
OpenVMS AXP Device Support: Reference

Order Number: AA-Q28PA-TE

March 1994

This manual provides reference information for the following manuals:

*Creating an OpenVMS AXP Step 2 Device Driver from an OpenVMS
VAX Device Driver*

*Creating an OpenVMS AXP Step 2 Device Driver from a Step 1 Device
Driver*

OpenVMS AXP Device Support: Developer's Guide

Revision/Update Information: This is a new manual.

Software Version: OpenVMS AXP Version 6.1

**Digital Equipment Corporation
Maynard, Massachusetts**

March 1994

Digital Equipment Corporation makes no representations that the use of its products in the manner described in this publication will not infringe on existing or future patent rights, nor do the descriptions contained in this publication imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

Possession, use, or copying of the software described in this publication is authorized only pursuant to a valid written license from Digital or an authorized sublicensor.

© Digital Equipment Corporation 1994. All rights reserved.

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: Alpha AXP, AXP, DEC, DECnet, DECwindows, Digital, HSC, OpenVMS, Q-bus, TURBOchannel, UNIBUS, VAX, VAXcluster, VAX DOCUMENT, VAX MACRO, VMScluster, the AXP logo, and the DIGITAL logo.

The following is a third-party trademark:

Futurebus/Plus is a registered trademark of Force Computers GMBH, Federal Republic of Germany.

Internet is a registered trademark of Internet, Inc.

This document is available on CD-ROM.

ZK6323

This document was prepared using VAX DOCUMENT Version 2.1.

Send Us Your Comments

We welcome your comments on this or any other OpenVMS manual. If you have suggestions for improving a particular section or find any errors, please indicate the title, order number, chapter, section, and page number (if available). We also welcome more general comments. Your input is valuable in improving future releases of our documentation.

You can send comments to us in the following ways:

- Internet electronic mail: `OPENVMSDOC@ZKO.MTS.DEC.COM`
- Fax: 603-881-0120 Attn: OpenVMS Documentation, ZK03-4/U08
- A completed Reader's Comments form (postage paid, if mailed in the United States), or a letter, via the postal service. Two Reader's Comments forms are located at the back of each printed OpenVMS manual. Please send letters and forms to:

Digital Equipment Corporation
Information Design and Consulting
OpenVMS Documentation
110 Spit Brook Road, ZK03-4/U08
Nashua, NH 03062-2698
USA

You may also use an online questionnaire to give us feedback. Print or edit the online file `SYSSHELP:OPENVMSDOC_SURVEY.TXT`. Send the completed online file by electronic mail to our Internet address, or send the completed hardcopy survey by fax or through the postal service.

Thank you.

Contents

Preface	xiii
1 Device Driver Entry Points	
Alternate Start-I/O Routine	1-2
Cancel-I/O Routine	1-4
Cancel Selective Routine	1-7
Channel Assign Routine	1-9
Cloned UCB Routine	1-11
Controller Initialization Routine	1-14
Driver Channel Grant Fork Routine Entry	1-17
Driver Device Timeout Routine Entry	1-18
Driver Resume from Interrupt Routine Entry	1-19
Start I/O Routine (Simple Fork, JSB Environment)	1-20
Driver Unloading Routine	1-21
FDT Upper-Level Action Routine	1-22
FDT Error-Handling Callback Routine	1-25
Interrupt Service Routine	1-28
Mount Verification Routine	1-31
Register Dumping Routine	1-33
Start-I/O Routine (Simple Fork, Call Environment)	1-35
Start-I/O Routine (Kernel Process)	1-38
Timeout Handling Code (Traditional)	1-40
Timeout Handling Code (Kernel Process)	1-42
Unit Delivery Routine	1-44
Unit Initialization Routine	1-46
2 System Routines	
ACP_STD\$ACCESS	2-6
ACP_STD\$ACCESSNET	2-8
ACP_STD\$DEACCESS	2-10
ACP_STD\$MODIFY	2-12
ACP_STD\$MOUNT	2-14
ACP_STD\$READBLK	2-16
ACP_STD\$WRITEBLK	2-18
COM_STD\$DELATTNAST	2-20
COM_STD\$DELATTNASTP	2-22
COM_STD\$DELCTRLAST	2-24

