module

IOSUBNPAG

input

Location	Contents
R4	Address of device's CSR
R5	Address of UCB
UCB\$L_IRP	Address of current IRP
IRP\$W_STS	IRP\$V_DIAGBUF set if a diagnostic buffer exists
IRP\$L_DIAGBUF	Address of diagnostic buffer, if one is present
UCB\$B_ERTCNT	Final error retry count
UCB\$L_DDB	Address of DDB
DDB\$L_DDT	Address of DDT
DDT\$L_REGDUMP	Address of driver's register dumping routine
EXE\$GQ_SYSTIME	Current system time (time at I/O request completion)

output

Location	Contents
R0, R1	Destroyed
R2	Address of DDT
R3	Address of IRP
R4	Address of device's CSR
R5	Address of UCB

synchronization

The caller of IOC\$DIAGBUFILL may be executing at or above fork IPL and must hold the corresponding fork lock in a VMS multiprocessing environment. IOC\$DIAGBUFILL returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

A device driver fork process calls IOC\$DIAGBUFILL at the end of I/O processing but before releasing the I/O channel. IOC\$DIAGBUFILL stores the I/O completion time and the final error retry count in the diagnostic buffer. (IOC\$INITIATE has already placed the I/O initiation time in the first quadword of the buffer.) IOC\$DIAGBUFILL then calls the driver's register dumping routine, which fills the remainder of the buffer, and returns to its caller.

Operating System Routines

IOC\$INITIATE

IOC\$INITIATE

Initiates the processing of the next I/O request for a device unit.

module

IOSUBNPAG

input

Location	Contents
R3	Address of IRP
R5	Address of UCB
CPU\$L_PHY_CPUID	CPU ID of local processor
IRP\$L_SVAPTE	Address of system buffer (buffered I/O) or system virtual address of the PTE that maps process buffer (direct I/O)
IRP\$W_BOFF	Byte offset of start of buffer
IRP\$L_BCNT	Size in bytes of transfer
IRP\$W_STS	IRP\$V_DIAGBUF set if a diagnostic buffer exists
IRP\$L_DIAGBUF	Address of diagnostic buffer, if one is present
EXE\$GQ_SYSTIME	Current system time (when I/O processing began)
UCB\$L_DDB	Address of DDB
UCB\$L_DDT	Address of DDT
UCB\$L_AFFINITY	Device's affinity mask
DDT\$L_START	Address of driver start-I/O routine

output

Location	Contents
R0, R1	Destroyed
UCB\$L_IRP	Address of IRP
UCB\$L_SVAPTE	IRP\$L_SVAPTE
UCB\$W_BOFF	IRP\$W_BOFF
UCB\$W_BCNT	IRP\$L_BCNT (low-order word)
UCB\$L_STS	UCB\$V_CANCEL and UCB\$V_TIMEOUT cleared
Diagnostic buffer	Current system time (first quadword)

synchronization

IOC\$INITIATE is called at fork IPL with the corresponding fork lock held in a VMS multiprocessing system. Within this context, it transfers control to the driver's start-I/O routine.

DESCRIPTION

IOC\$INITIATE creates the context in which a driver fork process services an I/O request. IOC\$INITIATE creates this context and activates the driver's start-I/O routine in the following steps:

- Checks the CPU ID of the local processor against the device's affinity mask to determine whether the local processor can initiate the I/O operation on the device. If it cannot, IOC\$INITIATE takes steps to initiate the I/O function on another processor in a VMS multiprocessing system. It then returns to its caller.
- Stores the address of the current IRP in UCB\$L_IRP.
- Copies the transfer parameters contained in the IRP into the UCB:
 - Copies the address of the system buffer (buffered I/O) or the system virtual address of the PTE that maps process buffer (direct I/O) from IRP\$L_SVAPTE to UCB\$L_SVAPTE
 - Copies the byte offset within the page from IRP\$W_BOFF to UCB\$W_BOFF
 - Copies the low-order word of the byte count from IRP\$L_BCNT to UCB\$W_BCNT
- Clears the cancel-I/O and timeout bits in the UCB status longword (UCB\$V_CANCEL and UCB\$V_TIMEOUT in UCB\$L_STS).
- If the I/O request specifies a diagnostic buffer, as indicated by IRP\$V_DIAGBUF in IRP\$W_STS, stores the system time in the first quadword of the buffer to which IRP\$L_DIAGBUF points (the \$QIO system service having already allocated the buffer).
- Transfers control to the driver's start-I/O routine.

Operating System Routines

IOC\$IOPPOST

IOC\$IOPPOST

Performs device-independent I/O postprocessing and delivers the results of an I/O request to a process.

module

IOCIOPPOST

input

Location	Contents
IRP\$L_PID	Process identification of the process that initiated the I/O request
IRP\$L_UCB	Address of UCB
IRP\$W_STS	IRP\$V_BUFIO set if buffered-I/O request, clear if direct-I/O request; IRP\$V_PHYSIO set if physical-I/O function; IRP\$V_EXTEND set if an IRPE is linked to this IRP; IRP\$V_KEY set if IRP\$L_KEYDESC contains the address of an encryption key buffer; IRP\$V_FUNC set if read function, clear if write function; IRP\$V_DIAGBUF set if diagnostic buffer exists; IRP\$V_MBXIO set if mailbox read function
IRP\$L_DIAGBUF	Address of diagnostic buffer, if one is present
IRP\$L_SVAPTE	Address of system buffer (buffered I/O) or system virtual address of the PTE that maps process buffer (direct I/O)
IRP\$W_BOFF	Byte offset of start of buffer
IRP\$L_BCNT	Size in bytes of transfer
IRP\$L_OBCNT	Original byte count for virtual I/O transfer
IRP\$L_IOST1	First I/O status longword
IRP\$W_CHAN	I/O request channel index number
IRP\$L_IOSB	Address of I/O status block, if specified
IRP\$B_RMOD	Access mode of I/O request; ACB\$V_QUOTA set if request specified AST
IRP\$B_EFN	Event flag number
UCB\$W_QLEN	Length of pending-I/O queue
UCB\$L_DEVCHAR	DEV\$V_FOD set if file-oriented device
PCB\$W_DIOCNT	Process's direct-I/O count
PCB\$W_BIOCNT	Process's buffered-I/O count
JIB\$L_BYTCNT	Job byte count quota
CCB\$W_IOC	Number of outstanding I/O requests on channel
CCB\$L_DIRP	Address of IRP for requested deaccess

output

Location	Contents
UCB\$W_QLEN	Decrement
PCB\$W_DIOCNT	Incremented for a direct-I/O request
PCB\$W_BIOCNT	Incremented for a buffered I/O request
JIB\$L_BYTCNT	Updated for buffered I/O request
CCB\$W_IOC	Decrement
CCB\$L_DIRP	Cleared if channel is idle

synchronization

IOC\$IOPOST executes in response to an interrupt granted at IPL\$_IOPOST. It performs some of its functions in a special kernel-mode AST that executes within process context at IPL\$_ASTDEL. It obtains and releases the various spin locks required to deallocate nonpaged pool and adjust process quotas.

DESCRIPTION

This interrupt service routine processes IRPs in the systemwide and local CPU I/O postprocessing queues, gaining control when the processor grants a software interrupt at IPL\$_IOPOST. When the I/O postprocessing queues are empty, IOC\$IOPOST dismisses the interrupt with an REI instruction.

IOC\$IOPOST performs several tasks to complete either a direct- or buffered-I/O request:

- For a *buffered-I/O* read request, it copies data from the system buffer to the process buffer. If it cannot write to the process buffer, it returns SS\$_ACCVIO status. For read and write requests, it releases the system buffer to nonpaged pool.
- For a *direct-I/O* request, it unlocks those process buffer pages that were locked for the I/O transfer. (If an IRPE exists, the unlocked pages include any defined in the IRPE area descriptors.)

IOC\$IOPOST performs the following tasks for *both* direct and buffered I/O requests:

- Decrements the device's pending-I/O queue length
- Adjusts direct-I/O or buffered-I/O quota use
- Sets an event flag if one was specified in the \$QIO system service call
- Copies I/O completion status from the IRP to the process's I/O status block (if one was specified in the \$QIO system service call).
- Queues a user mode AST (if specified) to the process
- Copies the diagnostic buffer (if specified) from system to process space and releases the system buffer
- Deallocates the IRP and any IRPEs

Note that many of these operations are performed within process context by the special kernel-mode AST IOC\$IOPOST queues to the process.

Operating System Routines

IOC\$LOADALTMAP

IOC\$LOADALTMAP

Loads a set of Q22-bus alternate map registers.

module [SYSLOA]MAPSUBxxx

macro LOADALT

input

Location	Contents
R5	Address of UCB
UCB\$W_BCNT	Number of bytes in transfer
UCB\$W_BOFF	Byte offset in first page of transfer
UCB\$L_SVAPTE	System virtual address of PTE for first page of transfer
UCB\$L_CRB	Address of CRB
CRB\$L_INTD+ VEC\$W_NUMALT	Number of alternate map registers allocated
CRB\$L_INTD+ VEC\$W_MAPALT	Number of first alternate map register allocated
CRB\$L_INTD+ VEC\$L_ADP	Address of ADP
ADP\$L_MR2ADDR	Address of the first Q22-bus alternate map register

output

Location	Contents
R0	SS\$_NORMAL, SS\$_INSFMAPREG, or SS\$_SSFAIL
R1, R2	Destroyed

synchronization

A driver fork process calls IOC\$LOADALTMAP at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment. IOC\$LOADALTMAP returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

A driver fork process calls IOC\$LOADALTMAP to load a previously-allocated set of alternate map registers with page-frame numbers (PFNs). This enables a device DMA transfer to or from the buffer indicated by the contents of UCB\$L_SVAPTE, UCB\$W_BCNT, and UCB\$W_BOFF.

IOC\$LOADALTMAP confirms that sufficient alternate map registers have been previously allocated. If not, it issues a UBMAPEXCED bugcheck. Otherwise, it loads the appropriate PFN into each map register and sets

Operating System Routines

IOC\$LOADALTMAP

the map register valid bit. It clears the last map register. This last invalid register prevents a transfer overrun.

If the VAX system does not support alternate map registers, the routine exits with SS\$_SSFAIL status.

Operating System Routines

IOC\$LOADMBAMAP

IOC\$LOADMBAMAP

Loads MASSBUS map registers.

module

LOADMREG

macro

LOADMBA

input

Location	Contents
R4	Address of MBA configuration register (MBA\$_CSR)
R5	Address of UCB
UCB\$_BCNT	Number of bytes in transfer
UCB\$_BOFF	Byte offset in first page of transfer
UCB\$_SVAPTE	System virtual address of PTE for first page of transfer
MBA\$_MAP	Address of first MASSBUS map register

output

Location	Contents
R0, R1, R2	Destroyed

synchronization

A driver fork process calls IOC\$LOADMBAMAP at fork IPL. IOC\$LOADMBAMAP returns control to its caller at the caller's IPL.

DESCRIPTION

Driver fork processes for DMA transfers call IOC\$LOADMBAMAP to load MASSBUS adapter map registers with page-frame numbers (PFNs).

IOC\$LOADMBAMAP uses the contents of UCB\$_SVAPTE, UCB\$_BCNT, and UCB\$_BOFF to determine the number of pages involved in the transfer. It then copies the page frame numbers from the page-table entries associated with this buffer into map registers, starting with map register 0. IOC\$LOADMBAMAP also loads the negated transfer size into the MASSBUS adapter's byte count register (MBA\$_BCR) and the byte offset of the transfer into the MASSBUS adapter's virtual address register (MBA\$_VAR). It clears the last map register. This last invalid register prevents a transfer overrun.

The driver must own the MASSBUS adapter, and thus its map registers, before it calls this routine.

IOC\$LOADUBAMAP, IOC\$LOADUBAMAPA

Load a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers.

module LOADMREG

macro LOADUBA

input

Location	Contents
R5	Address of UCB
UCB\$W_BCNT	Number of bytes in transfer
UCB\$W_BOFF	Byte offset in first page of transfer
UCB\$L_SVAPTE	System virtual address of PTE for first page of transfer
UCB\$L_CRB	Address of CRB
CRB\$L_INTD+ VEC\$B_NUMREG	Number of map registers allocated
CRB\$L_INTD+ VEC\$W_MAPREG	Number of first map register allocated
CRB\$L_INTD+ VEC\$B_DATAPATH	Data path specifier; VEC\$V_LWAE set if longword buffering is used, clear if quadword buffering is used
CRB\$L_INTD+ VEC\$L_ADP	Address of ADP
UBA\$L_MAP	Address of first UNIBUS or Q22-bus map register
UCB\$L_SVAPTE	System virtual address of PTE for the first page of the transfer

output

Location	Contents
R0, R1, R2	Destroyed

synchronization

A driver fork process calls IOC\$LOADUBAMAP or IOC\$LOADUBAMAPA at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment. Either routine returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

A driver fork process calls IOC\$LOADUBAMAP or IOC\$LOADUBAMAPA to load a previously-allocated set of map registers with page-frame numbers (PFNs). This enables a device DMA transfer to or from the buffer indicated by the contents of UCB\$L_SVAPTE, UCB\$W_BCNT, and UCB\$W_BOFF.

Operating System Routines

IOC\$LOADUBAMAP, IOC\$LOADUBAMAPA

Either IOC\$LOADUBAMAP or IOC\$LOADUBAMAPA confirms that sufficient map registers have been previously allocated. If not, it issues a UBMAPEXCED bugcheck. Otherwise, it loads into each map register the appropriate PFN and data-path number. It sets the map register valid bit and, if VEC\$V_LWAE is set in VEC\$B_DATAPATH, the longword-access-enable bit.

IOC\$LOADUBAMAP checks the low bit of UCB\$W_BOFF to determine whether the transfer is byte-aligned or word-aligned. If the low bit is set, it sets the byte-offset bit in each map register. Drivers for byte-aligned UNIBUS devices that must never set the byte-offset bit call IOC\$LOADUBAMAPA. Drivers for Q22-bus-only devices also call IOC\$LOADUBAMAPA as there is no byte-offset bit in a Q22-bus map register.

Both IOC\$LOADUBAMAP and IOC\$LOADUBAMAPA clear the last map register. This last invalid register prevents a transfer overrun.

IOC\$MOVFRUSER, IOC\$MOVFRUSER2

Move data from a user buffer to a device.

module

BUFFERCTL

input

Location	Contents
R0	Address of byte to be moved (IOC\$MOVFRUSER2 only)
R1	Address of driver's buffer
R2	Number of bytes to move
R5	Address of UCB
DPT\$B_FLAGS	Bit DPT\$V_SVP set (causing a system page-table entry (SPTE) to be allocated to the driver)
UCB\$L_SVAPTE	System virtual address of PTE that maps the first page of the buffer
UCB\$L_SVPN	System virtual page number of SPTE allocated to driver
UCB\$W_BOFF	Byte offset to start of transfer in page

output

R0	Next address of user's buffer
----	-------------------------------

synchronization

The caller of IOC\$MOVFRUSER or IOC\$MOVFRUSER2 may be executing at fork IPL or above and must hold the corresponding fork lock in a VMS multiprocessing environment. Either routine returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

A driver calls IOC\$MOVFRUSER and IOC\$MOVFRUSER2 to move data from a user buffer to a device that cannot itself map the user buffer to system virtual addresses (for instance, a non-DMA device).

In order to accomplish the move, IOC\$MOVFRUSER and IOC\$MOVFRUSER2 first map the user buffer using the system page-table entry (SPTE) the driver allocated in a DPTAB macro invocation. If an SPTE has not been allocated to the driver, these routines cause an access violation when they attempt to refer to the location addressed by the contents of the field UCB\$L_SVAPTE. (See the description of the DPTAB macro in Chapter 2 for information on how to allocate this SPTE.)

IOC\$MOVFRUSER2 is useful for moving blocks of data in several pieces, each piece beginning within a page rather than on a page boundary. To begin, the driver calls IOC\$MOVFRUSER. For each subsequent piece, the driver calls IOC\$MOVFRUSER2.

Operating System Routines

IOC\$MOVTOUSER, IOC\$MOVTOUSER2

IOC\$MOVTOUSER, IOC\$MOVTOUSER2

Move data from a device to a user buffer.

module

BUFFERCTL

input

Location	Contents
R0	User buffer address to which to move the byte (IOC\$MOVTOUSER2 only)
R1	Address of driver's buffer
R2	Number of bytes to move
R5	Address of UCB
DPT\$B_FLAGS	Bit DPT\$V_SVP set (causing a system page-table entry (SPTE) to be allocated to the driver)
UCB\$L_SVAPTE	System virtual address of PTE that maps the first page of the buffer
UCB\$L_SVPN	System virtual page number of SPTE allocated to driver
UCB\$W_BOFF	Byte offset to start of transfer in page

output

Location	Contents
R0	Next starting address of user's buffer

synchronization

The caller of IOC\$MOVTOUSER or IOC\$MOVTOUSER2 may be executing at fork IPL or above and must hold the corresponding fork lock in a VMS multiprocessing environment. Either routine returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

A driver calls IOC\$MOVTOUSER and IOC\$MOVTOUSER2 to move data from a device to a user buffer when the device itself (for instance, a non-DMA device) cannot map the user buffer to system virtual addresses.

In order to accomplish the move, IOC\$MOVTOUSER and IOC\$MOVTOUSER2 first map the user buffer using the system page-table entry (SPTE) the driver allocated in a DPTAB macro invocation. If an SPTE has not been allocated to the driver, these routines cause an access violation when they attempt to refer to the location addressed by the contents of the field UCB\$L_SVAPTE. (See the description of the DPTAB macro in Chapter 2 for information on how to allocate this SPTE.)

Operating System Routines IOC\$MOVTOUSER, IOC\$MOVTOUSER2

IOC\$MOVTOUSER2 is useful for moving blocks of data in several pieces, each piece beginning within a page rather than on a page boundary. It handles as many pages as you need. To begin, the driver calls IOC\$MOVTOUSER. For each subsequent buffer to move, the driver calls IOC\$MOVTOUSER2.

Operating System Routines

IOC\$PURGDATAP

IOC\$PURGDATAP

Purges the buffered data path and logs memory errors that may have occurred during an I/O transfer.

module [SYSLOA]LIOSUBxxx

macro PURDPR

input

Location	Contents
R5	Address of UCB

output

Location	Contents
R0	Bit 0 set if success, clear if failure
R1	Contents of data path after purge
R2	Address of start of the I/O bus map registers
R3	Address of CRB

synchronization

The caller of IOC\$PURGDATAP may be executing at fork IPL or above and must hold the corresponding fork lock in a VMS multiprocessing environment. It returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

All device drivers that support DMA transfers, including those on VAX systems that have no buffered data paths (such as the MicroVAX systems), call IOC\$PURGDATAP after a data transfer.

IOC\$PURGDATAP performs the following tasks:

- Obtains the start of adapter register space using the following chain of pointers:
UCB\$L_CRB → CRB\$L_INTD+VEC\$L_ADP → ADP\$L_CSR
- Extracts the caller's data path number (buffered or direct) from the CRB.
- Purges the data path if it is a buffered data path. Note that a purge of a direct data path (data path 0) is legal and always results in success status.
- Stores the contents of the data path register in R1. The driver's register dumping routine writes this value to the error message buffer.
- Clears any purge errors in the data path register.
- Places the appropriate return status in R0.

Operating System Routines

IOC\$PURGDATAP

- Determines the base of UNIBUS or Q22-bus map registers and writes the value into R2. The driver's register dumping routine writes this value to the error message buffer.
- In some machine implementations, checks for memory errors that might have occurred during the DMA operation and, if an error is detected, logs it.