COM_STD\$DELCTRLASTP	2-26
COM_STD\$DRVDEALMEM	2-28
COM_STD\$FLUSHATTNS	2-30
COM_STD\$FLUSHCTRLS	2-32
COM_STD\$POST, COM_STD\$POST_NOCNT	2-34
COM_STD\$SETATTNAST	2-36
COM_STD\$SETCTRLAST	2-39
ERL_STD\$ALLOCEMB	2-42
ERL_STD\$DEVICEATTN, ERL_STD\$DEVICERR, ERL_STD\$DEVICTMO	2-44
ERL_STD\$RELEASEMB	2-47
EXE\$BUS_DELAY	2-48
EXE\$DELAY	2-50
EXE\$KP_ALLOCATE_KPB	2-51
EXE\$KP_DEALLOCATE_KPB	2-54
EXE\$KP_END	2-56
EXE\$KP_FORK	2-58
EXE\$KP_FORK_WAIT	2-60
EXE\$KP_RESTART	2-62
EXE\$KP_STALL_GENERAL	2-64
EXE\$KP_START	2-67
EXE\$KP_STARTIO	2-70
EXE\$TIMEDWAIT_COMPLETE	2-72
EXE\$TIMEDWAIT_SETUP, EXE\$TIMEDWAIT_SETUP_10US	2-74
EXE_STD\$ABORTIO	2-76
EXE_STD\$ALLOCBUF, EXE_STD\$ALLOCIRP	2-79
EXE_STD\$ALTQUEPKT	2-82
EXE_STD\$CARRIAGE	2-84
EXE_STD\$CHK _{xxx} ACCES	2-85
EXE_STD\$FINISHIO	2-87
EXE\$ILLIOFUNC	2-90
EXE_STD\$INSERT_IRP	2-92
EXE_STD\$INSIOQ, EXE_STD\$INSIOQC	2-94
EXE_STD\$IORSNWAIT	2-96
EXE_STD\$LCLDSKVALID	2-98
EXE_STD\$MNTVERSIO	2-101
EXE_STD\$MODIFY	2-103
EXE_STD\$MODIFYLOCK	2-107
EXE_STD\$MOUNT_VER	2-113
EXE_STD\$ONEPARG	2-115
EXE_STD\$PRIMITIVE_FORK	2-117
EXE_STD\$PRIMITIVE_FORK_WAIT	2-119
EXE_STD\$QIOACPPKT	2-121
EXE_STD\$QIODRVPKT	2-123
EXE_STD\$QUEUE_FORK	2-126
EXE_STD\$QXQPPKT	2-127
EXE_STD\$READ	2-129

EXE_STD\$READCHK	2-133
EXE_STD\$READLOCK	2-137
EXE_STD\$SENSEMODE	2-143
EXE_STD\$SETCHAR, EXE_STD\$SETMODE	2-145
EXE_STD\$SNDEVMSG	2-148
EXE_STD\$WRITE	2-150
EXE_STD\$WRITECHK	2-154
EXE_STD\$WRITELOCK	2-158
EXE_STD\$WRTMAILBOX	2-164
EXE_STD\$ZEROPARM	2-166
IOC\$ALOALTMAP, IOC\$ALOALTMAPN, IOC\$ALOALTMAPSP	2-168
IOC\$ALOUBAMAP, IOC\$ALOUBAMAPN	2-169
IOC\$ALLOC_CNT_RES	2-170
IOC\$ALLOC_CRAB	2-174
IOC\$ALLOC_CRCTX	2-176
IOC\$ALLOCATE_CRAM	2-178
IOC\$CANCEL_CNT_RES	2-180
IOC\$CRAM_CMD	2-182
IOC\$CRAM_IO	2-185
IOC\$CRAM_QUEUE	2-187
IOC\$CRAM_WAIT	2-189
IOC\$DEALLOC_CNT_RES	2-191
IOC\$DEALLOC_CRAB	2-193
IOC\$DEALLOC_CRCTX	2-194
IOC\$DEALLOCATE_CRAM	2-195
IOC\$SKP_REQCHAN	2-196
IOC\$SKP_WFIKPCH, IOC\$SKP_WFIRLCH	2-198
IOC\$LOAD_MAP	2-200
IOC\$MAP_IO	2-202
IOC\$NODE_FUNCTION	2-204
IOC\$READ_IO	2-207
IOC\$UNMAP_IO	2-209
IOC\$WRITE_IO	2-210
IOC_STD\$ALTREQCOM	2-211
IOC_STD\$BROADCAST	2-213
IOC_STD\$CANCELIO	2-215
IOC_STD\$CLONE_UCB	2-217
IOC_STD\$COPY_UCB	2-219
IOC_STD\$CREDIT_UCB	2-221
IOC_STD\$CVT_DEVNAM	2-222
IOC_STD\$CVTLOGPHY	2-224
IOC_STD\$DELETE_UCB	2-226
IOC_STD\$DIAGBUFILL	2-227
IOC_STD\$FILSPT	2-229
IOC_STD\$GETBYTE	2-231
IOC_STD\$INITBUFWIND	2-233
IOC_STD\$INITIATE	2-235

IOC_STD\$LINK_UCB	2-238
IOC_STD\$MAPVBLK	2-240
IOC_STD\$MNTVER	2-242
IOC_STD\$MOVFRUSER, IOC_STD\$MOVFRUSER2	2-243
IOC_STD\$MOVTOUSER, IOC_STD\$MOVTOUSER2	2-246
IOC_STD\$PARSDEVNAM	2-249
IOC_STD\$POST_IRP	2-251
IOC_STD\$PTETOPFN	2-252
IOC_STD\$QNXTSEG1	2-254
IOC_STD\$PRIMITIVE_REQCHANH, IOC_STD\$PRIMITIVE_REQCHANL	2-256
IOC_STD\$PRIMITIVE_WFIKPCH, IOC_STD\$PRIMITIVE_WFIRLCH	2-259
IOC_STD\$RELCHAN	2-262
IOC_STD\$REQCOM	2-264
IOC_STD\$SEARCHDEV	2-267
IOC_STD\$SEARCHINT	2-269
IOC_STD\$SENSEDISK	2-271
IOC_STD\$SEVER_UCB	2-273
IOC_STD\$SIMREQCOM	2-274
IOC_STD\$THREADCRB	2-276
MMG_STD\$IOLOCK	2-278
MMG_STD\$UNLOCK	2-280
MT_STD\$CHECK_ACCESS	2-282
SCH_STD\$IOLOCKR	2-284
SCH_STD\$IOLOCKW	2-286
SCH_STD\$IOUNLOCK	2-288

3 Data Structures

3.1 ADP (Adapter Control Block)	3-3
3.1.1 BUSARRAY (Bus Array)	3-10
3.2 CCB (Channel Control Block)	3-12
3.3 CRAM (Controller Register Access Mailbox)	3-13
3.4 CRB (Channel Request Block)	3-17
3.5 VEC (Interrupt Transfer Vector Block)	3-19
3.6 DDB (Device Data Block)	3-20
3.7 DDT (Driver Dispatch Table)	3-22
3.8 DPT (Driver Prologue Table)	3-25
3.9 IDB (Interrupt Dispatch Block)	3-30
3.10 IRP (I/O Request Packet)	3-32
3.11 IRPE (I/O Request Packet Extension)	3-38
3.12 KPB (Kernel Process Block)	3-39
3.13 ORB (Object Rights Block)	3-46
3.14 UCB (Unit Control Block)	3-47
3.15 VLE (Vector List Extension)	3-68

4 MACRO-32 Driver Macros

CALL_ABORTIO	4-7
CALL_ALLOCBUF, CALL_ALLOCIRP	4-8
CALL_ALLOCEMB	4-9
CALL_ALTQUEPKT	4-10
CALL_ALTREQCOM	4-11
CALL_BROADCAST	4-12
CALL_CANCELIO	4-13
CALL_CARRIAGE	4-14
CALL_CHKxxxACCES	4-15
CALL_CLONE_UCB	4-16
CALL_COPY_UCB	4-17
CALL_CREDIT_UCB	4-18
CALL_CVTLOGPHY	4-19
CALL_CVT_DEVNAM	4-20
CALL_DELATTNAST	4-21
CALL_DELATTNASTP	4-22
CALL_DELCTRLAST	4-23
CALL_DELCTRLASTP	4-24
CALL_DELETE_UCB	4-25
CALL_DEVICEATTN, CALL_DEVICERR, CALL_DEVICTMO	4-26
CALL_DIAGBUFILL	4-27
CALL_DRVDEALMEM	4-28
CALL_FILSPT	4-29
CALL_FINISHIO, CALL_FINISHIOC, CALL_FINISHIO_NOIOST	4-30
CALL_FLUSHATTNS	4-31
CALL_FLUSHCTRLS	4-32
CALL_GETBYTE	4-33
CALL_INITBUFWIND	4-34
CALL_INITIATE	4-35
CALL_INSERT_IRP	4-36
CALL_IOLOCK	4-37
CALL_IOLOCKR	4-38
CALL_IOLOCKW	4-39
CALL_IORSNWAIT	4-40
CALL_IOUNLOCK	4-41
CALL_LINK_UCB	4-42
CALL_MAPVBLK	4-43
CALL_MNTVER	4-44
CALL_MNTVERSIO	4-45
CALL_MODIFYLOCK, CALL_MODIFYLOCK_ERR	4-46
CALL_MOUNT_VER	4-47
CALL_MOVFRUSER, CALL_MOVFRUSER2	4-48
CALL_MOVTOUSER, CALL_MOVTOUSER2	4-49
CALL_PARSDEVNAM	4-50
CALL_POST, CALL_POST_NOCNT	4-51