Operating System Routines

IOC\$RELALTMAP

IOC\$RELALTMAP

Releases a set of Q22-bus alternate map registers.

module [SYSLOA]MAPSUBxxx

macro RELALT

input

Location	Contents
R5	Address of UCB
UCB\$_CRB	Address of CRB
CRB\$_INTD+ VEC\$_ADP	Address of ADP
CRB\$_INTD+ VEC\$_MAPALT	Starting alternate map register number; VEC\$_ALTLOCK set indicates that alternate map registers have been permanently allocated to this controller
CRB\$_INTD+ VEC\$_NUMALT	Number of allocated alternate map registers
ADP\$_MR2QFL	Head of queue of UCBs waiting for alternate map registers
ADP\$_MR2NREGAR, ADP\$_MR2FREGAR, ADP\$_MR2ACTMDR	Alternate map register descriptor arrays

output

Location	Contents
R0	SS\$_NORMAL or SS\$_SSFAIL
R1, R2	Destroyed
ADP\$_MR2NREGAR, ADP\$_MR2FREGAR, ADP\$_MR2ACTMDR	Updated

synchronization A driver fork process calls IOC\$RELALTMAP at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment.

DESCRIPTION

A driver fork process calls IOC\$RELALTMAP to release a previously-allocated set of Q22-bus alternate map registers (registers 496 to 8191) and update the alternate map register descriptor arrays in the ADP. IOC\$RELALTMAPREG assumes that its caller is the current owner of the controller data channel.

Operating System Routines

IOC\$RELALTMAP

IOC\$RELALTMAP obtains the location and number of the allocated map registers from CRB\$L_INTD+VEC\$W_MAPALT and CRB\$L_INTD+VEC\$W_NUMALT, respectively. If VEC\$V_ALTLOCK is set in CRB\$L_INTD+VEC\$W_MAPALT, the alternate map registers have been permanently allocated to the controller and IOC\$RELALTMAP returns successfully to its caller.

After adjusting the alternate map register descriptor arrays, IOC\$RELALTMAP examines the alternate-map-register wait queue. If the queue is empty, IOC\$RELALTMAP returns successfully to its caller. If the queue contains waiting fork processes, IOC\$RELALTMAP dequeues the first process and calls IOC\$ALOALTMAP to attempt to allocate the set of map registers it requires.

If there are sufficient alternate map registers, IOC\$RELALTMAP restores R3 through R5 to the process and reactivates it. When this fork process returns control to IOC\$RELALTMAP, IOC\$RELALTMAP attempts to allocate map registers to the next waiting fork process. IOC\$RELALTMAP continues to allocate map registers in this manner until the alternate-map-register wait queue is empty or it cannot satisfy the requirements of the process at the head of the queue. In the latter event, IOC\$RELALTMAP reinserts the fork process's UCB in the queue and returns successfully to its caller.

If the VAX system does not support alternate map registers, IOC\$RELALTMAP exits with SS\$_SSFAIL status.

Operating System Routines

IOC\$RELCHAN

IOC\$RELCHAN

Releases device ownership of all controller data channels.

module IOSUBNPAG

macro RELCHAN

input

Location	Contents
R5	Address of UCB
UCB\$L_CRB	Address of CRB
CRB\$L_LINK	Address of secondary CRB
CRB\$B_MASK	CRB\$V_BSY set if the channel is busy
CRB\$L_INTD+VEC\$L_IDB	Address of IDB
IDB\$L_OWNER	Address of UCB of channel owner
CRB\$L_WQFL	Head of queue of UCBs waiting for the controller channel

output

Location	Contents
R0, R1, R2	Destroyed
IDB\$L_OWNER	Cleared if no driver is waiting for the channel
CRB\$B_MASK	CRB\$V_BSY cleared if no driver is waiting for the channel

synchronization

A driver fork process calls IOC\$RELCHAN at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment. IOC\$RELCHAN returns control to its caller after resuming execution of other fork processes waiting for a controller channel.

DESCRIPTION

A driver fork process calls IOC\$RELCHAN to release all controller data channel assigned to a device; it calls IOC\$RELSCHAN to release only the secondary data channel.

If the channel wait queue contains waiting fork processes, IOC\$RELCHAN dequeues a process, assigns the channel to that process, restores R3 and R5, moves the address of the CSR (IDB\$L_CSR) into R4, and reactivates the suspended fork process.

IOC\$RELDATAP

Releases a UNIBUS adapter's buffered data path.

module

IOSUBNPAG

macro

RELDPR

input

Location	Contents
R5	Address of UCB
UCB\$L_CRB	Address of CRB
CRB\$L_INTD+ VEC\$L_ADP	Address of ADP
CRB\$L_INTD+ VEC\$B_DATAPATH	Data path specifier; VEC\$V_PATHLOCK set if the data path has been permanently allocated to the controller
ADP\$L_DPQFL	Head of queue of UCBs waiting for a UNIBUS adapter buffered data path
ADP\$W_DPBITMAP	Data path bit map

output

Location	Contents
R0, R1, R2	Destroyed
ADP\$W_DPBITMAP	Bit representing data path set if the path is not allocated to another driver fork process
CRB\$L_INTD+ VEC\$B_DATAPATH	Bits 0 through 4 cleared if the path is not permanently allocated

synchronization

A driver fork process calls IOC\$RELDATAP at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment. IOC\$RELDATAP returns control to its caller after resuming execution of any other fork processes waiting for a buffered data path.

DESCRIPTION

A driver fork process must own a UNIBUS buffered data path when it calls IOC\$RELDATAP.

IOC\$RELDATAP obtains the number of the allocated data path from bits 0 through 4 of the data path specifier. If VEC\$V_PATHLOCK is set in the specifier, the data path has been permanently allocated to the controller and IOC\$RELDATAP returns to its caller.

Operating System Routines

IOC\$RELDATAP

If the data path wait queue contains waiting fork processes, IOC\$RELDATAP dequeues the first process, allocates the data path to it, restores R3 through R5, and reactivates it. Otherwise, it marks the path available by setting the corresponding bit in the data path bit map (ADP\$W_DPBITMAP), and returns to its caller.

If the bit map has been corrupted, IOC\$RELDATAP issues an INCONSTATE bugcheck.

IOC\$RELMAPREG

Releases a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers.

module IOSUBNPAG

macro RELMPR

input

Location	Contents
R5	Address of UCB
UCB\$_CRB	Address of CRB
CRB\$_INTD+ VEC\$_ADP	Address of ADP
CRB\$_INTD+ VEC\$_MAPREG	Starting map register number; VEC\$_MAPLOCK set indicates that map registers have been permanently allocated to this controller
CRB\$_INTD+ VEC\$_NUMREG	Number of allocated map registers
ADP\$_MRQFL	Head of queue of UCBs waiting for map registers
ADP\$_MRNREGARY, ADP\$_MRFREGARY, ADP\$_MRACTMDRS	Map register descriptor arrays

output

Location	Contents
R0	SS\$_NORMAL or SS\$_SSFAIL
R1, R2	Destroyed
ADP\$_MRNREGARY, ADP\$_MRFREGARY, ADP\$_MRACTMDRS	Updated

synchronization

A driver fork process calls IOC\$RELMAPREG at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment.

DESCRIPTION

A driver fork process calls IOC\$RELMAPREG to release a previously-allocated set of UNIBUS map registers or a set of the first 496 Q22-bus map registers. IOC\$RELMAPREG updates the alternate map register descriptor arrays in the ADP. IOC\$RELMAPREG assumes that its caller is the current owner of the controller data channel.

Operating System Routines

IOC\$RELMAPREG

IOC\$RELMAPREG obtains the location and number of the allocated map registers from CRB\$L_INTD+VEC\$W_MAPREG and CRB\$L_INTD+VEC\$B_NUMREG, respectively. If VEC\$V_MAPLOCK is set in CRB\$L_INTD+VEC\$W_MAPREG, the map registers have been permanently allocated to the controller and IOC\$RELMAPREG returns successfully to its caller.

After adjusting the map register descriptor arrays, IOC\$RELMAPREG examines the standard-map-register wait queue. If the queue is empty, IOC\$RELMAPREG returns successfully to its caller. If the queue contains waiting fork processes, IOC\$RELMAPREG dequeues the first process and calls IOC\$ALOUBAMAP to attempt to allocate the set of map registers it requires.

If there are sufficient map registers, IOC\$RELMAPREG restores R3 through R5 to the process and reactivates it. When this fork process returns control to IOC\$RELMAPREG, IOC\$RELMAPREG attempts to allocate map registers to the next waiting fork process. IOC\$RELMAPREG continues to allocate map registers in this manner until the standard-map-register wait queue is empty or it cannot satisfy the requirements of the process at the head of the queue. In the latter event, IOC\$RELMAPREG reinserts the fork process's UCB in the queue and returns successfully to its caller.

IOC\$RELSCHAN

Releases device ownership of only the secondary controller's data channel.

module

IOSUBNPAG

macro

RELSCHAN

input

Location	Contents
R5	Address of UCB
UCB\$L_CRB	Address of CRB
CRB\$L_LINK	Address of secondary CRB
CRB\$B_MASK	CRB\$V_BSY set if the channel is busy
CRB\$L_INTD+VEC\$L_IDB	Address of IDB
IDB\$L_OWNER	Address of UCB of channel owner
CRB\$L_WQFL	Head of queue of UCBs waiting for the controller channel

output

Location	Contents
R0, R1, R2	Destroyed
IDB\$L_OWNER	Cleared if no driver is waiting for the channel
CRB\$B_MASK	CRB\$V_BSY cleared if no driver is waiting for the channel

synchronization

A driver fork process calls IOC\$RELSCHAN at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment. IOC\$RELSCHAN returns control to its caller after resuming execution of other fork processes waiting for the secondary controller's channel.

DESCRIPTION

IOC\$RELSCHAN releases a secondary controller's data channel (for instance, the MASSBUS adapter's controller data channel). The caller retains ownership of the primary controller's data channel. A driver fork process calls IOC\$RELCHAN to release all controller data channels assigned to a device.

If the secondary channel's wait queue contains waiting fork processes, IOC\$RELSCHAN dequeues a process, assigns the channel to that process, restores R3 through R5, and reactivates the suspended process.

Operating System Routines

IOC\$REQALTMAP

IOC\$REQALTMAP

Allocates sufficient Q22-bus alternate map registers to accommodate a DMA transfer and, if unavailable, places the requesting fork process in an alternate-map-register wait queue.

module SYSLOA[MAPSUB]xxx

macro REQALT

input

Location	Contents
R5	Address of UCB
00(SP)	Return PC of caller
04(SP)	Return PC of caller's caller
UCB\$W_BCNT	Transfer byte count
UCB\$W_BOFF	Byte offset in page
UCB\$L_CRB	Address of CRB
CRB\$L_INTD+ VEC\$L_ADP	Address of ADP
CRB\$L_INTD+ VEC\$W_MAPALT	VEC\$V_ALTLOCK set indicates that alternate map registers have been permanently allocated to this controller
ADP\$W_MR2NREGAR, ADP\$W_MR2FREGAR, ADP\$L_MR2ACTMDR	Alternate map register descriptor arrays
ADP\$L_MR2QBL	Tail of queue of UCBs waiting for alternate map registers

output

Location	Contents
R0	SS\$_NORMAL or SS\$_SSFAIL
R1	Destroyed
R2	Address of ADP
CRB\$_INTD+ VEC\$_W_NUMALT	Number of alternate map registers allocated
CRB\$_INTD+ VEC\$_W_MAPALT	Starting alternate map register number
ADP\$_W_MR2NREGAR, ADP\$_W_MR2FREGAR, ADP\$_L_MR2ACTMDR	Updated
ADP\$_L_MR2QBL	Updated
UCB\$_L_FR3	R3 of caller
UCB\$_L_FR4	R4 of caller
UCB\$_L_FPC	00(SP)

synchronization

A driver fork process calls IOC\$REQALTMAP at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment.

DESCRIPTION

A driver fork process calls IOC\$REQALTMAP to allocate a contiguous set of Q22-bus alternate map registers (registers 496 to 8191) to service the DMA transfer described by UCB\$_W_BCNT and UCB\$_W_BOFF. IOC\$REQALTMAP calls IOC\$ALOALTMAP.

If alternate map registers have been permanently allocated to the controller, IOC\$REQALTMAP returns successfully to its caller without allocating map registers. Otherwise, it searches the alternate map register descriptor arrays for the required number of map registers.

IOC\$ALOALTMAP determines the required number of alternate map registers from the contents of UCB\$_W_BOFF and UCB\$_W_BCNT. It allocates one extra map register; this register is marked invalid when the driver fork process subsequently calls IOC\$LOADALTMAP, thus preventing a transfer overrun. If an odd number of map registers is required, IOC\$ALOALTMAP rounds this value up to an even multiple.

If sufficient alternate map registers are available, IOC\$REQALTMAP assigns them to its caller, records the allocation in the ADP and CRB, and returns successfully to its caller.

If IOC\$REQALTMAP cannot allocate a sufficient number of contiguous map registers, it saves process context by placing the contents of R3, R4, and the PC into the UCB fork block and the UCB into the alternate-map-register wait queue (ADP\$_L_MR2QBL). It then returns to its caller's caller.

If the VAX system does not support alternate map registers, IOC\$REQALTMAP exits with SS\$_SSFAIL status.

Operating System Routines

IOC\$REQCOM

IOC\$REQCOM

Completes an I/O operation on a device unit, requests I/O postprocessing of the current request, and starts the next I/O request waiting for the device.

module

IOSUBNPAG

macro

REQCOM

input

Location	Contents
R0	First longword of I/O status.
R1	Second longword of I/O status.
R5	Address of UCB.
UCB\$L_STS	UCB\$V_ERLOGIP set if error logging is in progress.
UCB\$B_ERTCNT	Final error count.
UCB\$B_ERTMAX	Maximum error retry count.
UCB\$L_EMB	Address of error message buffer.
UCB\$L_IRP	Address of IRP.
UCB\$B_DEVCLASS	DC\$_DISK and DC\$_TAPE devices are subject to mount verification checks.
UCB\$L_IOQFL	Device unit's pending-I/O queue.

output

Location	Contents
R0 through R3	Destroyed. Other registers (used by the driver's start-I/O routine) are destroyed if IOC\$INITIATE is called.
IRP\$L_IOST1	First longword of I/O status.
IRP\$L_IOST2	Second longword of I/O status.
UCB\$L_OPCNT	Incremented.
UCB\$L_IOQFL	Updated.
EMB\$W_DV_STS	UCB\$W_STS.
EMB\$B_DV_ERTCNT	UCB\$B_ERTCNT.
EMB\$B_DV_ERTCNT+1	UCB\$B_ERTMAX.
EMB\$Q_DV_IOSB	Quadword of I/O status.
UCB\$L_STS	UCB\$V_BSY and UCB\$V_ERLOGIP cleared.

synchronization

A driver fork process calls IOC\$REQCOM at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment. IOC\$REQCOM transfers control to IOC\$RELCHAN. If the fork process calls IOC\$REQCOM by means of the REQCOM macro (or a JMP instruction), IOC\$RELCHAN returns control to the caller of the driver fork process (for instance, the fork dispatcher).

DESCRIPTION

A driver fork process calls this routine after a device I/O operation and all device-dependent processing of an I/O request is complete.

IOC\$REQCOM performs the following tasks:

- If error logging is in progress for the device (as indicated by UCB\$V_ERLOGIP in UCB\$L_STS), writes into the error message buffer the status of the device unit, the error retry count for the transfer, the maximum error retry count for the driver, and the final status of the I/O operation. It then releases the error message buffer by calling ERL\$RELEASEMB.
- Increments the device unit's operations count (UCB\$L_OPCNT).
- If UCB\$B_DEVCLASS specifies a disk device (DC\$_DISK) or tape device (DC\$_TAPE) and error status is reported, performs a set of checks to determine if mount verification is necessary. Tape end-of-file errors (SS\$_ENDOFFILE) are exempt from these checks. For a tape device with success status, checks to determine if CRC must be generated.
- Writes final I/O status (R0 and R1) into IRP\$L_IOST1 and IRP\$L_IOST2.
- Inserts the IRP in systemwide I/O postprocessing queue.
- Requests a software interrupt from the local processor at IPL\$_IOPOST.
- Attempts to remove an IRP from the device's pending-I/O queue (at UCB\$L_IOQFL). If successful, it transfers control to IOC\$INITIATE to begin driver processing of this I/O request. If the queue is empty, it clears the unit busy bit (UCB\$V_BSY in UCB\$L_STS) to indicate that the device is idle.
- Exits by transferring control to IOC\$RELCHAN.

Operating System Routines

IOC\$REQDATAP, IOC\$REQDATAPNW

IOC\$REQDATAP, IOC\$REQDATAPNW

Request a UNIBUS adapter buffered data path and, optionally, if no path is available, place process in data-path wait queue.

module IOSUBNPAG

macro REQDPR

input

Location	Contents
R5	Address of UCB
00(SP)	Return PC of caller
04(SP)	Return PC of caller's caller
UCB\$L_CRB	Address of CRB
UCB\$L_CRB	Address of CRB
CRB\$L_INTD+ VEC\$L_ADP	Address of ADP
CRB\$L_INTD+ VEC\$B_DATAPATH	Data path specifier; VEC\$V_PATHLOCK set if the data path is permanently allocated to the controller
ADP\$W_DPBITMAP	Data path bit map

output

Location	Contents
R0	SS\$_NORMAL or bit 0 set (indicating error status)
CRB\$L_INTD+ VEC\$B_DATAPATH	Data path specifier
ADP\$W_DPBITMAP	Bit corresponding to allocated data path cleared

synchronization A driver fork process calls IOC\$REQDATAP or IOC\$REQDATAPNW at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment.

DESCRIPTION A driver fork process calls IOC\$REQDATAP or IOC\$REQDATAPNW to request a UNIBUS adapter buffered data path for a DMA transfer.

If a buffered data path is already permanently allocated to the controller, IOC\$REQDATAP or IOC\$REQDATAPNW returns successfully to its caller without allocating a data path. Otherwise, it searches the data path bit map for the first available data path.

Operating System Routines

IOC\$REQDATAP, IOC\$REQDATAPNW

If IOC\$REQDATAP or IOC\$REQDATAPNW locates a free data path, it writes the data path number into CRB\$L_INTD+VEC\$B_DATAPATH, updates the data path bit map (ADP\$W_DPBITMAP), and returns successfully to its caller. If the bit map has been corrupted, the routine issues an INCONSTATE bugcheck.

If IOC\$REQDATAP cannot allocate a data path, it saves process context by placing the contents of R3, R4, and the PC into the UCB fork block and the UCB into the data-path wait queue (ADP\$L_DPQBL). It then returns to its caller's caller. By contrast, if IOC\$REQDATAPNW cannot allocate a data path, it returns immediately to its caller with the low bit in R0 clear, indicating an error.

When called from a driver executing in a VAX system that does not provide buffered data paths, IOC\$REQDATAP and IOC\$REQDATAPNW return control after examining the data path bit map in the ADP.