CALL_POST_IRP	4-52
CALL_PTETOPFN	4-53
CALL_QIOACPPKT	4-54
CALL_QIODRVPKT	4-55
CALL_QNXTSEG1	4-56
CALL_QXQPPKT	4-57
CALL_READCHK, CALL_READCHKR	4-58
CALL_READLOCK, CALL_READLOCK_ERR	4-59
CALL_RELCHAN	4-60
CALL_RELEASEMB	4-61
CALL_REQCOM	4-62
CALL_SEARCHDEV	4-63
CALL_SEARCHINT	4-64
CALL_SETATTNAST	4-65
CALL_SETCTRLAST	4-66
CALL_SEVER_UCB	4-67
CALL_SIMREQCOM	4-68
CALL_SNDEVMSG	4-69
CALL_THREADCRB	4-70
CALL_UNLOCK	4-71
CALL_WRITECHK, CALL_WRITECHKR	4-72
CALL_WRITELOCK, CALL_WRITELOCK_ERR	4-73
CALL_WRTMAILBOX	4-74
CLASS_UNIT_INIT	4-75
CPUDISP	4-77
CRAM_ALLOC	4-78
CRAM_CMD	4-79
CRAM_DEALLOC	4-81
CRAM_IO	4-82
CRAM_QUEUE	4-83
CRAM_WAIT	4-84
DDTAB	4-85
DEVICELOCK	4-89
DPTAB	4-91
DPT_STORE	4-96
DPT_STORE_ISR	4-99
\$DRIVER_ALTSTART_ENTRY	4-100
\$DRIVER_CANCEL_ENTRY	4-101
\$DRIVER_CANCEL_SELECTIVE	4-102
\$DRIVER_CHANNEL_ASSIGN	4-103
\$DRIVER_CLONEDUCB	4-104
DRIVER_CODE	4-105
\$DRIVER_CRTLINIT	4-106
\$DRIVER_DELIVER_ENTRY	4-107
\$DRIVER_ERRRTN	4-108
\$DRIVER_FDT_ENTRY	4-109
\$DRIVER_MNTVER	4-110

\$DRIVER_REGDUMP	4-111
\$DRIVER_START_ENTRY.....	4-112
\$DRIVER_UNITINIT.....	4-113
DRIVER_DATA	4-114
\$FDTARGDEF	4-115
FDT_ACT.....	4-116
FDT_BUF	4-118
FDT_INI	4-119
FORK.....	4-120
FORK_ROUTINE.....	4-122
FORK_WAIT	4-123
FORKLOCK.....	4-125
IOFORK.....	4-127
IFNORD, IFNOWRT, IFRD, IFWRT	4-129
KP_ALLOCATE_KPB	4-132
KP_DEALLOCATE_KPB	4-133
KP_END	4-134
KP_RESTART	4-135
KP_REQCOM	4-136
KP_STALL_FORK, KP_STALL_IOFORK.....	4-137
KP_STALL_FORK_WAIT.....	4-138
KP_STALL_GENERAL	4-139
KP_STALL_REQCHAN	4-140
KP_STALL_WFIKPCH, KP_STALL_WFIRLCH.....	4-141
KP_START.....	4-142
KP_SWITCH_TO_KP_STACK	4-143
LOCK.....	4-144
RELCHAN	4-146
REQCHAN.....	4-147
REQCOM.....	4-149
REQPCHAN.....	4-150
SYSDISP	4-151
TBI_ALL	4-152
TBI_DATA_64	4-153
TBI_SINGLE	4-154
TBI_SINGLE_64	4-155
TIMEDWAIT	4-156
WFIKPCH, WFIRLCH.....	4-159

5 C Driver Macros

DEVICE_LOCK	5-2
DEVICE_UNLOCK	5-4
FORK.....	5-6
FORK_LOCK	5-7
FORK_UNLOCK	5-8
FORK_WAIT	5-10

IOFORK	5-11
RFI	5-12
WFIKPCH	5-13
WFIRLCH	5-14

Index

Figures

3-1	I/O Database	3-3
3-2	ADP List	3-4
3-3	ADP Hierarchy and System Configuration	3-5
3-4	Composition of Extended Unit Control Blocks	3-49

Tables

2-1	New, Changed, and Unsupported OpenVMS System Routines	2-1
2-2	Kernel Process Stall Jacket Routines and Scheduling Stall Routines	2-65
3-1	Contents of Adapter Control Block	3-6
3-2	Contents of Bus Array	3-10
3-3	Contents of Bus Array	3-11
3-4	Contents of Channel Control Block	3-12
3-5	Contents of Controller Register Access Mailbox	3-14
3-6	Contents of Channel Request Block	3-17
3-7	Contents of Interrupt Transfer Vector Block (VEC)	3-20
3-8	Contents of Device Data Block	3-21
3-9	Contents of Driver Dispatch Table	3-22
3-10	Contents of Driver Prologue Table	3-26
3-11	Contents of Interrupt Dispatch Block	3-30
3-12	Contents of I/O Request Packet (IRP)	3-33
3-13	Contents of I/O Request Packet Extension (IRPE)	3-38
3-14	Contents of Kernel Process Block (KPB)	3-40
3-15	Contents of KPB Debug Area	3-46
3-16	Contents of Object Rights Block	3-46
3-17	UCB Extensions and Sizes Defined in \$UCBDEF	3-48
3-18	Contents of Unit Control Block	3-50
3-19	Contents of UCB Error Log Extension	3-60
3-20	Contents of UCB Local Tape Extension	3-61
3-21	Contents of UCB Local Disk Extension	3-61
3-22	Contents of UCB Terminal Extension	3-62
3-23	Contents of the Vector List Extension	3-69
4-1	New, Changed, and Unsupported OpenVMS Driver Macros	4-1

Preface

OpenVMS AXP Version 6.1 introduces formal support for user-written device drivers and a new device driver interface known as the **Step 2** driver interface. The Step 2 driver interface replaces the temporary **Step 1** driver interface that was provided in OpenVMS AXP Versions 1.0 and 1.5.

This manual describes the entry points, system routines, data structures, and macros used in OpenVMS AXP device drivers.

Intended Audience

OpenVMS AXP Device Support: Reference is intended for software engineers who must prepare a Step 2 device driver to run on OpenVMS AXP Version 6.1 or engineers who want to write an OpenVMS AXP device driver in a high-level language.

This manual assumes that its reader is familiar with the components of OpenVMS VAX device drivers and Step 1 OpenVMS AXP device drivers. It also relies on a familiarity with the software interfaces within the OpenVMS operating system that support device drivers.

Document Structure

This manual contains the following sections:

- Chapter 1 provides specific information about how each driver entry point is defined and accessed in an OpenVMS AXP driver.
- Chapter 2 includes call-based OpenVMS system routines that support Step 2 OpenVMS AXP drivers and system routines these drivers may use.
- Chapter 3 describes the data structures in the I/O database.
- Chapter 4 documents the OpenVMS macros that have been changed or augmented to provide for Step 2 OpenVMS AXP drivers. It also introduces new macros these drivers may use.
- Chapter 5 describes OpenVMS C Driver macros.

Associated Document

Before using *OpenVMS AXP Device Support: Reference*, you should understand the information in the following manuals:

- *Creating an OpenVMS AXP Step 2 Device Driver from a Step 1 Device Driver*
- *Creating an OpenVMS AXP Step 2 Device Driver from an OpenVMS VAX Device Driver*
- *OpenVMS AXP Device Support: Developer's Guide*

Conventions

In this manual, every use of OpenVMS VAX means the OpenVMS VAX operating system, every use of OpenVMS AXP means the OpenVMS AXP operating system, and every use of OpenVMS means both the OpenVMS VAX operating system and the OpenVMS AXP operating system.

The following conventions are used in this manual:

Ctrl/ <i>x</i>	A sequence such as Ctrl/ <i>x</i> indicates that you must hold down the key labeled Ctrl while you press another key or a pointing device button.
PF1 <i>x</i>	A sequence such as PF1 <i>x</i> indicates that you must first press and release the key labeled PF1, then press and release another key or a pointing device button.
Return	In examples, a key name enclosed in a box indicates that you press a key on the keyboard. (In text, a key name is not enclosed in a box.)
...	A horizontal ellipsis in examples indicates one of the following possibilities: <ul style="list-style-type: none">• Additional optional arguments in a statement have been omitted.• The preceding item or items can be repeated one or more times.• Additional parameters, values, or other information can be entered.
. . .	A vertical ellipsis indicates the omission of items from a code example or command format; the items are omitted because they are not important to the topic being discussed.
()	In format descriptions, parentheses indicate that, if you choose more than one option, you must enclose the choices in parentheses.
[]	In format descriptions, brackets indicate optional elements. You can choose one, none, or all of the options. (Brackets are not optional, however, in the syntax of a directory name in an OpenVMS file specification, or in the syntax of a substring specification in an assignment statement.)
{ }	In format descriptions, braces surround a required choice of options; you must choose one of the options listed.
boldface text	<p>Boldface text represents the introduction of a new term or the name of an argument, an attribute, or a reason.</p> <p>Boldface text is also used to show user input in online versions of the manual.</p>
<i>italic text</i>	Italic text emphasizes important information, indicates variables, and indicates complete titles of manuals. Italic text also represents information that can vary in system messages (for example, Internal error <i>number</i>), command lines (for example, /PRODUCER= <i>name</i>), and command parameters in text.
UPPERCASE TEXT	Uppercase text indicates a command, the name of a routine, the name of a file, the name of a file protection code, or the abbreviation for a system privilege.

-
numbers

A hyphen in code examples indicates that additional arguments to the request are provided on the line that follows.

All numbers in text are assumed to be decimal, unless otherwise noted. Nondecimal radices—binary, octal, or hexadecimal—are explicitly indicated.

Device Driver Entry Points

This chapter describes the standard driver routines that OpenVMS AXP uses as entry points in a device driver program.

Unlike OpenVMS VAX, OpenVMS AXP does not support driver unloading routines and unsolicited interrupt handling routines.

This chapter also describes the Step 2 driver-entry-point macros that replace the `.JSB_ENTRY` directive used in the Step 1 OpenVMS AXP driver entry points. These macros perform the following operations:

1. Declare a call entry point.
2. Specify a register save list that consists of the registers that the Step 1 JSB interface was allowed to scratch. This save list augments the list of autopreserved registers detected by the MACRO-32 compiler. You can specify an alternative save list if you are certain that the default mask contains registers that are not used in the execution path initiated by the entry point.
3. Define symbolic AP offsets that correspond to the routine parameters.
4. Copy the input parameters into the registers that correspond to the input registers of the Step 1 JSB interface. You can disable this register loading by using an optional parameter.