Operating System Routines

IOC\$REQMAPREG

IOC\$REQMAPREG

Allocates sufficient UNIBUS map registers or a sufficient number of the first 496 Q22-bus map registers to accommodate a DMA transfer and, if unavailable, places process in standard-map-register wait queue.

module IOSUBNPAG

macro REQMPR

input

Location	Contents
R5	Address of UCB
00(SP)	Return PC of caller
04(SP)	Return PC of caller's caller
UCB\$W_BCNT	Transfer byte count
UCB\$W_BOFF	Byte offset in page
UCB\$L_CRB	Address of CRB
CRB\$L_INTD+ VEC\$L_ADP	Address of ADP
CRB\$L_INTD+ VEC\$W_MAPREG	VEC\$V_MAPLOCK set indicates that map registers have been permanently allocated to this controller
ADP\$W_MRNREGARY, ADP\$W_MRFREGARY, ADP\$L_MRACTMDRS	Map register descriptor arrays
ADP\$L_MRQBL	Tail of queue of UCBs waiting for map registers

Operating System Routines

IOC\$REQMAPREG

output

Location	Contents
R0	SS\$_NORMAL
R1	Destroyed
R2	Address of ADP
CRB\$_INTD+ VEC\$_B_NUMREG	Number of map registers allocated
CRB\$_INTD+ VEC\$_W_MAPREG	Starting map register number
ADP\$_W_MRNREGARY, ADP\$_W_MRFREGARY, ADP\$_L_MRACTMDRS	Updated
ADP\$_L_MRQBL	Updated
UCB\$_L_FR3	R3 of caller
UCB\$_L_FR4	R4 of caller
UCB\$_L_FPC	00(SP)

synchronization

A driver fork process calls IOC\$REQMAPREG at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment.

DESCRIPTION

A driver fork process calls IOC\$REQMAPREG to allocate a contiguous set of UNIBUS map registers or a set of the first 496 Q22-bus map registers to service the DMA transfer described by UCB\$_W_BCNT and UCB\$_W_BOFF. IOC\$REQMAPREG calls IOC\$ALOUBAMAP.

If map registers have been permanently allocated to the controller, IOC\$REQMAPREG returns successfully to its caller without allocating map registers. Otherwise, it searches the map register descriptor arrays for the required number of map registers.

IOC\$ALOUBAMAP determines the required number of map registers from the contents of UCB\$_W_BOFF and UCB\$_W_BCNT. It allocates one extra map register; this register is marked invalid when the driver fork process subsequently calls IOC\$LOADUBAMAP, thus preventing a transfer overrun. If an odd number of map registers is required, IOC\$ALOUBAMAP rounds this value up to an even multiple.

If sufficient map registers are available, IOC\$REQMAPREG assigns them to its caller, records the allocation in the ADP and CRB, and returns successfully to its caller.

If IOC\$REQMAPREG cannot allocate a sufficient number of contiguous map registers, it saves process context by placing the contents of R3, R4, and the PC into the UCB fork block and R5 into the standard-map-register wait queue (ADP\$_L_MRQBL). It then returns to its caller's caller.

Operating System Routines

IOC\$REQPCHANH, IOC\$REQPCHANL, IOC\$REQSCHANH, IOC\$REQSCHANL

IOC\$REQPCHANH, IOC\$REQPCHANL, IOC\$REQSCHANH, IOC\$REQSCHANL

Request a controller's primary or secondary data channel and, if unavailable, place process in channel wait queue.

module IOSUBNPAG

macro REQCHAN, REQCHAN

input

Location	Contents
R5	Address of UCB
00(SP)	Return PC of caller
04(SP)	Return PC of caller's caller
UCB\$L_CRB	Address of CRB
CRB\$L_LINK	Address of secondary CRB (IOC\$REQSCHANH and IOC\$REQSCHANL only)
CRB\$B_MASK	CRB\$V_BSY set if the channel is busy
CRB\$L_INTD+VEC\$L_IDB	Address of IDB
CRB\$L_WQFL	Head of queue of UCBs waiting for the controller channel
CRB\$L_WQBL	Tail of queue of UCBs waiting for the controller channel
IDB\$L_CSR	Address of device CSR

output

Location	Contents
R0, R1, R2	Destroyed
R4	Address of device CSR
IDB\$L_OWNER	Address of UCB
CRB\$L_WQFL	Updated
CRB\$L_WQBL	Updated

synchronization

A driver fork process calls IOC\$REQPCHANH, IOC\$REQPCHANL, IOC\$REQSCHANH, or IOC\$REQSCHANL holding the corresponding fork lock in a VMS multiprocessing environment.

Operating System Routines

IOC\$REQPCHANH, IOC\$REQPCHANL, IOC\$REQSCHANH, IOC\$REQSCHANL

DESCRIPTION

A driver fork process calls IOC\$REQPCHANH or IOC\$REQPCHANL to acquire ownership of the primary controller's data channel; it calls IOC\$REQSCHANH or IOC\$REQSCHANL to request the secondary controller's data channel (for instance, the MASSBUS adapter's controller data channel).

Each routine examines CRB\$V_BSY in CRB\$B_MASK. If the selected controller's data channel is idle, the routine grants the channel to the fork process, placing its UCB address in IDB\$L_OWNER and returning successfully with the device's CSR address in R4.

If the data channel is busy, the routine saves process context by placing the contents of R3 and the PC into the UCB fork block. (Note that IOC\$RELCHAN moves the contents of IDB\$L_CSR into R4 before resuming execution of a waiting fork process.) IOC\$REQPCHANH and IOC\$REQSCHANH then insert the UCB at the head of the channel wait queue (CRB\$L_WQFL); IOC\$REQPCHANL and IOC\$REQSCHANL insert the UCB at the tail of the queue (CRB\$L_WQBL). Finally, the routine returns control to its caller's caller.

Operating System Routines

IOC\$RETURN

IOC\$RETURN

Returns to its caller.

module

None.

input

None.

output

None.

synchronization

IOC\$RETURN executes at its caller's IPL and returns control to the caller at that IPL.

DESCRIPTION

IOC\$RETURN is a universal executive routine vector in the fixed portion of the VMS executive. It contains a single RSB instruction. When a driver invokes the DDTAB macro, the macro writes the address of IOC\$RETURN into routine address fields of the DDT that are not supplied in the macro invocation.

IOC\$VERIFYCHAN

Verifies an I/O channel number and translates it to a CCB address.

module IOSUBPAGD

input

Location	Contents
R0	Channel number (in low word)
CTL\$GL_CCBBASE	Base address of process CCB table
CCB\$B_AMOD	Access mode (plus 1) of process owning the channel

output

Location	Contents
R0	SS\$_NORMAL, SS\$_IVCHAN, or SS\$_NOPRIV
R1	Address of CCB
R2	Channel index number
R3	Destroyed

synchronization Because IOC\$VERIFYCHAN gains access to information stored in user process virtual address space, it should only be called from code originating at IPL\$_ASTDEL or below.

DESCRIPTION

Drivers call IOC\$VERIFYCHAN to validate a user-supplied channel number, construct a channel index, and obtain the address of the CCB to which the channel number points.

If the channel number is invalid or zero, or if the channel is unowned, IOC\$VERIFYCHAN returns SS\$_IVCHAN status to its caller.

If the access mode of the current process is less privileged than that indicated in CCB\$B_AMOD, IOC\$VERIFYCHAN returns SS\$_NORMAL!SS\$_NOPRIV status to its caller with the address of the CCB in R1.

Otherwise, IOC\$VERIFYCHAN returns successfully to its caller with the address of the CCB in R1.

Operating System Routines

IOC\$WFIKPCH, IOC\$WFIRLCH

IOC\$WFIKPCH, IOC\$WFIRLCH

Suspend a driver fork thread and fold its context into a fork block in anticipation of a device interrupt or timeout.

module

IOSUBNPAG

macro

WFIKPCH, WFIRLCH

input

Location	Contents
R3, R4	(Preserved)
R5	Address of UCB
R5	Address of UCB
00(SP)	Address following the JSB to IOC\$WFIKPCH or IOC\$WFIRLCH
04(SP)	Timeout value in seconds
08(SP)	IPL to which to lower before returning to the caller's caller
12(SP)	Return PC of caller's caller
EXE\$GL_ABSTIM	Absolute time

output

Location	Contents
UCB\$L_DUETIM	Sum of timeout value and EXE\$GL_ABSTIM
UCB\$V_INT	Set to indicate that interrupts are expected on the device
UCB\$V_TIM	Set to indicate device I/O is being timed
UCB\$V_TIMEOUT	Cleared to indicate that unit is not timed out
UCB\$L_FR3	R3
UCB\$L_FR4	R4
UCB\$L_FPC	00(SP)+2

synchronization

When it is called, IOC\$WFIKPCH or IOC\$WFIRLCH assumes that the local processor has obtained the appropriate synchronization with the device database:

- In a *uniprocessing* environment, the processor must be executing at device IPL or above.
- In a *multiprocessing* environment, the processor must own the appropriate device lock, as recorded in the unit control block (UCB\$L_DLCK) of the device unit from which the interrupt is expected. This requirement also presumes that the local processor is executing at the device IPL associated with the lock.

Operating System Routines

IOC\$WFIKPCH, IOC\$WFIRLCH

Before exiting, IOC\$WFIKPCH or IOC\$WFIRLCH achieves the following synchronization:

- In a *uniprocessing* environment, it lowers the local processor's IPL to the IPL saved on the stack.
- In a *multiprocessing* environment, it conditionally releases the device lock, so that if the caller of the driver fork thread (the caller's caller) previously owned the device lock, it will continue to hold it when the routine exits. IOC\$WFIKPCH or IOC\$WFIRLCH also lowers the local processor's IPL to the IPL saved on the stack.

DESCRIPTION

A driver fork process calls IOC\$WFIKPCH to wait for an interrupt while keeping ownership of the controller's data channel; IOC\$WFIRLCH, by contrast, releases the channel.

Either routine performs the following operations:

- Adds 2 to the address on the top of the stack to determine the address of the next instruction in the driver fork thread after the invocation of the WFIKPCH or WFIRLCH macro. (Note that the macro places the relative offset to the timeout handling routine in the word following the JSB to IOC\$WFIKPCH or IOC\$WFIRLCH.) It pops this address into the UCB fork block (UCB\$L_FPC) so that the driver's interrupt service routine can resume execution of the driver fork thread with a JSB instruction.
- Moves contents of R3 and R4 into the UCB fork block.
- Sets UCB\$V_INT to indicate an expected interrupt from the device unit.
- Sets UCB\$V_TIM to indicate that VMS should check for timeouts from the device unit.
- Determines the timeout due time from the timeout value, now at the top of the stack, and EXE\$GL_ABSTIM, and stores the result in UCB\$L_DUETIM.
- Clears UCB\$V_TIMEOUT to indicate that the unit has not timed out.
- In a multiprocessing environment, issues a DEVICEUNLOCK to conditionally release the device lock associated with the device unit and to lower IPL to the IPL saved on the stack. These actions presume that the DEVICELOCK macro has been issued prior to the wait-for-interrupt invocation.
- Returns to the caller of the driver fork thread (that is, its caller's caller) whose address is now at the top of the stack.

In the course of processing, IOC\$WFIKPCH or IOC\$WFIRLCH explicitly removes the longwords at 00(SP) through 08(SP) from the stack and implicitly removes the longword at 12(SP) by exiting with an RSB instruction.

Operating System Routines

IOC\$WFIKPCH, IOC\$WFIRLCH

Note that IOC\$WFIRLCH exits by transferring control to IOC\$RELCHAN. IOC\$RELCHAN releases the controller data channel and executes the RSB instruction. Because the release of the channel occurs at fork IPL, an interrupt service routine cannot reliably distinguish between operations initiated by IOC\$WFIKPCH and IOC\$WFIRLCH by examining the ownership of the CRB.

LDR\$ALLOC_PT

Allocates the specified number of system page-table entries (SPTEs).

module

PTALLOC

input

Location	Contents
R2	Number of SPTEs to be allocated
LDR\$GL_SPTBASE	Base of system page table
LDR\$GL_FREE_PT	Offset to first free SPTE

output

Location	Contents
R0	SS\$_NORMAL, SS\$_INSFSPTS, or SS\$_BADPARAM
R1	Address of first allocated SPTE
R2	Number of allocated system page-table entries

synchronization

Because LDR\$ALLOC_PT executes at IPL\$_SYNCH and obtains the MMG spin lock in a VMS multiprocessing environment, its caller cannot be executing above IPL\$_SYNCH or hold any higher ranked spin locks. (For instance, a driver fork process executing at IPL\$_SYNCH holding the IOLOCK8 fork lock can call LDR\$ALLOC_PT.) LDR\$ALLOC_PT returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

LDR\$ALLOC_PT allocates the number of system page-table entries (SPTEs) specified in R2. LDR\$ALLOC_PT adjusts the pool of free SPTEs to reflect the allocation of the SPTEs.

A generic VAXBI device driver calls LDR\$ALLOC_PT if it must map the device's node window space. It is the caller's responsibility to fill in each allocated SPTE with a page-frame number (PFN), set its valid bit, and otherwise initialize it.

If R2 contains a zero, LDR\$ALLOC_PT returns SS\$_BADPARAM status in R0 and clears R1. If there are no free SPTEs, it returns SS\$_INSFSPTS status to its caller.

Operating System Routines

LDR\$DEALLOC_PT

LDR\$DEALLOC_PT

Deallocates the specified system page-table entries (SPTEs).

module

PTALLOC

input

Location	Contents
R1	Address of first SPTe to be deallocated
R2	Number of SPTEs to be deallocated
LDR\$GL_SPTBASE	Base of system page table
LDR\$GL_FREE_PT	Offset to first free SPTe

output

Location	Contents
R0	SS\$_NORMAL, SS\$_BADPARAM, or LOADER\$_PTE_NOT_EMPTY
R1	Address of first allocated SPTe
R2	Destroyed

synchronization

Because LDR\$DEALLOC_PT executes at IPL\$_SYNCH and obtains the MMG spin lock in a VMS multiprocessing environment, its caller cannot be executing above IPL\$_SYNCH or hold any higher ranked spin locks. (For instance, a driver fork process executing at IPL\$_SYNCH holding the IOLOCK8 fork lock can call LDR\$DEALLOC_PT.) LDR\$DEALLOC_PT returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

LDR\$DEALLOC_PT deallocates the number of system page-table entries (SPTEs) specified in R2, starting at the one indicated by the contents of R1. LDR\$DEALLOC_PT adjusts the pool of free SPTEs to reflect the addition of the deallocated SPTEs.

If R2 contains a zero, LDR\$DEALLOC_PT returns SS\$_BADPARAM status in R0 and clears R1.

It is the caller's responsibility to ensure that the SPTEs to be deallocated are empty.⁵ If they are not, LDR\$DEALLOC_PT returns LOADER\$_PTE_NOT_EMPTY status in R0.

⁵ Modifications to valid SPTEs require that these SPTEs be flushed from the system's translation buffers. See the description of the INVALIDATE_TB macro in Chapter 2.

MMG\$UNLOCK

Unlocks process pages previously locked for a direct-I/O operation.

module

IOLOCK

input

Location	Contents
R1	Number of buffer pages to unlock
R3	System virtual address of PTE for the first buffer page

output

None.

synchronization

Because MMG\$UNLOCK raises IPL to IPL\$_SYNCH, and obtains the MMG spin lock in a VMS multiprocessing environment, its caller cannot be executing above IPL\$_SYNCH or hold any higher ranked spin locks. MMG\$UNLOCK returns control to its caller at the caller's IPL. The caller retains any spin locks it held at the time of the call.

DESCRIPTION

Drivers rarely use MMG\$UNLOCK. At the completion of a direct-I/O transfer, IOC\$IOPOST automatically unlocks the pages of both the user buffer and any additional buffers specified in region 1 (if defined) and region 2 (if defined) for all the IRPEs linked to the packet undergoing completion processing.

However, driver FDT routines do use MMG\$UNLOCK when an attempt to lock IRPE buffers for a direct-I/O transfer fails. The buffer-locking routines called by such a driver—EXE\$READLOCKR, EXE\$WRITELOCKR, and EXE\$MODIFYLOCKR—all perform coroutine calls back to the driver if an error occurs. When called as a coroutine, the driver must unlock all previously locked regions using MMG\$UNLOCK, and deallocate the IRPE (using EXE\$DEANONPAGED), before returning to the buffer-locking routine.

Operating System Routines

SMP\$ACQNOIPL

SMP\$ACQNOIPL

Acquires a device lock, assuming the local processor is already running at the IPL appropriate for acquisition of the lock.

module

SPINLOCKS

macro

DEVICELock

input

Location

Contents

R0

Address of device lock

output

Location

Contents

R0

Address of device lock

synchronization

Upon entry, the local processor *must* be executing at the synchronization IPL of the device lock, as it is, for instance, when responding to a device interrupt.

SMP\$ACQNOIPL exits with the IPL unchanged and the device lock held.

DESCRIPTION

The DEVICELock macro calls SMP\$ACQNOIPL when NOSETIPL is specified as its **condition** argument.

SMP\$ACQNOIPL attempts to acquire the requested device lock, allowing the acquisition to succeed if the local processor already holds the lock or if the lock is unowned.

If the lock is unowned, the routine increments by 1 a counter that records the acquisition level. Each additional (or nested) acquisition of this lock by the owning processor again increments this counter.

If the lock is owned by another processor, the local processor spin waits until the lock is released.

SMP\$ACQUIRE

Acquires a fork lock or spin lock and enforces the appropriate IPL synchronization on the local processor.

module

SPINLOCKS

macro

FORKLOCK, LOCK

input

Location	Contents
R0	Fork lock or spin lock index

output

Location	Contents
R0	Fork lock or spin lock index

synchronization

When calling SMP\$ACQUIRE, the local processor should be executing at an IPL less than or equal to the synchronization IPL of the lock. The routine, if necessary, immediately raises IPL to the synchronization IPL of the lock. Violations of IPL synchronization in a full-checking multiprocessing environment result in a SPLIPLHIGH bugcheck.

In a full-checking multiprocessing environment, if it must spin wait for the requested lock to be released by another processor, SMP\$ACQUIRE temporarily restores the original IPL for the duration of the wait. If the original IPL was less than IPL\$_RESCHED, the spin wait occurs at IPL\$_RESCHED.

SMP\$ACQUIRE exits with IPL at the synchronization IPL of the lock and the fork lock or spin lock held.

DESCRIPTION

The FORKLOCK and LOCK macros call SMP\$ACQUIRE.

In a full-checking multiprocessing environment, SMP\$ACQUIRE, having ensured that IPL has been set to the lock's synchronization IPL, verifies that the local processor does not currently hold any higher-ranked locks. If a higher-ranked lock is held, SMP\$ACQUIRE issues an SPLACQERR bugcheck.

Otherwise SMP\$ACQUIRE attempts to acquire the requested lock, allowing the acquisition to succeed if the local processor already holds the lock or if the lock is unowned.

Operating System Routines

SMP\$ACQUIRE

If the lock is unowned, the routine increments by 1 a counter that records the acquisition level. Each additional (or nested) acquisition of this lock by the owning processor again increments this counter.

If the lock is owned by another processor, the local processor spin waits until the lock is released.

SMP\$ACQUIREL

Acquires a device lock and enforces the appropriate IPL synchronization on the local processor.

module

SPINLOCKS

macro

DEVICELOCK

input

Location	Contents
R0	Address of device lock

output

Location	Contents
R0	Address of device lock

synchronization

When calling SMP\$ACQUIREL, the local processor should be executing at an IPL less than or equal to the synchronization IPL of the device lock. The routine, if necessary, immediately raises IPL to the synchronization IPL of the device lock. Violations of IPL synchronization result in a SPLIPLHIGH bugcheck if full-checking multiprocessing is enabled.

In a full-checking multiprocessing environment, if it must spin wait for the requested lock to be released by another processor, SMP\$ACQUIREL temporarily restores the original IPL for the duration of the wait. If the original IPL was less than IPL\$_RESCHED, the spin wait occurs at IPL\$_RESCHED. SMP\$ACQUIREL exits with IPL at the device lock's synchronization IPL and the device lock held.

DESCRIPTION

The DEVICELOCK macro calls SMP\$ACQUIREL when NOSETIPL is *not* specified as its **condition** argument.

SMP\$ACQUIREL, having ensured that IPL has been set to the device lock's synchronization IPL, attempts to acquire the requested device lock, allowing the acquisition to succeed if the local processor already holds the lock or if the lock is unowned.