For more information about optimizing the use of the driver-entry-point macros, see the *OpenVMS AXP Device Support: Developer's Guide*.

OpenVMS AXP 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.

Format

ALTSTART (irp, ucb)

Arguments

Argument	Type	Access	Mechanism
irp	IRP	input	reference
ucb	UCB	input	reference

irp

I/O request packet for the current I/O request

ucb

Unit control block of the device that is the target of the I/O request

Essentials

Identifying the Routine

Specify the address of the alternate start-I/O routine in the **altstart** argument to the DDTAB macro. This macro places the procedure value of the routine into the DDT.

Declaring the Entry Point

Use:

```
$DRIVER_ALTSTART_ENTRY [preserve=<R2,R3,R4,R5>] [,fetch=YES]
```

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO-32 compiler) across the call to the alternate start-I/O routine

fetch=YES, the default, loads the addresses of the IRP and UCB into R3 and R5, respectively; **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver alternate start-I/O routine that uses this macro can access any of its arguments by using a symbolic name of the form **ALTARG\$_argument-name**.

Called by

Called by routine EXE_STDSALTQUEPKT in module SYSQIOREQ. A driver FDT routine typically is the caller of EXE_STDSALTQUEPKT.

Context

An alternate start-I/O routine begins execution at fork IPL, holding the corresponding fork lock. It must return control to EXE_STDSALTQUEPKT in this context.

OpenVMS AXP Device Driver Entry Points Alternate Start-I/O Routine

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.

Exit mechanism

The alternate start-I/O routine completes I/O requests by calling COM_STDSPOST. 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_STDSALTQUEPKT, which relinquishes fork level synchronization and returns to the driver FDT routine that called it. The FDT routine performs any required postprocessing and returns the SSS_FDT_COMPL status to its caller.

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.

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 use separate fork blocks for each request.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate the entry point of the routine with a \$DRIVER_ALTSTART_ENTRY macro, indicating which registers must be saved and restored across routine execution.
- You must replace direct CSR access (for instance, by means of a MOVL instruction) with CSR access by means of a CRAM.
- You should examine the routine's use of suspension mechanisms (for instance, its forking, wait-for-interrupt, and resource-wait semantics) to determine whether it needs to be adapted to use the kernel process services. Typically a driver that makes subroutine calls before suspending itself (and relies on the previous context of these subroutines remaining intact on the stack), must be adapted to use the kernel process services.
- If the routine need not be converted to a kernel process, you should replace any calls to EXE\$FORK, EXE\$FORK_WAIT, EXE\$IOFORK, IOC\$WFIKPCH, IOC\$WFIRLCH, IOC\$REQCHANH, and IOC\$REQCHANL with invocations of the appropriate suspension macro or with calls to EXE_STDS\$PRIMITIVE_FORK, IOC_STDS\$PRIMITIVE_WFIKPCH, IOC_STDS\$PRIMITIVE_WFIRLCH, IOC_STDS\$PRIMITIVE_REQCHANH, or IOC_STDS\$PRIMITIVE_REQCHANL.

OpenVMS AXP 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.

Format

CANCEL (chan, irp, pcb, ucb, reason)

Arguments

Argument	Type	Access	Mechanism
chan	integer	input	value
irp	IRP	input	reference
pcb	PCB	input	reference
ucb	UCB	input	reference
reason	integer	input	value

chan

Channel index number.

irp

I/O request packet, if any, for device (contents of UCB\$SL_IRP).

pcb

Process control block of process for which the I/O request is being canceled.

ucb

Unit control block.

reason

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

Essentials

Identifying the Routine

Supply the address of the cancel-I/O routine in the **cancel** argument of the DDTAB macro. The macro places the procedure value of this routine into DDT. Many drivers specify the system routine IOC_STD\$CANCELIO as their cancel-I/O routine.

Declaring the Entry Point

Use:

```
$DRIVER_CANCEL_ENTRY [preserve=<R2,R3,R4>] [,fetch=YES]
```

OpenVMS AXP Device Driver Entry Points Cancel-I/O Routine

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO-32 compiler) across the call to the cancel I/O routine

fetch=YES, the default, loads the channel index number into R2, the cancellation reason into R8, and the addresses of the IRP, PCB, and UCB into R3, R4, and R5, respectively. **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver cancel-I/O routine that uses this macro can access any of its arguments by using a symbolic name of the form **CANARG\$_argument-name**.

Called by

System 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 (UCBSL_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

Context

A cancel-I/O routine begins execution at fork IPL, holding the corresponding fork lock. It must return control to its caller in this context.

A cancel-I/O routine executes in kernel mode in the context of the caller of the \$CANCEL, \$DALLOC, or \$DASSGN system service.