If the lock is unowned, the routine increments by 1 a counter that records the acquisition level. Each additional (or nested) acquisition of this lock by the owning processor again increments this counter.

If the lock is owned by another processor, the local processor spin waits until the lock is released.

Operating System Routines

SMP\$RELEASE

SMP\$RELEASE

Releases all acquisitions of a fork lock or spin lock by the local processor and makes the lock available for acquisition by other processors.

module

SPINLOCKS

macro

FORKUNLOCK, UNLOCK

input

Location	Contents
R0	Fork lock or spin lock index

output

Location	Contents
R0	Fork lock or spin lock index

synchronization

Upon entry, the local processor must be executing at or above the IPL at which the lock was originally obtained. This IPL must be greater than IPL\$_ASTDEL. Violations of IPL synchronization in a full-checking multiprocessing environment result in a SPLIPLLOW bugcheck. At exit, IPL is unchanged and the lock is released.

DESCRIPTION

The FORKUNLOCK and UNLOCK macros call SMP\$RELEASE when the **condition=RESTORE** argument is not specified.

SMP\$RELEASE first verifies that the local processor owns the specified lock. If this is not the case, the procedure issues an SPLRELERR bugcheck. Otherwise, SMP\$RELEASE initializes the ownership count of the lock and releases the lock.

SMP\$RELEASEL

Releases all acquisitions of a device lock by the local processor and makes the lock available for acquisition by other processors.

module

SPINLOCKS

macro

DEVICEUNLOCK

input

Location	Contents
R0	Address of device lock

output

Location	Contents
R0	Address of device lock

synchronization

Upon entry, the local processor must be executing at or above the IPL at which the device lock was originally obtained. This IPL must be greater than IPL\$_ASTDEL. Violations of IPL synchronization in a full-checking multiprocessing environment result in a SPLIPLLOW bugcheck. At exit, IPL is unchanged and the device lock is released.

DESCRIPTION

The DEVICEUNLOCK macro calls SMP\$RELEASEL when the **condition=RESTORE** argument is not specified.

SMP\$RELEASEL first verifies that the local processor owns the specified device lock. If this is not the case, the procedure issues an SPLRELERR bugcheck. Otherwise, SMP\$RELEASEL initializes the ownership count of the device lock and releases the lock.

Operating System Routines

SMP\$RESTORE

SMP\$RESTORE

Releases a single acquisition of a fork lock or spin lock held by the local processor.

module

SPINLOCKS

macro

FORKUNLOCK, UNLOCK

input

Location	Contents
R0	Fork lock or spin lock index

output

Location	Contents
R0	Fork lock or spin lock index

synchronization

Upon entry, the local processor must be executing at or above the IPL at which the lock was originally obtained. This IPL must be greater than IPL\$_ASTDEL. Violations of IPL synchronization in a full-checking multiprocessing environment result in a SPLIPLLOW bugcheck. At exit, IPL is unchanged and the lock may or may not be still held.

DESCRIPTION

The FORKUNLOCK and UNLOCK macros call SMP\$RESTORE when RESTORE is specified as the **condition** argument.

SMP\$RESTORE first verifies that the local processor owns the specified lock. If this is not the case, the procedure issues an SPLRSTERR bugcheck. Otherwise, SMP\$RESTORE proceeds to decrement the ownership count of the lock. If the ownership count of the lock drops to its initial state, the procedure releases the lock and makes it available to other processors.

SMP\$RESTOREL

Releases a single acquisition of a device lock held by the local processor.

module

SPINLOCKS

macro

DEVICEUNLOCK

input

Location	Contents
R0	Address of device lock

output

Location	Contents
R0	Address of device lock

synchronization

Upon entry, the local processor must be executing at or above the IPL at which the device lock was originally obtained. This IPL must be greater than IPL\$_ASTDEL. Violations of IPL synchronization in a full-checking multiprocessing environment result in a SPLIPLLOW bugcheck. At exit, IPL is unchanged and the device lock may or may not be still held.

DESCRIPTION

The DEVICEUNLOCK macro calls SMP\$RESTOREL when RESTORE is specified as its **condition** argument.

SMP\$RESTOREL first verifies that the local processor owns the specified device lock. If this is not the case, the procedure issues an SPLRSTERR bugcheck. Otherwise, SMP\$RESTOREL proceeds to decrement the ownership count of the device lock. If the ownership count of the device lock drops to its initial state, the procedure releases the lock and makes it available to other processors.

4 Device Driver Entry Points

This chapter describes the standard driver routines and their environment that VMS uses as entry points in a device driver program. The standard entry routines are:

- Alternate start-I/O
- Cancel-I/O
- Cloned UCB
- Controller initialization
- Driver unloading
- FDT
- Interrupt service
- Register dumping
- Start-I/O
- Timeout handling
- Unit delivery
- Unit initialization
- Unsolicited interrupt service

Device Driver Entry Points

Alternate Start-I/O Routine

Alternate Start-I/O Routine

Initiates activity on a device that can support multiple, concurrent I/O operations and synchronizes access to its UCB.

specified in Specify the address of the alternate start-I/O routine in the **altstart** argument to the DDTAB macro. This macro places the address into DDT\$L_ALTSTART.

called by Called by routine EXE\$ALTQUEPKT in module SYSQIOREQ. A driver FDT routine generally is the caller of EXE\$ALTQUEPKT.

synchronization An alternate start-I/O routine begins execution at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment. It must return control to its EXE\$ALTQUEPKT in this context.

context Because an alternate start-I/O routine gains control in fork process context, it can access only those virtual addresses that are in system (S0) space.

register usage An alternate start-I/O routine must preserve the contents of all registers except R0 through R5.

input

Location	Contents
R3	Address of IRP
R5	Address of UCB

exit The alternate start-I/O routine completes I/O requests by calling the routine COM\$POST. This routine places each IRP in the I/O postprocessing queue and returns control to the driver. The driver can then fetch another IRP from an internal queue. If no IRPs remain, the driver returns control to EXE\$ALTQUEPKT, which relinquishes fork level synchronization and returns to the driver FDT routine that called it. The FDT routine performs any postprocessing and transfers control to the routine EXE\$QIORETURN.

DESCRIPTION

An alternate start-I/O routine initiates requests for activity on a device that can process two or more I/O requests simultaneously. Because the method by which the alternate start-I/O routine is invoked bypasses the unit's pending-I/O queue (UCB\$L_IOQFL) and the device busy flag (UCB\$V_BSY in UCB\$L_STS), the routine is activated regardless of whether the device unit is busy with another request.

Device Driver Entry Points

Alternate Start-I/O Routine

As a result, the driver that incorporates an alternate start-I/O routine must use its own internal I/O queues (in a UCB extension, for instance) and maintain synchronization with the unit's pending-I/O queue. In addition, if the routine processes more than one IRP at a time, it must employ separate fork blocks for each request.

Device Driver Entry Points

Cancel-I/O Routine

Cancel-I/O Routine

Prevents further device-specific processing of the I/O request currently being processed on a device.

specified in

Supply the address of the cancel-I/O routine in the **cancel** argument of the DDTAB macro. The macro places this address into DDT\$L_CANCEL. Many drivers specify the system routine IOC\$CANCELIO as their cancel-I/O routine.

called by

VMS routines call a driver's cancel-I/O routine under the following circumstances:

- When a process issues a Cancel-I/O-on-Channel system service (\$CANCEL)
- When a process deallocates a device, causing the device's reference count (UCB\$W_REFC) to become zero (that is, no process I/O channels are assigned to the device)
- When a process deassigns a channel from a device, using the \$DASSGN system service
- When the command interpreter performs cleanup operations as part of image termination by canceling all pending I/O requests for the image and closing all image-related files open on process I/O channels

synchronization

A cancel-I/O routine begins execution at fork IPL, holding the corresponding fork lock in a VMS multiprocessing environment. It must return control to its caller in this context.

context

A cancel-I/O routine executes in kernel mode in process context.

register usage

A cancel-I/O routine must preserve the contents of all registers except R4 and R5.

Device Driver Entry Points

Cancel-I/O Routine

input

Location	Contents
R2	Channel index number
R3	Contents of UCB\$L_IRP (address of current IRP, if any, for device)
R4	Address of PCB of the process for which the I/O request is being canceled
R5	Address of UCB
R8	Reason for cancellation, one of the following: CAN\$C_CANCEL Called by \$CANCEL system service CAN\$C_DASSGN Called by \$DASSGN or \$DALLOC system service

exit

The cancel-I/O routine issues an RSB instruction to return to its caller.

DESCRIPTION

A driver's cancel-I/O routine must perform the following tasks:

- 1 Confirm that the device is busy by examining the device-busy bit in the UCB status longword (UCB\$V_BSY in UCB\$L_STS).
- 2 Confirm that the PID of the request the device is servicing (IRP\$L_PID) matches that of the process requesting the cancellation (PCB\$L_PID).
- 3 Confirm that the channel-index number of the request the device is servicing (IRP\$W_CHAN) matches that specified in the cancel-I/O request.
- 4 Cause to be completed (canceled) as quickly as possible all active I/O requests on the specified channel that were made by the process that has requested the cancellation. The cancel-I/O routine usually accomplishes this by setting UCB\$V_CANCEL in the UCB\$L_STS. When the next interrupt or timeout occurs for the device, the driver's start-I/O routine detects the presence of an active but canceled I/O request by testing this bit and takes appropriate action, such as completing the request without initiating any further device activity. Other driver routines, such as the timeout handling routine, check the cancel-I/O bit to determine whether to retry the I/O operation or abort it.

Device Driver Entry Points

Cloned UCB Routine

Cloned UCB Routine

Performs device-specific initialization and verification of a cloned UCB.

specified in

Specify the address of a cloned UCB routine in the **cloneducb** argument of the DDTAB macro. The macro places this address into DDT\$L_CLONEDUCB. Only drivers for template devices, such as mailboxes, specify a cloned UCB routine.

called by

EXE\$ASSIGN calls the driver's cloned UCB routine when an Assign I/O Channel system service request (\$ASSIGN) specifies a template device (that is, bit UCB\$V_TEMPLATE in UCB\$L_STS is set).

synchronization

A cloned UCB routine executes at IPL\$_ASTDEL, holding the I/O database mutex (IOC\$GL_MUTEX).

context

A cloned UCB routine executes in kernel mode in process context.

register usage

A cloned UCB routine must preserve the contents of R2 and R4.

input

Location	Contents
R0	SS\$_NORMAL
R2	Address of cloned UCB
R3	Address of DDT
R4	Address of current PCB
R5	Address of template UCB
UCB\$L_FQFL(R2)	Address of UCB\$L_FQFL(R2)
UCB\$L_FQBL(R2)	Address of UCB\$L_FQFL(R2)
UCB\$L_FPC(R2)	0
UCB\$L_FR3(R2)	0
UCB\$L_FR4(R2)	0
UCB\$W_BUFQUO(R2)	0
UCB\$L_ORB(R2)	Address of cloned ORB
UCB\$L_LINK(R2)	Address of next UCB in DDB chain
UCB\$L_IOQFL(R2)	Address of UCB\$L_IOQFL(R2)
UCB\$L_IOQBL(R2)	Address of UCB\$L_IOQFL(R2)
UCB\$W_UNIT(R2)	Device unit number
UCB\$W_CHARGE(R2)	Mailbox byte quota charge (UCB\$W_SIZE)
UCB\$W_REFC(R2)	0

Device Driver Entry Points

Cloned UCB Routine

Location	Contents
UCB\$L_STS(R2)	UCB\$V_DELETEUCB set, UCB\$V_ONLINE set
UCB\$W_DEVSTS(R2)	UCB\$V_DELMBX set if DEV\$V_MBX is set in UCB\$L_DEVCHAR(R2)
UCB\$L_OPCNT(R2)	0
UCB\$L_SVAPTE(R2)	0
UCB\$W_BOFF(R2)	0
UCB\$W_BCNT(R2)	0
UCB\$L_ORB(R2)	Address of cloned ORB
ORB\$L_OWNER of template ORB	UIC of current process
ORB\$L_ACL_MUTEX of template ORB	FFFF ₁₆
ORB\$B_FLAGS of template ORB	ORB\$V_PROT ₁₆ set
ORB\$W_PROT of template ORB	0
ORB\$L_ACL_COUNT of template ORB	0
ORB\$L_ACL_DESC of template ORB	0
ORB\$R_MIN_CLASS of template ORB	0 in first longword

exit

A cloned UCB routine issues an RSB instruction to return control to EXE\$ASSIGN. If the routine returns error status in R0, EXE\$ASSIGN undoes the process of UCB cloning and completes with failure status in R0.

DESCRIPTION

When a process requests that a channel be assigned to a template device, EXE\$ASSIGN does not assign the channel to the template device itself. Rather, it creates a copy of the template device's UCB and ORB, initializing and clearing certain fields as appropriate.

The driver's cloned UCB routine verifies the contents of these fields and completes their initialization.

Device Driver Entry Points

Controller Initialization Routine

Controller Initialization Routine

Prepares a controller for operation.

specified in

Use the `DPT_STORE` macro to place the address of the controller initialization routine into `CRB$L_INTD+VEC$L_INITIAL`.

called by

`SYSGEN` calls a driver's controller initialization routine when processing a `CONNECT` command. Also, `VMS` calls this routine if the device, controller, processor, or adapter to which the device is connected experiences a power failure.

synchronization

`VMS` calls a controller initialization routine at `IPL$POWER`. If it must lower `IPL`, the controller initialization routine cannot explicitly do so. Rather, it must fork. Because `SYSGEN` calls the unit initialization routine immediately after the controller initialization returns control to it, the driver's initialization routines must synchronize their activities. If the controller initialization routine forks, the unit initialization routine must be prepared to execute before the controller initialization routine completes.

The portion of the controller initialization that services power failure cannot acquire any spin locks. As a result, the routine cannot fork to perform power failure servicing.

context

Because a controller initialization routine executes within system context, it can refer only to those virtual addresses that reside in system (`S0`) space.

register usage

A controller initialization routine must preserve the contents of all registers except `R0`, `R1`, and `R2`.

input

Location	Contents
<code>R4</code>	Address of device's CSR
<code>R5</code>	Address of IDB associated with the controller
<code>R6</code>	Address of DDB associated with the controller
<code>R8</code>	Address of controller's CRB

exit

The controller initialization routine returns control to its caller with an `RSB` instruction.

Device Driver Entry Points Controller Initialization Routine

DESCRIPTION

Some controllers require initialization when the system's driver-loading routine loads the driver and when the system is recovering from a power failure. Depending on the device, a controller initialization routine performs any and all of the following actions:

- Determine whether it is being called as a result of a power failure by examining the power bit (UCB\$V_POWER in UCB\$L_STS) in the UCB. A controller initialization routine may want to perform or avoid specific tasks when servicing a power failure.
- Clear error-status bits in device registers.
- Enable controller interrupts.
- Allocate resources that must be permanently allocated to the controller.
- If the controller is dedicated to a single-unit device, such as a printer, fill in IDB\$L_OWNER and set the online bit (UCB\$V_ONLINE in UCB\$L_STS).
- For generic VAXBI devices, initialize BIIC and device hardware.

Device Driver Entry Points

Driver Unloading Routine

Driver Unloading Routine

A driver specifies a driver unloading routine if there is any device-specific work to do when the driver is unloaded and reloaded.

specified in Specify the address of the driver unloading routine in the **unload** argument of the DPTAB macro. The driver-loading procedure puts the relative address of this routine in DPT\$W_UNLOAD.

called by SYSGEN calls the driver unloading routine, if it exists, when executing a RELOAD command.

synchronization SYSGEN calls a driver unloading routine at IPL\$_POWER. The driver unloading routine cannot lower IPL.

context The driver unloading routine executes in process context.

register usage The driver unloading routine can use all registers.

input

Location	Contents
R6	Address of DDB
R10	Address of DPT

exit The driver unloading routine returns exits with an RSB instruction. If it returns a success code (bit 0 set) in R0, SYSGEN proceeds to load the new version of the driver. If it returns a failure code (bit 0 clear), SYSGEN neither unloads the old version of the driver nor loads the new version.

DESCRIPTION Because the driver unloading routine cannot lower IPL from IPL\$_POWER or obtain spin locks, it is of limited usefulness. It cannot safely modify I/O database fields, but can use COM\$DRVDEALMEM to return system buffers allocated by the driver to nonpaged pool.

FDT Routines

Perform any device-dependent activities needed to prepare the I/O database to process an I/O request.

specified in

Use the FUNCTAB macro to specify the set of FDT routines that preprocess requests for I/O activity of a given type. Specify the names of the routines in the order in which you want them to execute for each type of I/O operation.

called by

The \$QIO system service calls a driver's FDT routines from the module SYSQIOREQ.

synchronization

FDT routines are called at IPL\$_ASTDEL and must exit at IPL\$_ASTDEL. FDT routines must not lower IPL below IPL\$_ASTDEL. If they raise IPL, they must lower it to IPL\$_ASTDEL before passing control to any other code. Similarly, before exiting they must release any spin locks they may acquire in a VMS multiprocessing environment.

context

FDT routines execute in the context of the process that requested the I/O activity. If an FDT routine alters the stack, it must restore the stack before returning control to the caller of the routine.

register usage

FDT routines must preserve the contents of R3 through R8, the AP, and the FP.

input

Location	Contents
R0	Address of FDT routine being called
R3	Address of IRP
R4	Address of PCB of the requesting process
R5	Address of UCB of the device on which I/O activity is requested
R6	Address of CCB that describes the user-specified process-I/O channel
R7	Number of the bit that specifies the code for the requested I/O function
R8	Address of entry in the function decision table that dispatched control to this FDT routine
AP	Address of first function-dependent argument (p1) specified in the \$QIO request

Device Driver Entry Points

FDT Routines

exit

In a set of FDT routines associated with an I/O function, each, except the last, must return control to its caller by means of an RSB instruction. The last must exit using one of the following mechanisms:

Exit Mechanism	Function
JMP EXE\$ABORTIO	Aborts an I/O request and returns status to the caller of the \$QIO system service in R0.
JSB EXE\$ALTQUEPKT	Queues an IRP to the driver's alternate start-I/O routine without checking the status of the device.
JMP EXE\$FINISHIO	Completes the processing of an I/O request, returning status to the caller of the \$QIO system service. (EXE\$FINISHIO takes the status information from R0 and R1 and returns it in the IOSB specified in the call to \$QIO.)
JMP EXE\$FINISHIOC	Completes the I/O processing of an I/O request, returning status to the caller of the \$QIO system service. (EXE\$FINISHIOC takes the status information from R0 and returns it in the IOSB specified in the call to \$QIO, clearing the second longword of the IOSB.)
JMP EXE\$QIODRVPKT	Inserts an IRP into a device's pending-I/O queue if the device is busy, or starts I/O activity if the device is idle.

DESCRIPTION

FDT routines validate the function-dependent arguments to a \$QIO system service request and prepare the I/O database to service the request. For each function that a device supports, a set of FDT routines must provide preprocessing of requests for that function. For a function that does not involve an I/O transfer, a set of FDT routines may complete its processing. Otherwise FDT routines can abort the request, pass it to the next FDT routine in the set, or pass it to a VMS routine that delivers it to the driver.

Interrupt Service Routine

Processes interrupts generated by a device.

specified in

UNIBUS, Q22-bus, and generic VAXBI devices require an interrupt service routine for each interrupt vector the device has. Use the `DPT_STORE` macro to place the address of the interrupt service routine into `CRB$L_INTD+VEC$L_ISR`.

If the device has two interrupt vectors, use the `DPT_STORE` macro to place the address of the second interrupt service routine into `CRB$L_INTD2+VEC$L_ISR`.

Tape devices on the MASSBUS require an interrupt service routine that interrogates the tape formatter (the controller) to determine which drive needs attention and whether the interrupt is unsolicited.

Disk devices on the MASSBUS use the interrupt service routine provided by VMS and do not need to provide their own interrupt service routine.

called by

The interrupt service routine is called either by the VMS interrupt dispatcher (for direct-vector adapters) or by an adapter interrupt service routine (for non-direct-vector adapters).

synchronization

A driver's interrupt service routine is called, executes, and returns at device IPL. In a VMS multiprocessing environment, the interrupt service routine must obtain the device lock associated with its device IPL. It performs this acquisition as soon as it obtains the address of the UCB of the interrupting device. It must release this device lock before dismissing the interrupt.

context

At the execution of a driver's interrupt service routine, the processor is running in kernel mode on the interrupt stack. As a result, an interrupt service routine can reference only those virtual addresses that reside in system (S0) space.

Device Driver Entry Points

Interrupt Service Routine

register usage

If an interrupt service routine uses R6 through R11, the AP, or the FP, it must first save the contents of those registers, restoring their contents before exiting by means of the REI instruction. MASSBUS drivers must also preserve the contents of R0 and R1.

input

Location	Contents
00(SP)	Address of longword that contains the address of the IDB
04(SP) to 24(SP)	<i>For UNIBUS, Q22-bus, and generic VAXBI devices,</i> the contents of R0 through R5 at the time of the interrupt
28(SP)	<i>For UNIBUS, Q22-bus, and generic VAXBI devices,</i> PC at the time of the interrupt
32(SP)	<i>For UNIBUS, Q22-bus, and generic VAXBI devices,</i> PSL at the time of the interrupt
04(SP) to 16(SP)	<i>For MASSBUS devices,</i> the contents of R2 through R5 at the time of the interrupt
20(SP)	<i>For MASSBUS devices,</i> PC at the time of the interrupt
24(SP)	<i>For MASSBUS devices,</i> PSL at the time of the interrupt

exit

Before an interrupt service routine transfers control to the suspended driver, it must restore the contents of R3 and R4 from the UCB. It then transfers control to the address saved in UCB\$L_FPC.

When it regains control (after the suspended driver forks), an interrupt service routine removes the address of the pointer to the IDB from the top of the stack and restores the registers VMS saved when dispatching the interrupt (R0 through R5 for UNIBUS, Q22-bus, and generic VAXBI interrupt service routines, R2 through R5 for MASSBUS interrupt service routines). Finally, an interrupt service routine dismisses the interrupt with an REI instruction.

DESCRIPTION

An interrupt service routine performs the following functions:

- 1 Determines whether the interrupt is expected
- 2 Processes or dismisses unexpected interrupts
- 3 Activates the suspended driver so it can process expected interrupts

For MASSBUS devices, a VMS interrupt service routine performs these functions.

Register Dumping Routine

Copies the contents of a device's registers to an error message buffer or a diagnostic buffer.

specified in	Specify the name of the register dumping routine in the regdmp argument of the DDTAB macro. This macro places the address of the routine into DDT\$L_REGDUMP.								
called by	The VMS error logging routines (ERL\$DEVICERR, ERL\$DEVICTMO, and ERL\$DEVICEATTN) and diagnostic buffer filling routine (IOC\$DIAGBUFILL) call the register dumping routine.								
synchronization	VMS calls a register dumping routine at the same IPL at which the driver called the VMS routine ERL\$DEVICERR, ERL\$DEVICTMO, ERL\$DEVICEATTN, or IOC\$DIAGBUFILL. A register dumping routine must not change IPL.								
context	A register dumping routine executes within the context of an interrupt service routine or a driver fork process, using the kernel-mode stack. As a result, it can only refer to those virtual addresses that reside in system (S0) space.								
register usage	The register dumping routine preserves the contents of all registers except R0 through R2. If it uses the stack, the register dumping routine must restore the stack before passing control to another routine, waiting for an interrupt, or returning control to its caller.								
input	<table><thead><tr><th style="text-align: left;">Location</th><th style="text-align: left;">Contents</th></tr></thead><tbody><tr><td>R0</td><td>Address of buffer into which a register dumping routine copies the contents of device registers</td></tr><tr><td>R4</td><td>Address of device's CSR (if the driver invoked the WFIKPCH macro to wait for an interrupt or timeout)</td></tr><tr><td>R5</td><td>Address of UCB</td></tr></tbody></table>	Location	Contents	R0	Address of buffer into which a register dumping routine copies the contents of device registers	R4	Address of device's CSR (if the driver invoked the WFIKPCH macro to wait for an interrupt or timeout)	R5	Address of UCB
Location	Contents								
R0	Address of buffer into which a register dumping routine copies the contents of device registers								
R4	Address of device's CSR (if the driver invoked the WFIKPCH macro to wait for an interrupt or timeout)								
R5	Address of UCB								
exit	The register dumping routine issues an RSB instruction to return to its caller.								

Device Driver Entry Points

Register Dumping Routine

DESCRIPTION

A register dumping routine fills the indicated buffer as follows:

- 1 Writes a longword value representing the number of device registers to be written into the buffer
- 2 Moves device register longword values into the buffer following the register count longword

Start-I/O Routine

Activates a device to process a requested I/O function.

specified in

Specify the name of the start-I/O routine in the **start** argument of the DDTAB macro. This macro places the address of the routine into DDT\$L_START.

called by

The start-I/O routine is called by IOC\$INITIATE and IOC\$REQCOM in module IOSUBNPAG.

synchronization

A start-I/O routine is placed into execution at fork IPL, holding the associated fork lock in a VMS multiprocessing environment. It must relinquish control of the processor in the same context.

For many devices, the start-I/O routine raises IPL to IPL\$_POWER to check that a power failure has not occurred on the device prior to loading the device's registers. The start-I/O routine initiates device activity at device IPL, after acquiring the corresponding device lock in a VMS multiprocessing environment. An invocation of the WFIKPCH or WFIRLCH macro to wait for a device interrupt releases this device lock.

context

Because a start-I/O routine gains control of the processor in the context of a fork process, it can refer only to those addresses that reside in system (S0) space.

register usage

A start-I/O routine must preserve the contents of all registers except R0, R1, R2, and R4. If the start-I/O routine uses the stack, it must restore the stack before completing the request, waiting for an interrupt, or requesting system resources.

input

Location	Contents
R3	Address of IRP
R5	Address of UCB
UCB\$W_BCNT	Number of bytes to be transferred, copied from the low-order word of IRP\$L_BCNT
UCB\$W_BOFF	Byte offset into first page of direct-I/O transfer; for buffered-I/O transfers, number of bytes to be charged to the process allocating the buffer
UCB\$L_SVAPTE	For a <i>direct-I/O</i> transfer, virtual address of first page-table entry (PTE) of I/O-transfer buffer; for <i>buffered-I/O</i> transfer, address of buffer in system address space

Device Driver Entry Points

Start-I/O Routine

exit

The start-I/O routine suspends itself whenever it must wait for a required resource, such as a controller data channel or UNIBUS/Q22-bus map registers. To do so, it invokes a VMS macro (such as REQCHAN or REQMPR) that saves its context in the UCB fork block, places the UCB in a resource wait queue, and returns control to the caller of the start-I/O routine.

The start-I/O routine also suspends itself when it issues a WFIKPCCH or WFIRLCH macro to initiate device activity. These macros also store the driver's context in the UCB fork block to be restored when the device interrupts or times out.

The start-I/O routine is again suspended if it forks to complete servicing of a device interrupt. The IOFORK macro places driver context in the UCB fork block, inserts the fork block into a processor-specific fork queue, and requests a software interrupt from the processor at the corresponding fork IPL. After issuing the IOFORK macro, the routine issues an RSB instruction, returning control to the driver's interrupt service routine.

The routine completes the processing of an I/O request by invoking the REQCOM macro. In addition to initiating device-independent postprocessing of the current request, the REQCOM macro also attempts to start the next request waiting for a device unit. If there are no waiting requests, the macro returns control to the caller of the start-I/O routine. This is often the VMS fork dispatcher.

DESCRIPTION

A driver's start-I/O routine activates a device and waits for a device interrupt or timeout. After a device interrupt, the driver's interrupt service routine returns control to the start-I/O routine at device IPL, holding the associated device lock in a VMS multiprocessing environment.

The start-I/O routine usually forks at this time to perform various device-dependent postprocessing tasks, and returns control to the interrupt service routine.

Timeout Handling Routine

Takes whatever action is necessary when a device has not yet responded to a request for device activity and the time allowed for a response has expired.

specified in

Specify the address of the timeout handling routine in the **except** argument to the WFIKPCH or the WFIRLCH macro.

called by

The WFIKPCH and WFIRLCH macros use this entry point, but only when the name of a timeout handling routine is provided in their **except** argument. These macros are used in the driver's start-I/O routine; thus, strictly speaking, the driver itself is the only entity that uses this entry point.

Routines in the VMS module TIMESCHDL call the timeout handling routine at the request of the WFIKPCH and WFIRLCH macros.

synchronization

A timeout handling routine is called at device IPL and must return to its caller at device IPL. In a VMS multiprocessing environment, the processor holds both the fork lock and device lock associated with the device at the time of the call.

After taking whatever device-specific action is necessary at device IPL, a timeout handling routine can lower IPL to fork IPL to perform less critical activities. Because its caller restores IPL to fork IPL (and releases the device lock in a VMS multiprocessing environment), if a timeout handling routine does lower IPL, it can do so only by forking or by performing the following steps:

- Issue a DEVICEUNLOCK macro to lower to fork level
- Perform timeout handling activities possible at the lower IPL
- Issue a DEVICELOCK macro to again obtain the device lock and raise to device IPL
- Issue an RSB instruction to return to its caller

context

Because a timeout handling routine executes in the context of a fork process, it can access only those virtual addresses that refer to system (S0) space.

register usage

A timeout handling routine can use R0, R1, and R2 freely, but must preserve the contents of all other registers. If a timeout handling routine uses the stack, it must restore the stack before completing or canceling the current I/O request, waiting for an interrupt, or returning control to its caller.

Device Driver Entry Points

Timeout Handling Routine

input

Location	Contents
R3	Contents of R3 when the last invocation of WFIKPCH or WFIRLCH took place
R4	Contents of R4 when the last invocation of WFIKPCH or WFIRLCH took place
R5	Address of UCB of the device
UCB\$L_STS	UCB\$V_INT and UCB\$V_TIM clear; UCB\$V_TIMOUT set

exit

The timeout handling routine issues an RSB instruction to return to its caller.

DESCRIPTION

There are no outputs required from a timeout handling routine, but, depending on the characteristics of the device, the timeout handling routine might cancel or retry the current I/O request, send a message to the operator, or take some other action.

Before calling a timeout handling routine, VMS places the device in a state in which no interrupt is expected (by clearing the bit UCB\$V_INT in field UCB\$L_STS). If the requested interrupt occurs after this routine is called, it will appear to be an unsolicited interrupt. Many drivers handle this situation by disabling interrupts while the timeout handling routine executes.

Unit Delivery Routine

For controllers that can control a variable number of device units, determines which specific devices are present and available for inclusion in the system's configuration.

specified in	Specify the name of the unit delivery routine in the deliver argument to the DPTAB macro. The macro puts the relative address of this routine in DPT\$W_DELIVER.														
called by	SYSGEN's AUTOCONFIGURE command calls the unit delivery routine once for each unit the controller is capable of controlling. This value is specified in the defunits argument to the DPTAB macro.														
synchronization	The unit delivery routine is called at IPL\$_POWER. It must not lower IPL.														
context	The unit delivery routine executes in the context of the process within which SYSGEN executes.														
register usage	The unit delivery routine can use R0, R1, and R2 freely, but must preserve the contents of all other registers.														
input	<table style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left; width: 30%;">Location</th> <th style="text-align: left;">Contents</th> </tr> </thead> <tbody> <tr> <td>R3</td> <td>Address of IDB; 0 if none exists</td> </tr> <tr> <td>R4</td> <td>Address of device's CSR</td> </tr> <tr> <td>R5</td> <td>Number of unit that the unit delivery routine must decide to configure or not to configure</td> </tr> <tr> <td>R6</td> <td>Address of start of the UNIBUS adapter's or Q22-bus's I/O space (UNIBUS/Q22-bus devices); address of MBA configuration register (MASSBUS devices)</td> </tr> <tr> <td>R7</td> <td>Address of AUTOCONFIGURE command's configuration control block (ACF)</td> </tr> <tr> <td>R8</td> <td>Address of ADP</td> </tr> </tbody> </table>	Location	Contents	R3	Address of IDB; 0 if none exists	R4	Address of device's CSR	R5	Number of unit that the unit delivery routine must decide to configure or not to configure	R6	Address of start of the UNIBUS adapter's or Q22-bus's I/O space (UNIBUS/Q22-bus devices); address of MBA configuration register (MASSBUS devices)	R7	Address of AUTOCONFIGURE command's configuration control block (ACF)	R8	Address of ADP
Location	Contents														
R3	Address of IDB; 0 if none exists														
R4	Address of device's CSR														
R5	Number of unit that the unit delivery routine must decide to configure or not to configure														
R6	Address of start of the UNIBUS adapter's or Q22-bus's I/O space (UNIBUS/Q22-bus devices); address of MBA configuration register (MASSBUS devices)														
R7	Address of AUTOCONFIGURE command's configuration control block (ACF)														
R8	Address of ADP														
exit	A unit delivery routine issues an RSB instruction to return control to the SYSGEN autoconfiguration facility. If the routine returns error status in R0, SYSGEN does not configure the unit.														

DESCRIPTION The unit delivery routine determines which units on a controller should be configured. For instance, a unit delivery routine can prevent the creation of UCBs for devices that do not respond to a test for their presence.

Device Driver Entry Points

Unit Initialization Routine

Unit Initialization Routine

Prepares a device for operation and, in the case of a device on a dedicated controller, initializes the controller.

specified in

You can specify a unit initialization routine in two ways, either of which will suffice for all but a few specific devices.

- Specify the address of the unit initialization routine **unitinit** argument of the DDTAB macro. This macro places the address of the routine into DDT\$L_UNITINIT. MASSBUS device drivers *must* use this method.
- Use the DPT_STORE macro to place the address of the unit initialization routine into CRB\$L_INTD+VEC\$L_UNITINIT.

called by

SYSGEN calls a driver's unit initialization routine when processing a CONNECT command. VMS calls a unit initialization routine when the device, the controller, the processor, or the adapter to which the device is connected undergoes power failure recovery.

synchronization

VMS calls a unit initialization routine at IPL\$ POWER. If it must lower IPL, the controller initialization routine cannot explicitly do so. Rather, it must fork. Because SYSGEN calls the unit initialization routine immediately after the controller initialization returns control to it, the driver's initialization routines must synchronize their activities. If the controller initialization routine forks, the unit initialization routine must be prepared to execute before the controller initialization routine completes.

The portion of the unit initialization that services power failure cannot acquire any spin locks. As a result, the routine cannot fork to perform power failure servicing.

context

Because VMS calls it in system context, a unit initialization routine can only refer to those virtual addresses that reside in system (S0) space.

register usage

A unit initialization routine must preserve the contents of all registers except R0, R1, and R2.

Device Driver Entry Points

Unit Initialization Routine

input

Location	Contents
R3	Address of primary CSR.
R4	Address of secondary CSR, if it exists. (If it does not, the contents of R4 are the same as those of R3.)
R5	Address of UCB.

exit

The unit initialization routine returns control to its caller with an RSB instruction.

DESCRIPTION

Depending on the device, a unit initialization routine performs any or all of the following tasks:

- 1 Determines whether it is being called as a result of a power failure by examining the power bit (UCB\$V_POWER in UCB\$L_STS) in the UCB. A unit initialization routine may want to perform or avoid specific tasks when servicing a power failure.
- 2 Clears error-status bits in device registers.
- 3 Enables controller interrupts.
- 4 Sets the online bit (UCB\$V_ONLINE in UCB\$L_STS).
- 5 Allocates resources that must be permanently allocated to the device or, for some devices, the controller.
- 6 If the device has a dedicated controller, as some printers do, fills in IDB\$L_OWNER.
- 7 For dedicated VAXBI controllers, initializes BIIC and device hardware.
- 8 For multiunit VAXBI controllers, tests for the existence of the unit for which it was called and returns success or failure status to SYSGEN.

Device Driver Entry Points

Unsolicited Interrupt Service Routine

Unsolicited Interrupt Service Routine

Services an interrupt from a MASSBUS disk that is not the result of a driver's request.

specified in	Specify the name of the unsolicited interrupt service routine in the unsolic argument to the DDTAB macro. This macro places the address of the routine into DDT\$L_UN SOLINT.						
called by	The MASSBUS adapter's interrupt service routine (MBA\$INT in module ADPERRSUB of the SYSLOA facility) calls a driver's unsolicited interrupt service routine.						
synchronization	An unsolicited interrupt service routine is called, executes, and returns at device IPL.						
context	Because the unsolicited interrupt service routine executes in kernel mode on the interrupt stack, it can only refer to those addresses that reside in system (S0) space.						
register usage	The unsolicited interrupt service routine must not alter the contents of registers R6 through R11, the AP, or the FP.						
input	<table><thead><tr><th>Location</th><th>Contents</th></tr></thead><tbody><tr><td>R4</td><td>Address of MBA's configuration register</td></tr><tr><td>R5</td><td>Address of UCB</td></tr></tbody></table>	Location	Contents	R4	Address of MBA's configuration register	R5	Address of UCB
Location	Contents						
R4	Address of MBA's configuration register						
R5	Address of UCB						
exit	An unsolicited interrupt service routine issues an RSB instruction to return control to the MASSBUS adapter's interrupt service routine.						

DESCRIPTION Only drivers of MASSBUS disks must provide unsolicited interrupt service routines. All other devices detect unsolicited interrupts in their interrupt service routines.

The routine that handles these unsolicited interrupts must determine the nature of the interrupt and act accordingly, depending on the characteristics of the device and controller. Examples of such unsolicited interrupts include disks being placed on line or taken off line.

Index

A

- ACB\$V_QUOTA • 3-7, 3-10
 - ACB (AST control block) • 1-38, 1-86, 3-2, 3-4
 - contents • 3-6
 - Accessibility of memory
 - See Buffer
 - Access violation
 - See SS\$_ACCVIO
 - ACF (configuration control block) • 1-2 to 1-4
 - ACL (access rights list) • 1-45
 - ACP (ancillary control process) • 1-12, 1-39, 1-40, 1-74
 - See also XQP
 - class • 1-28
 - default • 1-28
 - ACP_MULTIPLE parameter • 1-28
 - Adapter dispatch table • 1-6, 1-7
 - address • 1-7
 - ADP\$_CSR • 3-82
 - ADP\$_DPQFL • 3-87
 - ADP\$_MBASCB • 1-7
 - ADP\$_MBASPT • 1-8
 - ADP\$_ADPTYPE • 2-3
 - ADP\$_DPBITMAP • 3-96
 - ADP (adapter control block) • 1-4 to 1-11
 - address • 1-26, 1-36
 - alternate map register allocation information • 1-10
 - alternate map register wait queue • 1-10
 - data path allocation information • 1-9
 - data path wait queue • 1-7
 - fields supporting ADPDISP macro • 2-3
 - map register allocation information • 1-9
 - map register wait queue • 1-8
 - size • 1-4
 - ADPDISP macro • 2-2 to 2-4
 - examples • 2-4
 - Affinity
 - See Device affinity
 - Allocation class • 1-28
 - Alternate map registers • 1-8, 1-26 to 1-27, 2-3
 - allocating • 3-63 to 3-64
 - allocating permanent • 1-26
 - loading • 2-44, 3-74 to 3-75
 - number of active • 1-10, 1-11
 - number of disabled • 1-11
 - Alternate map registers (Cont.)
 - releasing • 2-53, 3-84 to 3-85
 - requesting • 2-58, 3-92 to 3-93
 - Alternate map register wait queue • 1-10, 3-93
 - Alternate start I/O routine • 3-17
 - address • 1-30, 4-2
 - context • 4-2
 - entry point • 4-2
 - exit method • 4-2
 - input • 4-2
 - register usage • 4-2
 - synchronization requirements • 4-2
 - ARB (access rights block) • 1-42
 - AST (asynchronous system trap) • 3-6 to 3-7
 - See also Attention AST
 - control • 1-86
 - delivering • 3-2, 3-11
 - for aborted I/O request • 3-11
 - out of band • 1-86
 - process-requested • 3-7, 3-10, 3-73
 - queuing • 3-73
 - special kernel-mode • 1-12
 - user specified • 1-39
 - Asynchronous event notification • 2-70, 2-73 to 2-90
 - Asynchronous SCSI data transfer mode
 - enabling • 2-88
 - AT\$_GENBI • 1-33
 - AT\$_MBA • 1-33
 - AT\$_UBA • 1-33
 - Attached processor
 - See Secondary processor
 - Attention AST
 - See also AST
 - blocking • 1-82, 1-83
 - delivering • 3-2
 - disabling • 3-6 to 3-7
 - enabling • 3-6 to 3-7
 - flushing • 3-4
 - Autoconfiguration
 - See also System Generation Utility
-

B

- BADDALRQSZ bugcheck • 3-3, 3-19

Index

BIIC (backplane interconnect interface chip)
 self test • 2-5

BIOLM (buffered I/O limit) quota
 for mailbox • 1-73

BI_NODE_RESET macro • 2-5

BOOTED processor state • 1-16

Boot stack • 1-15

BOOT_REJECTED processor state • 1-16

BR level
 relation to SCB vectors • 1-9

Buffer
 allocating • 3-12 to 3-13, 3-14, 3-15, 3-22 to 3-23
 allocating a physically contiguous • 3-16
 deallocating • 3-3, 3-19
 locking • 1-42, 1-43, 3-31 to 3-33, 3-34 to 3-36, 3-40 to 3-42, 3-45 to 3-47, 3-54 to 3-55, 3-58 to 3-60
 locking multiple areas • 3-34, 3-45, 3-58
 moving data to from system to user • 3-80 to 3-81
 moving data to from user to system • 3-79
 testing accessibility of • 2-39 to 2-40, 3-31 to 3-33, 3-34 to 3-36, 3-40 to 3-42, 3-43 to 3-44, 3-45 to 3-47, 3-54 to 3-55, 3-56 to 3-57, 3-58 to 3-60
 unlocking • 3-109

Buffered data path • 1-8
 allocating permanent • 1-26
 odd transfer • 1-8
 purging • 3-82 to 3-83
 releasing • 2-55, 3-87
 requesting • 2-60, 3-96 to 3-97

Buffered I/O • 1-40, 1-41, 1-79
 chained • 1-40
 complex • 1-40
 postprocessing • 3-72

Bugcheck
 BADDALRQSZ • 3-3, 3-19
 ILLQBUSCFG • 1-22
 INCONSTATE • 3-88, 3-97
 SPLACQERR • 3-111
 SPLIPLHIGH • 3-111, 3-113
 SPLIPLLOW • 3-114, 3-115, 3-116, 3-117
 SPLRELEERR • 3-114, 3-115
 SPLRSTERR • 3-116, 3-117
 UBMAPEXCED • 3-74, 3-78
 UNSUPRTCPU • 2-10

BYTCNT (byte count) quota
 crediting • 3-18
 debiting • 3-12, 3-20 to 3-21, 3-22 to 3-23
 system maximum • 3-20, 3-22

BYTCNT (byte count) quota (Cont.)
 verifying • 3-20 to 3-21, 3-22 to 3-23

Byte count quota
 See BYTCNT

Byte limit
 See BYTLM

BYTLM (byte limit) quota
 crediting • 3-18
 debiting • 3-12, 3-20 to 3-21, 3-22 to 3-23

C

Cache control block • 1-83

Caching • 1-75

Cancel I/O routine • 1-30
 address • 4-4
 context • 4-4
 entry point • 4-4
 exit method • 4-5
 flushing ASTs in • 3-4
 input • 4-5
 register usage • 4-4
 synchronization requirements • 4-4

Card reader • 1-76

Carriage control • 1-74

CASE macro • 2-6
 example • 2-6

CCB\$B_AMOD • 3-103

CCB (channel control block) • 1-11 to 1-12
 address • 3-103

Channel index number • 3-68, 3-103, 4-5

Class driver entry vector table • 1-34

Class driver vector table • 1-89
 address • 2-8
 relocating • 2-7

CLASS_CTRL_INIT macro • 1-89, 2-7

CLASS_GETNXT service routine • 1-89, 2-8

CLASS_PUTNXT service routine • 1-89, 2-8

CLASS_UNIT_INIT macro • 2-8

Cloned UCB routine • 1-78
 address • 1-31, 4-6
 context • 4-6
 exit method • 4-7
 input • 4-6
 register usage • 4-6
 synchronization requirements • 4-6

COM\$DELATTNAST • 3-2

COM\$DRVDEALMEM • 3-3

COM\$FLUSHATTNS • 3-4, 3-6

- COM\$POST • 3–5, 4–2
 - COM\$POST_NOCNT • 3–5
 - COM\$SETATTNAST • 3–6 to 3–7
 - Connection
 - breaking • 2–73
 - obtaining characteristics of • 2–75 to 2–76
 - requesting • 2–70 to 2–71
 - setting characteristics of • 2–88 to 2–89
 - Connection characteristics buffer • 2–88
 - Controller initialization routine
 - address • 1–25, 2–26, 4–8
 - context • 4–8
 - entry point • 4–8
 - exit method • 4–8
 - forking • 1–21
 - for terminal port driver • 2–7
 - functions • 4–9
 - input • 4–8
 - register usage • 4–8
 - synchronization requirements • 4–8
 - Coroutine • 3–35, 3–46, 3–59, 3–109
 - CPU\$_PHY_CPUID • 3–70
 - CPU\$_SWIQFL • 3–26, 3–30
 - CPU\$_WORK_IFQ • 1–17
 - CPU (per-CPU database) • 1–12 to 1–19
 - locating • 2–31
 - CPUDISP macro • 2–9 to 2–11
 - CPU ID • 1–17, 3–70
 - CRB\$_INTD • 1–22 to 1–27
 - CRB\$_WQFL • 3–86, 3–91
 - CRB (channel request block) • 1–19 to 1–27
 - fork block • 1–21
 - initializing • 2–25
 - periodic wakeup of • 1–22
 - primary • 1–73
 - reinitializing • 2–25
 - secondary • 1–22
 - CSR (control and status register)
 - address • 1–36
 - bad address • 1–36
 - CTL\$_GL_CCBASE • 3–103
-
- D**
-
- Data path • 1–25 to 1–26
 - autopurging • 1–8, 2–3
 - buffered • 1–8, 2–3
 - direct • 2–3
 - purging • 2–51, 3–82 to 3–83
 - Data path allocation bit map • 1–9
 - Data path register
 - purge error • 3–83
 - Data path wait queue • 1–7, 3–88, 3–97
 - Data storage
 - device specific • 1–41, 1–68, 2–22
 - Data structure • 1–1
 - defining bit field within • 2–102 to 2–103
 - defining field within • 2–14, 2–15, 2–16
 - initializing • 2–24 to 2–26
 - Data transfer
 - byte aligned • 2–3, 3–78
 - byte count • 1–79, 1–83
 - byte offset • 1–79, 3–77
 - mapping local buffer for SCSI port • 2–77 to 2–79
 - negative byte count • 3–32, 3–35, 3–41, 3–43, 3–46, 3–55, 3–56, 3–59
 - starting address • 1–79
 - unmapping local buffer • 2–91
 - word aligned • 3–78
 - zero byte count • 3–32, 3–41, 3–55
 - Data transfer mode
 - as controlled by a third-party SCSI class driver • 2–88
 - asynchronous • 2–88
 - determining setting of • 2–75
 - synchronous • 2–88
 - \$DCDEF macro • 1–76, 2–3, 2–21
 - DDB (device data block) • 1–27 to 1–28
 - address • 1–74
 - initializing • 2–25
 - reinitializing • 2–25
 - DDT\$_ALTSTART • 4–2
 - DDT\$_CANCEL • 4–4
 - DDT\$_CLONEDUCB • 4–6
 - DDT\$_REGDUMP • 4–15
 - DDT\$_START • 4–17
 - DDT\$_UNITINIT • 4–22
 - DDT\$_UNSOLINT • 4–24
 - DDT (driver dispatch table) • 1–29 to 1–31, 3–102
 - address • 1–28, 1–80, 2–25
 - creating • 2–12 to 2–13
 - DDTAB macro • 2–12 to 2–13, 3–102
 - example • 2–13
 - \$DEFEND macro • 1–70, 2–15
 - example • 2–16
 - \$DEFINI macro • 1–70, 2–16
 - example • 2–16
 - \$DEF macro • 1–70, 2–14
 - example • 2–16
 - DEV\$_ELG • 3–8

Index

- \$DEVDEF macro • 1-74, 1-75
- Device
 - allocation class • 1-28
 - associated mailbox • 1-77
 - bus • 1-76
 - card reader • 1-76
 - cluster accessible • 1-73
 - cluster available • 1-75
 - directory structured • 1-74
 - disk • 1-76, 3-51, 3-95
 - dual ported • 1-74, 1-75
 - file structured • 1-28, 1-74
 - input • 1-75
 - line printer • 1-76
 - mailbox • 1-75, 1-76
 - mounted • 1-75, 1-78
 - mounted foreign • 1-75
 - network • 1-74
 - output • 1-75
 - random access • 1-75
 - real time • 1-75, 1-76
 - record oriented • 1-74
 - reference count • 1-79
 - sequential block-oriented • 1-74
 - shareable • 1-75
 - spooled • 1-74
 - synchronous communications • 1-76
 - tape • 1-76, 3-95
 - terminal • 1-74, 1-76
 - timed out • 1-78
 - workstation • 1-76
- Device affinity • 1-75, 3-71
- Device allocation lock • 1-73
- Device characteristics • 1-74 to 1-75
 - retrieving • 3-49
 - setting • 3-50 to 3-51
 - specifying • 2-25
- Device class • 1-76
 - specifying • 2-25
- Device controller • 1-19
 - multiunit • 1-36, 1-74, 1-77
 - number of units created for • 2-22
 - number of units supported by • 1-34, 1-36, 1-37, 2-22
 - reinitializing • 2-22
 - single unit • 1-36
 - status • 1-21
- Device controller data channel
 - See also Secondary controller data channel
 - obtaining ownership of • 1-36, 2-62, 3-100 to 3-101
- Device controller data channel (Cont.)
 - releasing • 2-54, 3-86
 - releasing before waiting for interrupt • 3-105
 - relinquishing ownership • 2-104
 - retaining ownership • 2-104
 - retaining while waiting for interrupt • 3-105
- Device controller data channel wait queue • 1-21, 3-86, 3-91, 3-101
- Device database
 - synchronizing access to • 2-17 to 2-18
- Device driver
 - branching on adapter characteristics • 2-2 to 2-4
 - branching on processor type • 2-9 to 2-11
 - entry points • 1-29, 4-1 to 4-24
 - for generic VAXBI device • 3-107
 - implementing a conditional wait • 2-92, 2-94
 - loading • 1-33
 - machine independence • 2-2 to 2-4, 2-9 to 2-11
 - name • 1-28, 1-34, 2-22
 - program sections • 2-13, 2-21
 - size • 1-33
 - suspending • 1-73
 - unloading • 1-33, 2-22
- Device interrupt
 - direct-vector • 1-7, 1-8, 1-25, 2-3
 - expected • 1-77, 3-105
 - multilevel Q22-bus • 1-22
 - non-direct-vector • 1-7, 1-25
 - unsolicited • 1-30
 - waiting for • 2-105, 3-104 to 3-106
- Device IPL • 1-77, 2-17 to 2-18
 - specifying • 2-25
- Device lock • 1-68, 1-77, 3-105
 - acquisition IPL • 3-113
 - address • 1-22, 1-36, 1-74
 - multiple acquisition of • 2-19, 3-117
 - obtaining • 2-17 to 2-18, 3-110, 3-113
 - releasing • 2-19 to 2-20, 3-115
 - restoring • 2-19, 3-117
- DEVICELOCK macro • 2-17 to 2-18, 2-66, 2-104, 3-110, 3-113
 - example • 2-18, 2-20, 2-66
- Device name • 1-28
- Device registers
 - accessing • 1-25, 1-36, 2-17 to 2-18
 - saving the value of • 4-16
- Device type • 1-76
 - specifying • 2-25
- Device unit • 1-68
 - allocating • 1-74, 1-75, 1-77
 - autoconfiguring • 2-22
 - busy indicator • 1-78

- Device unit (Cont.)
 - deaccessing • 1–12
 - deallocating • 1–78
 - error retry count • 1–79
 - marking available • 1–75
 - marking on line • 1–78
 - number • 1–77
 - operations count • 3–95
 - reference count • 4–4
 - reinitializing • 2–22
 - status • 1–77 to 1–79
 - DEVICEUNLOCK macro • 2–19 to 2–20, 2–66, 3–115, 3–117
 - example • 2–18, 2–20, 2–66
 - issued by IOC\$WFIKPC and IOC\$WFIRLCH • 3–105
 - Diagnostic buffer • 1–40, 1–42, 1–79, 1–83, 3–71
 - copied to process space • 3–73
 - filling • 3–69
 - size • 1–30
 - Direct data path
 - odd transfer • 1–8
 - Direct I/O • 1–40, 1–79
 - additional buffer regions for • 1–42 to 1–44
 - checking accessibility of process buffer for • 3–43 to 3–44, 3–56 to 3–57
 - locking a process buffer for • 3–31 to 3–33, 3–34 to 3–36, 3–40 to 3–42, 3–45 to 3–47, 3–54 to 3–55, 3–58 to 3–60
 - postprocessing • 3–72
 - unlocking process buffer • 3–109
 - Directory sequence number • 1–82, 1–83
 - Direct-vector interrupt • 1–7, 1–8, 1–25, 2–3
 - Disconnect feature
 - determining setting of • 2–75
 - enabling • 2–88
 - Disk driver • 1–78, 1–79
 - See also MBA, MASSBUS
 - ECC correction routine for • 3–67
 - using local disk UCB extension • 1–69, 1–82 to 1–84
 - DMA transfer
 - for modify operation • 3–31 to 3–33, 3–34 to 3–36
 - for read operation • 3–40 to 3–42, 3–45 to 3–47
 - for write operation • 3–54 to 3–55, 3–58 to 3–60
 - DPT\$V_SVP • 1–79, 2–21, 3–79, 3–80
 - DPT\$W_DELIVER • 4–21
 - DPT\$W_UNLOAD • 4–10
 - DPT (driver prologue table) • 1–31 to 1–35, 1–74, 1–76
 - creating • 2–21 to 2–26
 - DPT (driver prologue table) (Cont.)
 - initialization table • 1–33, 2–25 to 2–26
 - reinitialization table • 2–25, 2–25 to 2–26
 - DPTAB macro • 1–69, 2–21 to 2–23
 - example • 2–23
 - DPT_STORE macro • 2–24 to 2–26
 - example • 2–23
 - Driver unloading routine • 2–22, 2–26
 - address • 1–34, 4–10
 - context • 4–10
 - exit method • 4–10
 - functions • 4–10
 - input • 4–10
 - register usage • 4–10
 - synchronization requirements • 4–10
 - DSBINT macro • 2–27
 - Dual path UCB extension • 1–69
 - Dual ported device • 1–74
 - DYN\$C_BUFIO • 3–12, 3–22
 - DYN\$C_IRP • 3–12
 - DZ11 controller • 1–21
 - DZ32 controller • 1–21
-
- ## E
-
- ECC error correction • 1–78, 1–79, 1–83, 2–21, 3–67
 - ECC position register • 1–83
 - ECRB (Ethernet controller data block) • 2–2
 - EMB\$W_DV_STS • 3–94
 - EMB spin lock • 3–8
 - ENBINT macro • 2–28
 - Encryption key • 1–42
 - Entry point
 - specifying in driver tables • 2–13
 - \$EQLST macro • 2–29 to 2–30
 - example • 2–30, 2–103
 - ERL\$DEVICEATTN • 3–8 to 3–9, 4–15
 - ERL\$DEVICERR • 1–30, 1–80, 1–81, 3–8 to 3–9, 4–15
 - ERL\$DEVICTMO • 1–30, 1–80, 1–81, 3–8 to 3–9, 4–15
 - ERL\$RELEASEMB • 3–95
 - Error
 - servicing within driver • 3–82 to 3–83
 - Error log allocation buffer • 3–8
 - Error logging • 1–79 to 1–80, 3–8 to 3–9
 - enabling • 1–75
 - error log sequence number • 1–42
 - inhibiting • 3–8
 - in progress • 1–77

Index

Error logging (Cont.)

- performed by IOC\$REQCOM • 3-95
- Error logging routine • 1-30
- Error log in progress bit
 - See UCB\$V_ERLOGIP
- Error log UCB extension • 1-69, 1-80 to 1-81
- Error message buffer • 1-81, 1-83, 3-82
 - allocating • 3-8
 - filling • 3-9
 - releasing • 3-95
 - size • 3-8
 - specifying size • 1-30
 - written into by IOC\$REQCOM • 3-95
- Event flag • 1-39
 - handling for aborted I/O request • 3-11
- EXE\$ABORTIO • 1-40, 3-7, 3-10 to 3-11, 3-33, 3-42, 3-44, 3-46, 3-50, 3-51, 3-55, 3-57, 3-59, 4-12
- EXE\$ALLOCBUF • 3-12 to 3-13
- EXE\$ALLOCIRP • 1-42, 1-44, 3-12 to 3-13
- EXE\$ALONONPAGED • 3-13, 3-14, 3-61
- EXE\$ALONPAGVAR • 3-15
- EXE\$ALOPHYCNTG • 3-16
- EXE\$ALTQUEPKT • 1-30, 3-5, 3-17, 4-2, 4-12
- EXE\$ASSIGN • 1-11, 1-12, 4-6
- EXE\$CANCEL • 3-68
- EXE\$CREDIT_BYTCNT • 3-18
- EXE\$CREDIT_BYTCNT_BYTLM • 3-18
- EXE\$DASSGN • 1-12
- EXE\$DEANONPAGED • 3-3, 3-13, 3-19
- EXE\$DEBIT_BYTCNT • 3-20 to 3-21
- EXE\$DEBIT_BYTCNT_ALO • 3-22 to 3-23
- EXE\$DEBIT_BYTCNT_BYTLM • 3-20 to 3-21
- EXE\$DEBIT_BYTCNT_BYTLM_ALO • 3-22 to 3-23
- EXE\$DEBIT_BYTCNT_BYTLM_NW • 3-20 to 3-21
- EXE\$DEBIT_BYTCNT_NW • 3-20 to 3-21
- EXE\$FINISHIO • 1-41, 3-24 to 3-25, 3-49, 3-50, 3-51, 4-12
- EXE\$FINISHIOC • 1-41, 3-24 to 3-25, 4-12
- EXE\$FORK • 1-21, 2-32, 3-26
- EXE\$FORKDSPTH • 1-73
- EXE\$GB_CPUTYPE • 2-10
- EXE\$GL_ABSTIM • 1-22
- EXE\$GL_INTSTK
 - replaced by CPU\$L_INTSTK • 1-12
- EXE\$GQ_1ST_TIME • 3-29
- EXE\$GQ_SYSTIME • 2-52, 3-69
- EXE\$INSERTIRP • 1-38, 1-39, 1-76, 3-27, 3-28, 3-38
- EXE\$INSIOQ • 1-77, 3-28, 3-38
- EXE\$INSIOQC • 3-28
- EXE\$INSTIMQ • 3-29

- EXE\$IOFORK • 1-72, 1-73, 3-30
- EXE\$MODIFY • 3-31 to 3-33
- EXE\$MODIFYLOCK • 3-32, 3-34 to 3-36
- EXE\$MODIFYLOCKR • 1-43, 3-32, 3-34 to 3-36, 3-109
- EXE\$ONEPARM • 1-41, 3-37
- EXE\$QIO • 1-12, 1-30, 1-37 to 1-40, 1-42
- EXE\$QIOACPPKT • 1-74
- EXE\$QIODRVPKT • 3-32, 3-33, 3-37, 3-38, 3-41, 3-51, 3-55, 3-62, 4-12
- EXE\$QIORETURN • 3-39
- EXE\$READ • 1-41, 3-40 to 3-42
- EXE\$READCHK • 3-43 to 3-44
- EXE\$READCHKR • 3-32, 3-35, 3-41, 3-43 to 3-44, 3-46
- EXE\$READLOCK • 3-41, 3-45 to 3-47
- EXE\$READLOCKR • 1-43, 3-41, 3-45 to 3-47, 3-109
- EXE\$RMVTIMQ • 3-48
- EXE\$SENSEMODE • 3-49
- EXE\$SETCHAR • 3-50 to 3-51
- EXE\$SETMODE • 3-50 to 3-51
- EXE\$SNDEVMMSG • 3-52 to 3-53
- EXE\$TIMEOUT • 1-74, 1-77, 1-79
- EXE\$WRITE • 1-41, 3-54 to 3-55
- EXE\$WRITECHK • 3-56 to 3-57
- EXE\$WRITECHKR • 3-55, 3-56 to 3-57, 3-59
- EXE\$WRITELOCK • 3-55, 3-58 to 3-60
- EXE\$WRITELOCKR • 1-43, 3-55, 3-58 to 3-60, 3-109
- EXE\$WRTMAILBOX • 3-52, 3-61
- EXE\$ZEROPARM • 1-41, 3-62

F

FDT (function decision table)

- address • 1-30
- creating • 2-37 to 2-38
- size • 1-31

FDT routine

- adjusting process quotas in • 3-12
- allocating IRPE in • 1-42
- completing an I/O operation in • 3-24 to 3-25
- context • 4-11
- entry point • 4-11
- exit method • 4-12
- for direct I/O • 3-31 to 3-33, 3-40 to 3-42, 3-54 to 3-55
- register usage • 4-11
- returning to the system service dispatcher • 3-39

FDT routine (Cont.)
 setting attention ASTs in • 3–6
 specifying • 4–11
 synchronization requirements • 4–11
 unlocking process buffers in • 3–109

File structured device • 1–74

FIND_CPU_DATA macro • 2–31
 example • 2–31

Fork block • 2–104, 3–26, 3–30, 3–104 to 3–106
 in CRB • 1–21
 in UCB • 1–72 to 1–73

Fork database
 accessing • 2–33 to 2–34

Fork dispatcher • 2–33

Forking • 2–32, 2–43, 3–26, 3–30
 from controller initialization routine • 4–8
 from driver unloading routine • 4–10
 from unit initialization routine • 4–22

Fork IPL • 1–73, 2–33 to 2–34

Fork lock • 1–21, 1–68
 acquisition IPL • 3–111
 multiple acquisition of • 2–35, 3–116
 obtaining • 2–33 to 2–34, 3–111 to 3–112
 releasing • 2–35 to 2–36, 3–114
 restoring • 2–35, 3–116

Fork lock index • 1–73
 placing in UCB\$B_FLCK • 2–25

FORKLOCK macro • 2–33 to 2–34, 3–111
 example • 2–34

FORK macro • 2–32, 3–26

Fork process
 creating • 2–32, 2–43, 3–26, 3–30
 creation by IOC\$INITIATE • 3–70 to 3–71
 suspending • 2–104, 3–104 to 3–106

Fork queue • 1–17, 1–72, 3–26, 3–30

FORKUNLOCK macro • 2–35 to 2–36, 3–114, 3–116
 example • 2–34

Full duplex device driver • 4–2
 I/O completion for • 3–5

FUNCTAB macro • 2–37 to 2–38
 example • 2–38

H

HWCLK spin lock • 3–29, 3–48

I/O adapter
 configuration register • 1–6
 data path register • 2–51
 number of address bits • 1–8, 2–3
 type • 1–7, 1–33, 2–3, 2–21

I/O database • 1–1, 1–2
 creation • 1–33, 2–25

I/O function code • 1–39

I/O postprocessing • 1–41
 device-independent • 3–72 to 3–73
 for aborted I/O request • 3–10
 for full duplex device driver • 3–5
 for I/O request involving no device activity • 3–24 to 3–25

I/O postprocessing queue • 1–17, 1–79, 3–5, 3–95

I/O request
 aborting • 3–10 to 3–11
 canceling • 1–30, 1–78, 3–68
 completing • 3–94 to 3–95
 outstanding on channel • 1–12
 status • 1–40
 with no parameters • 3–62
 with one parameter • 3–37

I/O status block
 See IOSB

IDB\$_OWNER • 3–86, 3–100

IDB\$_NO_CSR • 1–36

IDB (interrupt dispatch block) • 1–35 to 1–37
 creation • 2–22
 size • 2–22

IFNORD macro • 2–39 to 2–40

IFNOWRT macro • 2–39 to 2–40

IFRD macro • 2–39 to 2–40
 example • 2–40

IFWRT macro • 2–39 to 2–40

ILLQBUSCFG bugcheck • 1–22

Image termination • 4–4

INCONSTATE bugcheck • 3–88, 3–97

Initialization table • 1–34, 2–25

Initiator
 completing an operation (in AEN mode) • 2–74
 enabling selection of • 2–70, 2–73 to 2–90
 receiving data from target (in AEN mode) • 2–80
 sending bytes to target (in AEN mode) • 2–83

INIT processor state • 1–16

Input device • 1–75

Interprocessor interrupt • 1–16

Index

- Interrupt
 - blocking • 2-27, 2-65
 - interprocessor • 1-16
 - requesting a software • 2-67
- Interrupt dispatcher • 1-7, 1-9
 - for MASSBUS • 4-24
 - for UNIBUS • 1-25
- Interrupt service routine • 1-73
 - address • 1-25, 2-26, 4-13
 - context • 4-13
 - entry point • 4-13
 - exit method • 4-14
 - for MASSBUS device • 4-13
 - for unsolicited interrupt • 4-24
 - functions • 4-14
 - input • 4-14
 - register usage • 4-14
 - specifying more than one • 4-13
 - synchronization requirements • 4-13
- Interrupt stack
 - address • 1-16
- INVALIDATE_TB macro • 2-41 to 2-42
- IO\$V_INHERLOG • 3-8
- IO\$_SENSECHAR function
 - servicing • 3-49
- IO\$_SENSEMODE function
 - servicing • 3-49
- IO\$_SETCHAR function
 - servicing • 3-50 to 3-51
- IO\$_SETMODE function
 - servicing • 3-50 to 3-51
- IOC\$ALOALTMAP • 1-10, 3-63 to 3-64, 3-93
- IOC\$ALOALTMAPN • 3-63 to 3-64
- IOC\$ALOALTMAPSP • 3-63 to 3-64
- IOC\$ALOUBAMAP • 3-65 to 3-66, 3-90, 3-99
- IOC\$ALOUBAMAPN • 3-65 to 3-66
- IOC\$APPLYECC • 1-83, 3-67
- IOC\$CANCELIO • 1-77, 3-68, 4-4
- IOC\$DIAGBUFILL • 1-30, 1-42, 3-69
- IOC\$GL_CRBTMOUT • 1-22
- IOC\$GL_DEVLIST • 1-27
- IOC\$GL_MUTEX • 4-6
- IOC\$GW_MAXBUF • 3-20, 3-22
- IOC\$INITIATE • 1-30, 1-40, 1-41, 1-77, 1-79, 3-28, 3-38, 3-69, 3-70 to 3-71, 3-95, 4-17
- IOC\$IOPOST • 1-41, 1-42, 1-43, 3-72 to 3-73
 - unlocking process buffers • 3-109
- IOC\$LOADALTMAP • 2-44, 3-74 to 3-75
- IOC\$LOADMBAMAP • 2-45, 3-76
- IOC\$LOADUBAMAP • 1-26, 2-46, 3-77 to 3-78
- IOC\$LOADUBAMAPA • 3-77 to 3-78
- IOC\$MNTVER • 1-30
- IOC\$MOVFRUSER • 2-21, 3-79
- IOC\$MOVFRUSER2 • 3-79
- IOC\$MOVTOUSER • 2-21, 3-80 to 3-81
- IOC\$MOVTOUSER2 • 3-80 to 3-81
- IOC\$PURGDATAP • 1-26, 2-51, 3-82 to 3-83
- IOC\$RELALTMAP • 1-10, 1-73, 2-53, 3-84 to 3-85
- IOC\$RELCHAN • 1-21, 1-36, 1-73, 2-54, 3-86, 3-95
 - called by IOC\$WFIRLCH • 3-106
- IOC\$RELDATAP • 1-7, 1-9, 1-73, 2-55, 3-87
- IOC\$RELMAPREG • 1-8, 1-9, 1-25, 1-26, 1-73, 2-56, 3-89 to 3-90
- IOC\$RELSCHAN • 1-21, 1-22, 1-36, 2-57, 3-91
- IOC\$REQALTMAP • 1-10, 1-73, 2-58, 3-92 to 3-93
- IOC\$REQCOM • 1-30, 1-38, 1-41, 1-76, 1-77, 1-79, 1-81, 2-59, 3-13, 3-94 to 3-95, 4-17
- IOC\$REQDATAP • 1-7, 1-9, 1-26, 1-73, 2-60, 3-96 to 3-97
- IOC\$REQDATAPNW • 3-96 to 3-97
- IOC\$REQMAPREG • 1-8, 1-9, 1-25, 1-26, 1-73, 2-61, 3-98 to 3-99
- IOC\$REQPCHANH • 1-21, 1-36, 1-73, 2-62, 3-100 to 3-101
- IOC\$REQPCHANL • 1-21, 1-36, 1-73, 2-62, 3-100 to 3-101
- IOC\$REQSCHANH • 1-21, 1-22, 1-36, 2-63, 3-100 to 3-101
- IOC\$REQSCHANL • 1-21, 1-22, 1-36, 1-73, 2-63, 3-100 to 3-101
- IOC\$RETURN • 2-13, 3-102
- IOC\$SEARCHDEV • 1-74
- IOC\$VERIFYCHAN • 3-103
- IOC\$WFIKPCH • 1-73, 1-77, 1-79, 3-104 to 3-106
- IOC\$WFIRLCH • 1-77, 1-79, 3-104 to 3-106
- IOFORK macro • 2-43, 3-30
- IOSB (I/O status block) • 1-39, 1-41, 3-5, 3-10, 3-73, 3-95
- IPL\$_ASTDEL • 3-10, 3-12, 3-31, 3-34, 3-37, 3-38, 3-40, 3-43, 3-49, 3-50, 3-56, 3-62, 3-73, 3-103, 3-114, 3-116, 3-117, 4-6, 4-11
- IPL\$_EMB • 3-8
- IPL\$_IOPOST • 3-5, 3-10, 3-25, 3-73, 3-95
- IPL\$_MAILBOX • 3-52, 3-61
- IPL\$_POOL • 3-14, 3-15
- IPL\$_POWER • 4-8, 4-10
- IPL\$_QUEUEAST • 3-2, 3-3
- IPL\$_RESCHED • 2-31, 3-111, 3-113
- IPL\$_TIMER • 3-29, 3-48
- IPL (interrupt priority level)
 - See also Device IPL, Fork IPL
 - lowering • 2-97, 3-26, 3-30

IPL (interrupt priority level) (Cont.)
 modifying • 2-17 to 2-18, 2-19 to 2-20, 2-27,
 2-28, 2-33 to 2-34, 2-35 to 2-36, 2-47 to
 2-48, 2-65, 2-96
 raising • 2-49, 2-65
 saving • 2-17, 2-33, 2-47, 2-64
 IRP\$B_CARCON • 1-41, 3-32, 3-41, 3-55
 IRP\$B_PRI • 3-27
 IRP\$L_BCNT • 3-32, 3-35, 3-41, 3-43, 3-46, 3-55,
 3-56, 3-59, 3-70, 3-71, 3-72
 IRP\$L_DIAGBUF • 3-69, 3-70, 3-71
 IRP\$L_IOST2 • 3-32, 3-41, 3-55
 IRP\$L_KEYDESC • 3-72
 IRP\$L_MEDIA • 1-41, 3-37, 3-51, 3-62
 IRP\$L_PID • 3-68, 4-5
 IRP\$L_SVAPTE • 3-33, 3-35, 3-41, 3-46, 3-55,
 3-59, 3-70, 3-71
 IRP\$V_BUFIO • 3-72
 IRP\$V_DIAGBUF • 3-69, 3-70, 3-71, 3-72
 IRP\$V_EXTEND • 3-72
 IRP\$V_FUNC • 3-32, 3-35, 3-41, 3-43, 3-46
 IRP\$V_KEY • 3-72
 IRP\$V_MBXIO • 3-72
 IRP\$V_PHYSIO • 3-72
 IRP\$W_BOFF • 3-33, 3-35, 3-41, 3-46, 3-55, 3-59,
 3-70, 3-71, 3-72
 IRP\$W_CHAN • 3-68, 4-5
 IRP (I/O request packet) • 1-37 to 1-42
 current • 1-77
 deallocation • 3-73
 dequeuing from UCB • 1-38
 insertion in pending-I/O queue • 3-27, 3-28
 size • 1-37
 unlocking buffers specified in • 3-109
 IRPE (I/O request packet extension) • 1-40, 1-42 to
 1-44, 3-72
 address • 1-42
 allocating • 1-42
 deallocation • 1-43, 3-73, 3-109
 unlocking buffers specified in • 3-73, 3-109

J

JIB\$L_BYTCNT • 3-12, 3-18, 3-20, 3-22
 JIB\$L_BYTLM • 3-12, 3-18, 3-20, 3-22
 JIB\$V_BYTCNT_WAITERS • 3-18
 JIB spin lock • 3-18, 3-20, 3-23
 Job controller • 1-78
 sending a message to • 3-53, 3-61

Job quota

byte count • 3-12, 3-18, 3-20 to 3-21, 3-22 to
 3-23
 byte limit • 3-12, 3-18, 3-20 to 3-21, 3-22 to
 3-23

L

LDR\$ALLOC_PT • 3-107
 LDR\$DEALLOC_PT • 3-108
 LDR\$GL_FREE_PT • 3-107, 3-108
 LDR\$GL_SPTBASE • 3-107, 3-108
 LOADALT macro • 2-44, 3-74
 LOADER\$_PTE_NOT_EMPTY status • 3-108
 LOADMBA macro • 2-45, 3-76
 LOADUBA macro • 2-46, 3-77
 Local disk UCB extension • 1-69, 1-82 to 1-84
 required for error logging • 3-9
 required for IOC\$APPLYECC routine • 3-67
 Local tape UCB extension • 1-69, 1-81 to 1-82
 required for error logging • 3-9
 Lock ID • 1-73
 LOCK macro • 2-47 to 2-48, 3-111
 Lock manager • 1-73
 LOCK_SYSTEM_PAGES macro • 2-49
 Logical I/O function
 translation to physical function • 3-31, 3-40, 3-54
 Longword access enable bit
 See VEC\$_LWAE
 Longword-aligned random-access mode • 1-26
 Lookaside list
 See Nonpaged pool
 Loopback mode • 1-91
 LWAE (longword access enable) bit
 See VEC\$_LWAE

M

Macro

format • 2-1
 Mailbox • 1-75, 1-76, 1-77
 associated with device • 1-77
 buffered I/O quota for • 1-73
 I/O function • 1-40
 in shared memory • 1-78
 marked for deletion • 1-78
 permanent • 1-78

Index

Mailbox (Cont.)

- sending a message to • 3-52 to 3-53, 3-61
- MAILBOX spin lock • 3-52, 3-61
- Map registers • 1-8, 1-25, 1-26, 2-3
 - allocating • 3-65 to 3-66
 - allocating permanent • 1-25
 - byte offset bit • 3-77
 - loading • 2-46, 3-77 to 3-78
 - number of active • 1-9, 1-10
 - number of disabled • 1-10
 - of MBA • 2-45, 3-76
 - releasing • 2-56, 3-89 to 3-90
 - requesting • 2-61, 3-98 to 3-99
- Map register wait queue • 1-8, 3-90, 3-99
- MBA\$INT • 4-24
- MBA\$_BCR • 3-76
- MBA\$_MAP • 3-76
- MBA\$_VAR • 3-76
- MBA (MASSBUS adapter)
 - registers
 - map • 2-45, 3-76
 - releasing secondary data channel • 3-91
- Media ID • 1-80
- Memory
 - See also Buffer, Nonpaged pool
 - detecting parity errors in • 2-51
 - testing accessibility of • 2-39 to 2-40
- MMG\$IOLock • 3-33, 3-35, 3-41, 3-46, 3-55, 3-59
- MMG\$UNLOCK • 1-43, 3-109
- MMG spin lock • 3-16, 3-107, 3-108, 3-109
- Mount verification • 1-40, 1-78
- Mount verification routine • 1-30, 1-31
- Multilevel device interrupt dispatching • 1-22
- Multiprocessor state • 1-16
- Mutex
 - for ACL • 1-45
 - for I/O database • 4-6

N

- Network device • 1-74
- Nexus ID • 1-6
- Node ID • 1-6
- Non-direct-vector interrupt • 1-7, 1-25
- Nonpaged pool
 - allocating • 3-12 to 3-13, 3-14, 3-15, 3-22 to 3-23
 - deallocating • 3-3, 3-19
 - lookaside list • 3-13, 3-14
 - variable region • 3-15

O

- Object
 - protection • 1-45
- OPCOM process
 - sending a message to • 3-53, 3-61
- Operator device • 1-74
- ORB (object rights block) • 1-44 to 1-46
 - address • 1-73
 - cloned • 4-7
- Output device • 1-75

P

- Page table entry
 - allocating • 3-107
 - deallocating • 3-108
 - modifying • 2-41
- Paging I/O function • 1-40
- PCB\$_PID • 3-68, 4-5
- PCB\$_SSRWAIT • 3-12, 3-20, 3-22
- PCB\$_ASTCNT • 3-4, 3-6, 3-10
- PDT (port descriptor table) • 1-80
- Pending-I/O queue • 1-38, 1-76, 3-27, 3-28, 3-37, 3-38, 3-73, 3-95
 - bypassing • 3-17
 - length • 1-79, 3-28
- Per-CPU database
 - See CPU
- Performance
 - stack time • 1-17
- Physical I/O function • 1-40, 3-72
- PID (process identification number) • 1-74
- POOL spin lock • 3-14, 3-15, 3-19
- Poor man's lockdown • 2-49 to 2-50, 2-97
- Port
 - DMA buffer • 2-77 to 2-79
 - resetting • 2-82
- Port command buffer
 - allocating • 2-69
 - deallocating • 2-72
- Port driver entry vector table • 1-34
- Port driver vector table • 1-89
 - address • 2-8
 - creating • 2-99, 2-100
 - defining entry in • 2-98
 - relocating • 2-7
- PORT_MAINT initiate routine • 1-90

Power failure
 occurring when device is busy • 1-78
 Power failure recovery procedure • 1-25, 1-26, 1-74
 PR\$_SID processor register • 1-17
 PR\$_SIRR processor register • 2-67
 Process
 See also Process quota
 current • 1-15
 privilege mask • 1-42
 Process I/O channel • 1-11, 1-40
 deassigning • 4-4
 reference count • 1-77, 1-78
 validating • 3-103
 Processor state
 See Multiprocessor state
 Processor status longword
 See PSL
 Processor subtype • 2-9
 Processor type • 2-9
 Process quota
 charging • 1-41, 4-17
 PSL (processor status longword)
 Z condition code • 3-27
 PURDPR macro • 2-51, 3-82

Q

Q22-bus • 2-3
 device interrupt dispatching • 1-22
 QUEUEAST spin lock • 3-7
 Quota
 See Process quota, Job quota

R

Random access device • 1-75
 Read check
 enabling • 1-75
 Read function • 1-40, 1-41
 postprocessing for • 3-72
 READ_SYSTIME macro • 2-52
 example • 2-52
 Real time device • 1-75, 1-76
 Record oriented device • 1-74
 Register dumping routine • 1-30, 1-83, 2-51, 3-9,
 3-69, 3-82
 address • 4-15

Register dumping routine (Cont.)

 context • 4-15
 entry point • 4-15
 exit method • 4-15
 functions • 4-16
 input • 4-15
 register usage • 4-15
 synchronization requirements • 4-15
 Reinitialization table • 1-34, 2-25
 RELALT macro • 2-53, 3-84
 RELCHAN macro • 2-54, 3-86
 RELDPR macro • 2-55, 3-87
 RELMPR macro • 2-56, 3-89
 RELSCHAN macro • 2-57, 3-91
 Remote terminal UCB extension • 1-75
 REQALT macro • 3-92
 REQCOM macro • 2-59, 3-94
 REQDPR macro • 2-60, 3-96
 REQMPR macro • 2-61, 3-98
 REQPCAN macro • 2-62, 3-100
 REQSCHAN macro • 2-63, 3-100
 Resource wait mode • 3-12, 3-20, 3-22
 Resource wait queue
 See also Alternate map register wait queue,
 Device controller data channel wait queue
 See also Map register wait queue, Secondary data
 channel wait queue, Data path wait queue
 buffered data path • 3-88
 RUN processor state • 1-16

S

SAVIPL macro • 2-64
 SCB (system control block) • 1-7
 SCDRP\$_BCNT • 2-78, 2-85
 SCDRP\$_CMD_PTR • 2-85
 SCDRP\$_SCSI_FLAGS • 2-78
 SCDRP\$_STS_PTR • 2-85, 2-86
 SCDRP\$_SVAPTE • 2-78
 SCDRP\$_SVA_SPTTE • 2-79
 SCDRP\$_SVA_USER • 2-78, 2-79, 2-85
 SCDRP\$_TRANS_CNT • 2-86
 SCDRP\$_W_BOFF • 2-78
 SCDRP\$_W_FUNC • 2-85
 SCDRP\$_W_MAPREG • 2-79
 SCDRP\$_W_NUMREG • 2-79
 SCDRP\$_W_PAD_BCNT • 2-85
 SCDRP\$_W_STS • 2-78
 SCDRP (SCSI class driver request packet) • 1-46 to
 1-54

Index

- SCDT (SCSI connection descriptor table) • 1–54 to 1–60
- SCH\$POSTEF • 1–39
- SCHED spin lock • 3–19
- SCS (system communications services) • 1–33
- SCSI bus
 - releasing in AEN operation • 2–81
 - resetting • 2–82
 - sensing phase of • 2–87
 - setting phase of • 2–90
- SCSI class driver request packet
 - See SCDRP
- SCSI command
 - determining timeout setting for • 2–76
 - disabling retry • 2–75, 2–88
 - enabling retry • 2–75
 - sending to SCSI device • 2–84 to 2–86
 - setting disconnect timeout for • 2–76, 2–89
 - setting DMA timeout for • 2–76, 2–89
 - setting phase change timeout for • 2–76, 2–89
 - terminating • 2–68
- SCSI command byte
 - buffering • 2–69
- SCSI connection descriptor table
 - See SCDT
- SCSI port descriptor table
 - See SPDT
- Secondary controller data channel • 2–57
 - obtaining ownership of • 2–63, 3–100 to 3–101
 - releasing • 3–91
- Secondary controller data channel wait queue • 3–91, 3–101
- Set device characteristics function • 1–76
- Set device mode function • 1–76
- SETIPL macro • 2–65
 - example • 2–66
- Set mode function • 1–76
- Shareable device • 1–75
- SHOW DEVICE command • 1–80
- SMP\$ACQNOIPL • 2–17
- SMP\$ACQUIRE • 2–34, 2–47
- SMP\$ACQUIREL • 2–17
- SMP\$AR_IPLVEC • 2–33, 3–26, 3–30
- SMP\$AR_SPNLKVEC • 1–66, 2–34, 2–47, 2–96
- SMP\$RELEASE • 2–35, 2–96
- SMP\$RELEASEL • 2–19
- SMP\$RESTORE • 2–35, 2–96
- SMP\$RESTOREL • 2–19
- SOFTINT macro • 2–67, 3–26, 3–30
- SPDT (SCSI port descriptor table) • 1–60 to 1–66
- SPI\$ABORT_COMMAND macro • 2–68
- SPI\$ALLOCATE_COMMAND_BUFFER macro • 2–69
- SPI\$CONNECT macro • 2–70 to 2–71
- SPI\$DEALLOCATE_COMMAND_BUFFER macro • 2–72
- SPI\$DISCONNECT macro • 2–73
- SPI\$FINISH_COMMAND macro • 2–74
- SPI\$GET_CONNECTION_CHAR macro • 2–75 to 2–76, 2–88
- SPI\$MAP_BUFFER macro • 2–77 to 2–79
- SPI\$RECEIVE_BYTES macro • 2–80
- SPI\$RELEASE_BUS macro • 2–81
- SPI\$SEND_BYTES macro • 2–83
- SPI\$SEND_COMMAND macro • 2–84 to 2–86
- SPI\$SENSE_PHASE macro • 2–87
- SPI\$SET_CONNECTION_CHAR macro • 2–88 to 2–89
- SPI\$SET_PHASE macro • 2–90
- SPI\$UNMAP_BUFFER macro • 2–91
- SPI (SCSI port interface) • 2–68 to 2–90
 - calling protocol for • 2–68
 - extensions to • 2–73 to 2–90
- Spin lock
 - acquisition IPL • 1–67, 3–111
 - acquisition PC list • 1–68
 - dynamic • 1–68
 - multiple acquisition of • 2–96, 3–116
 - obtaining • 2–47 to 2–48, 3–111 to 3–112
 - ownership • 1–67, 1–68
 - rank • 1–67
 - releasing • 2–96, 3–114
 - restoring • 2–96, 3–116
 - static • 1–68
 - system • 1–68
- Spin wait • 1–68, 3–110, 3–112, 3–113
- SPL\$IPL • 1–77
- SPL (spin lock data structure) • 1–66 to 1–68
- SPLACQERR bugcheck • 3–111
- \$SPLCODDEF macro • 2–23, 2–25
- SPLIPLHIGH bugcheck • 3–111, 3–113
- SPLIPLLOW bugcheck • 3–114, 3–115, 3–116, 3–117
- SPLRELERR bugcheck • 3–114, 3–115
- SPLRSTERR bugcheck • 3–116, 3–117
- Spooled device • 1–74
- SPTREQ parameter • 3–16
- SS\$_ACCVIO • 3–32, 3–33, 3–35, 3–41, 3–43, 3–46, 3–50, 3–51, 3–55, 3–56, 3–59, 3–73
- SS\$_BADPARAM • 3–32, 3–35, 3–41, 3–43, 3–46, 3–55, 3–56, 3–59, 3–107
- SS\$_EXQUOTA • 3–6, 3–20, 3–22
- SS\$_ILLIOFUNC • 3–51
- SS\$_INSFMAPREG • 3–64

SS\$_INSFMEM • 3-6, 3-12, 3-14, 3-15, 3-16, 3-52, 3-61
 SS\$_INSFSPTS • 3-16, 3-107
 SS\$_INSFWSL • 3-33, 3-35, 3-41, 3-46, 3-59
 SS\$_IVCHAN • 3-103
 SS\$_MBFULL • 3-52, 3-61
 SS\$_MBTOOSML • 3-52, 3-61
 SS\$_NOPRIV • 3-52, 3-61, 3-103
 SS\$_SSFALL • 3-64, 3-75, 3-85, 3-93
 Start I/O routine
 See also Alternate start I/O routine
 activating • 3-28
 address • 1-30, 4-17
 checking for zero length buffer • 3-32, 3-41, 3-55
 context • 4-17
 entry point • 4-17
 exit method • 4-18
 input • 4-17
 register usage • 4-17
 synchronization requirements • 4-17
 transferring control to • 3-38, 3-70 to 3-71
 STOPPED processor state • 1-16
 STOPPING processor state • 1-16
 Subcontroller • 1-33
 Swapping I/O function • 1-40
 Symbol list
 defining • 2-29 to 2-30
 Synchronous communications device • 1-76
 Synchronous SCSI data transfer mode
 determining REQ-ACK offset setting • 2-75
 determining transfer period setting • 2-75
 enabling • 2-88
 setting REQ-ACK offset • 2-88
 setting transfer period • 2-88
 SYS\$ALLOC • 1-74, 1-77
 SYS\$ASSIGN • 1-11, 1-77, 1-78
 for template device • 4-6
 SYS\$CANCEL • 1-30, 4-4
 SYS\$DALLOC • 1-30, 1-77, 4-4
 SYS\$DASSGN • 1-30, 1-77, 4-4
 SYS\$QIO • 1-37
 device-dependent arguments of • 1-41
 SYS\$QIOW • 1-37
 System buffer
 See Buffer, Nonpaged pool
 System Generation Utility (SYSGEN)
 AUTOCONFIGURE command • 1-2, 1-34, 1-68, 2-22, 4-21
 CONNECT command • 1-7, 1-26, 1-36, 1-44, 1-68, 2-22, 4-8, 4-22
 /NUMVEC qualifier • 1-23

System Generation Utility (SYSGEN) (Cont.)

RELOAD command • 4-10

System page-table entry

allocating • 3-107

allocating permanent • 1-33, 1-79, 2-21, 3-79, 3-80

deallocating • 3-108

System resource

accessing • 2-47 to 2-48

System time • 3-69

reading • 2-52

T

Tape driver • 1-74, 4-13

using local tape UCB extension • 1-69, 1-81 to 1-82

Target

enabling selection from • 2-70, 2-73 to 2-90

Template UCB • 1-78

Terminal • 1-74, 1-76

See also Terminal controller, Terminal class driver, Terminal port driver, Terminal UCB extension

detached • 1-75

I/O function for • 1-40

redirected • 1-75

Terminal class driver

binding to port driver • 2-8

Terminal controller • 1-21

Terminal port driver • 2-7

binding to class driver • 2-8

control flags • 1-89

Terminal UCB extension • 1-69, 1-84 to 1-91

remote • 1-75

Third-party SCSI class driver

receiving notification of asynchronous events on target • 2-70, 2-73 to 2-90

Time

reading system • 2-52

TIMEDWAIT macro • 2-92 to 2-93

See also TIMEWAIT macro

example • 2-93

Timeout • 1-78, 2-104

detecting • 1-79

disabling • 2-43, 3-30

due time • 1-79

expected • 1-77, 3-105

for SCSI device • 2-89

Timeout handling routine • 2-104, 4-5

address • 4-19

Index

Timeout handling routine (Cont.)

- context • 4–19
- entry point • 4–19
- exit method • 4–20
- functions • 4–20
- input • 4–20
- register usage • 4–19
- synchronization requirements • 4–19

Timeout interval • 2–104

Timer queue • 3–29, 3–48

Timer queue element

- See TQE

TIMER spin lock • 3–29, 3–48

TIMEWAIT macro • 2–94

- See also TIMEDWAIT macro
- example • 2–95

TIMOUT processor state • 1–16

TQEB_RQTYPE • 3–48

TQEQ_TIME • 3–29

TQE (timer queue element)

- expiration time • 3–29
- inserting in timer queue • 3–29
- removing in timer queue • 3–48

Translation buffer

- invalidating • 2–41 to 2–42

\$TTYMACS macro • 2–7, 2–8, 2–98, 2–99, 2–100

\$TTYUCBDEF macro • 1–69

U

UBMAPEXCED bugcheck • 3–74, 3–78

UCB\$B_DEVCLASS • 2–25, 3–51

UCB\$B_DEVTYPE • 2–25, 3–51

UCB\$B_DIPL • 2–25

UCB\$B_ERTCNT • 3–69, 3–94

UCB\$B_FIPL • 1–73, 2–33

UCB\$B_FLCK • 2–25, 2–33

UCB\$L_AFFINITY • 3–71

UCB\$L_DEVCHAR • 2–25

UCB\$L_DUETIM • 3–104, 3–105

UCB\$L_EMB • 3–8

UCB\$L_IOQFL • 3–28

UCB\$L_IRP • 3–71

UCB\$L_OPCNT • 3–5, 3–24, 3–94

- adjusted by IOC\$REQCOM • 3–95

UCB\$L_ORB • 1–44

UCB\$L_SVAPTE • 1–40, 3–71, 3–79

UCB\$L_SVPN • 2–21, 3–67, 3–79

UCB\$L_TT_CLASS • 2–8

UCB\$L_TT_PORT • 2–8

UCB\$Q_DEVDEPEND • 3–49, 3–51

UCB\$V_BSY • 3–28, 3–68, 4–5

UCB\$V_CANCEL • 3–68, 3–71, 4–5

UCB\$V_ECC • 3–67

UCB\$V_ERLOGIP • 3–8, 3–95

UCB\$V_ONLINE • 1–36

UCB\$V_TEMPLATE • 4–6

UCB\$V_TIM • 2–43, 3–30, 3–104

UCB\$V_TIMOUT • 3–71, 3–104

UCB\$W_BCNT • 1–41, 1–79, 3–64, 3–66, 3–71

UCB\$W_BOFF • 1–41, 1–79, 3–64, 3–66, 3–71

UCB\$W_BUFQUO

- in mailbox UCB • 3–61

UCB\$W_DEVBUFSIZ • 3–51

- in mailbox UCB • 3–61

UCB\$W_EC1 • 3–67

UCB\$W_EC2 • 3–67

UCB\$W_ERRCNT • 3–8

UCB\$W_QLEN • 3–28

UCB\$W_REFC • 4–4

UCB (unit control block) • 1–12, 1–68 to 1–91

- as template • 1–78
- cloned • 1–31, 1–78
- creation • 1–37, 1–68
- dual path extension • 1–69
- error log extension • 1–69, 1–80 to 1–81
- extending • 1–69 to 1–70
- local disk extension • 1–69, 1–82 to 1–84, 3–9, 3–67
- local tape extension • 1–69, 1–81 to 1–82, 3–9
- logical • 1–87
- physical • 1–86
- reference count • 1–78
- remote terminal extension • 1–75
- size • 1–33, 1–69 to 1–70, 1–72, 2–22
- terminal extension • 1–69, 1–84 to 1–91

\$UCBDEF macro • 1–69

Unit delivery routine • 1–2

- address • 1–34, 2–22, 4–21
- context • 4–21
- entry point • 4–21
- exit method • 4–21
- functions • 4–21
- input • 4–21
- register usage • 4–21
- synchronization requirements • 4–21

Unit initialization routine

- address • 1–26, 1–30, 2–26, 4–22
- context • 4–22
- entry point • 4–22

Unit initialization routine (Cont.)

- exit method • 4–23
- for MASSBUS device • 1–26
- functions • 4–23
- input • 4–23
- of terminal port driver • 2–8
- register usage • 4–22
- synchronization requirements • 4–22
- UNLOCK macro • 2–96, 3–114, 3–116
- UNLOCK_SYSTEM_PAGES macro • 2–97
- Unsolicited interrupt service routine • 1–30
 - address • 4–24
 - context • 4–24
 - entry point • 4–24
 - exit method • 4–24
 - input • 4–24
 - register usage • 4–24
 - synchronization requirements • 4–24
- UNSUPRTCPU bugcheck • 2–10

V

- VAXBI node
 - mapping window space of • 3–107
- VCB (volume control block) • 1–74, 1–78
- VEC\$L_INITIAL • 4–8
- VEC\$L_ISR • 4–13
- VEC\$L_UNITINIT • 4–22
- VEC\$Q_DISPATCH • 1–25
- VEC\$V_LWAE • 3–78
- VEC\$V_MAPLOCK • 3–90
- VEC\$V_PATHLOCK • 3–87
- VEC (interrupt transfer vector) • 1–9, 1–22 to 1–27
 - multiple • 1–23
- \$VECEND macro • 2–99
 - example • 2–100
- \$VECINI macro • 2–98, 2–100
 - example • 2–100
- \$VEC macro • 2–98
 - example • 2–100
- _VIELD macro • 1–70, 2–102 to 2–103
 - example • 2–103
- \$VIELD macro • 2–102 to 2–103
- Virtual I/O function • 1–40, 1–41
- Volume • 1–78

W

- WCB (window control block) • 1–12, 1–39
- WFIKPCH macro • 2–66, 2–104 to 2–105, 3–104, 4–19
- WFIRLCH macro • 2–104 to 2–105, 3–104, 4–19
- Working set limit • 3–35, 3–41
 - insufficient • 3–33
- Workstation device • 1–76
- Write check
 - enabling • 1–75

X

- XQP (extended QIO processor) • 1–12, 1–74
 - default • 1–28

How to Order Additional Documentation

Technical Support

If you need help deciding which documentation best meets your needs, call 800-343-4040 before placing your electronic, telephone, or direct mail order.

Electronic Orders

To place an order at the Electronic Store, dial 800-DEC-DEMO (800-332-3366) using a 1200- or 2400-baud modem. If you need assistance using the Electronic Store, call 800-DIGITAL (800-344-4825).

Telephone and Direct Mail Orders

Your Location	Call	Contact
Continental USA, Alaska, or Hawaii	800-DIGITAL	Digital Equipment Corporation P.O. Box CS2008 Nashua, New Hampshire 03061
Puerto Rico	809-754-7575	Local Digital subsidiary
Canada	800-267-6215	Digital Equipment of Canada Attn: DECdirect Operations KAO2/2 P.O. Box 13000 100 Herzberg Road Kanata, Ontario, Canada K2K 2A6
International	_____	Local Digital subsidiary or approved distributor
Internal ¹	_____	USASSB Order Processing - WMO/E15 <i>or</i> U.S. Area Software Supply Business Digital Equipment Corporation Westminster, Massachusetts 01473

¹For internal orders, you must submit an Internal Software Order Form (EN-01740-07).

Reader's Comments

VMS Device Support
Reference Manual
AA-PBPXA-TE

Please use this postage-paid form to comment on this manual. If you require a written reply to a software problem and are eligible to receive one under Software Performance Report (SPR) service, submit your comments on an SPR form.

Thank you for your assistance.

I rate this manual's:	Excellent	Good	Fair	Poor
Accuracy (software works as manual says)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Completeness (enough information)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Clarity (easy to understand)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Organization (structure of subject matter)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Figures (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Examples (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Index (ability to find topic)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Page layout (easy to find information)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

I would like to see more/less _____

What I like best about this manual is _____

What I like least about this manual is _____

I found the following errors in this manual:

Page	Description
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____

Additional comments or suggestions to improve this manual:

I am using **Version** _____ of the software this manual describes.


Name/Title _____ Dept. _____
Company _____ Date _____
Mailing Address _____
_____ Phone _____

- 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

DIGITAL EQUIPMENT CORPORATION
Corporate User Publications—Spit Brook
ZK01-3/J35 110 SPIT BROOK ROAD
NASHUA, NH 03062-9987



- Do Not Tear - Fold Here -

Cut Along Dotted Line