Exit mechanism

The cancel-I/O routine returns 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 (UCBSV_BSY in UCBSL_STS).
2. Confirm that the process ID (PID) of the request the device is servicing (IRPSL_PID) matches that of the process requesting the cancellation (PCBSL_PID).
3. Confirm that the channel-index number of the request the device is servicing (IRPSL_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 UCBSV_CANCEL in the UCBSL_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

OpenVMS AXP Device Driver Entry Points Cancel-I/O Routine

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.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate the entry point of a cancel-I/O routine with a `$DRIVER_CANCEL_ENTRY` macro, indicating which registers must be saved and restored across routine execution.

Cancel Selective Routine

Performs additional processing on a list of I/O requests that have been canceled.

Format

status=CANCEL_SELECTIVE (pcb, ucb, chan, iosb_vector, iosb_count)

Arguments

Argument	Type	Access	Mechanism
pcb	PCB	input	reference
ucb	UCB	input	reference
chan	integer	input	value
iosb_vector	address	input	value
iosb_count	integer	input	value

pcb

Process control block of process for which the I/O request is being canceled.

ucb

Unit control block.

chan

Channel index number.

iosb_vector

Vector of address of I/O status blocks (IOSBs), or zero.

iosb_count

Number of addresses in the IOSB vector.

Essentials

Identifying the Routine

Supply the address of the cancel selective routine in the **cancel_selective** argument of the DDTAB macro. The macro places the procedure value of this routine into DDT.

Declaring the Entry Point

Use:

```
$DRIVER_CANCEL_SELECTIVE_ENTRY [preserve=<>] [,fetch=YES]
```

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO-32 compiler) across the call to the cancel selective routine.

OpenVMS AXP Device Driver Entry Points

Cancel Selective Routine

fetch=YES, the default, loads `SS$_UNSUPPORTED` status into R0, the IOSB vector into R7, the IOSB count into R8, and the addresses of the PCB and UCB into R4 and R5, respectively, **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver cancel selective routine that uses this macro can access any of its arguments by using a symbolic name of the form `CANSARG$_argument-name`.

Called by

`EX$CANCEL_SELECTIVE` calls a driver's cancel selective routine.

Context

A cancel selective routine is called at device IPL, holding the corresponding device lock and the appropriate fork lock. The channel control block (CCB) is locked in memory. It must return control to `EX$CANCEL_SELECTIVE` in this context.

Exit mechanism

The cancel selective routine returns to its caller.

Description

Reserved to Digital.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note that you must indicate the entry point of a cancel-I/O routine with a `$DRIVER_CANCEL_SELECTIVE_ENTRY` macro, indicating which registers must be saved and restored across routine execution.

Channel Assign Routine

Performs specialized operations when a channel is assigned to a non-network device.

Format

CHANNEL_ASSIGN (ucb, ccb)

Arguments

Argument	Type	Access	Mechanism
ucb	UCB	input	reference
ccb	CCB	input	reference

ucb
Unit control block.

ccb
Channel control block.

Essentials

Identifying the Routine

Supply the address of the channel assign routine in the **channel_assign** argument of the DDTAB macro. The macro places the procedure value of this routine into DDT.

Declaring the Entry Point

Use:

```
$DRIVER_CHANNEL_ASSIGN_ENTRY [preserve=<>] [,fetch=YES]
```

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO-32 compiler) across the call to the cancel selective routine.

fetch=YES, the default, loads the addresses of UCB and CCB into R5 and R8, respectively, **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver channel assign routine that uses this macro can access any of its arguments by using a symbolic name of the form **CHANARG\$_argument-name**.

Called by

EXE\$ASSIGN_LOCAL (in module SYSASSIGN) calls a driver's channel assign routine.

Context

A channel assign routine is called in kernel mode at IPL 0.

OpenVMS AXP Device Driver Entry Points Channel Assign Routine

Exit mechanism

The channel assign routine returns to its caller.

Description

Reserved to Digital.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note that you must indicate the entry point of a channel assign routine with a `$DRIVER_CHANNEL_ASSIGN_ENTRY` macro, indicating which registers must be saved and restored across routine execution.

Cloned UCB Routine

Completes the initialization of the UCB cloned when a channel is requested for a template device.

Format

status = CLONEDUCB (cloned_ucb, ddt, pcb, template_ucb)

Arguments

Argument	Type	Access	Mechanism
cloned_ucb	UCB	input	reference
ddt	DDT	input	reference
pcb	PCB	input	reference
template_ucb	UCB	input	reference

cloned_ucb

Cloned unit control block. Fields of the cloned UCB have been initialized as follows:

Field	Value
UCB\$\$_FQFL	Address of UCB\$\$_FQFL
UCB\$\$_FQBL	Address of UCB\$\$_FQFL
UCB\$\$_FPC	0
UCB\$\$_FR3	0
UCB\$\$_FR4	0
UCB\$\$_BUFQUO	0
UCB\$\$_LINK	Address of next UCB in DDB chain
UCB\$\$_IOQFL	Address of UCB\$\$_IOQFL
UCB\$\$_IOQBL	Address of UCB\$\$_IOQFL
UCB\$\$_UNIT	Device unit number
UCB\$\$_CHARGE	Mailbox byte quota charge (UCB\$\$_SIZE)
UCB\$\$_REFC	0
UCB\$\$_STS	UCB\$\$_DELETEUCB set, UCB\$\$_ONLINE set
UCB\$\$_DEVSTS	UCB\$\$_DELMBX set if DEV\$\$_MBX is set in UCB\$\$_DEVCHAR

OpenVMS AXP Device Driver Entry Points Cloned UCB Routine

Field	Value
UCB\$_OPCNT	0
UCB\$_SVAPTE	0
UCB\$_BOFF	0
UCB\$_BCNT	0
UCB\$_ORB	Address of object rights block (ORB) for the cloned UCB

The cloned UCB ORB is initialized using the template UCB ORB. You can modify the ORB on the template UCB using the DCL SET SECURITY command.

ddt

Driver dispatch table.

pcb

Process control block of the current process.

template_ucb

Template unit control block.

Essentials

Identifying the Routine

Specify the address of a cloned UCB routine in the **cloneducb** argument of the DDTAB macro. The macro places the procedure value of the routine into the DDT. Only drivers for template devices, such as mailboxes, specify a cloned UCB routine.

Declaring the Entry Point

Use:

```
$DRIVER_CLONEDUCB_ENTRY [preserve=<R3>] [,fetch=YES]
```

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO-32 compiler) across the call to the cloned UCB routine

fetch=YES, the default, loads SSS_NORMAL status into R0, and the addresses of the cloned UCB, DDT, PCB, and template UCB into R2, R3, R4, and R5, respectively; **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver cloned UCB routine that uses this macro can access any of its arguments by using a symbolic name of the form CLONEARG\$_**argument-name**.

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\$_TEMPLATE in UCB\$_STS is set).

Context

A cloned UCB routine executes at IPL\$_ASTDEL, holding the I/O database mutex (IOC\$_MUTEX).

A cloned UCB routine executes in kernel mode in the context of the process that called the \$ASSIGN system service.

OpenVMS AXP Device Driver Entry Points Cloned UCB Routine

Exit mechanism

A cloned UCB routine must return control and status 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.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note that you must indicate the entry point of a cloned UCB with a \$DRIVER_CLONEDUCB_ENTRY macro, indicating which registers must be saved and restored across routine execution.

OpenVMS AXP Device Driver Entry Points

Controller Initialization Routine

Controller Initialization Routine

Prepares a controller for operation.

Format

status = CTRLINIT (idb, ddb, crb)

Arguments

Argument	Type	Access	Mechanism
idb	IDB	input	reference
ddb	DDB	input	reference
crb	CRB	input	reference

idb

Interrupt dispatch block associated with the controller.

ddb

Device data block associated with the controller.

crb

Controller request block.

Essentials

Identifying the Routine

Specify the address of a controller initialization routine in the **ctrlinit** argument of the DDTAB macro. The macro places the procedure value of this routine into the DDT.

Declaring the Entry Point

Use:

```
$DRIVER_CTRLINIT_ENTRY [preserve=<R2>] [,fetch=YES]
```

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO-32 compiler) across the call to the controller initialization routine.

fetch=YES, the default, loads SSS_NORMAL status into R0, the address of the IDB into R4 and R5, and the addresses of the DDB and CRB into R6 and R8, respectively; **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver controller initialization routine that uses this macro can access any of its arguments by using a symbolic name of the form CTRLARG\$_**argument-name**.

OpenVMS AXP Device Driver Entry Points Controller Initialization Routine

Called by

The driver-loading procedure calls a driver's controller initialization routine when processing a CONNECT command. Also, the system calls this routine if the device, controller, processor, or adapter to which the device is connected experiences a power failure.

Context

OpenVMS 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 the driver-loading procedure 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.

Because a controller initialization routine executes within system context, it can refer only to those virtual addresses that reside in system (S0) space.

Exit mechanism

The controller initialization routine returns success or failure status to its caller.

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:

- Determines whether it is being called as a result of a power failure by examining the power bit (UCBSV_POWER in UCBSL_STS) in the UCB. A controller initialization routine may want to perform or avoid specific tasks when servicing a power failure.
- Clears error-status bits in device registers.
- Enables controller interrupts.
- Allocates resources that must be permanently allocated to the controller.
- If the controller is dedicated to a single-unit device, such as a printer, fills in IDB\$PS_OWNER and set the online bit (UCBSV_ONLINE in UCBSL_STS).
- Initializes the interrupt vectors of devices with programmable interrupt vectors.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate the entry point of the routine with a \$DRIVER_CTRLINIT_ENTRY macro to indicate which registers must be saved and restored across routine execution.

OpenVMS AXP Device Driver Entry Points Controller Initialization Routine

- An OpenVMS VAX device driver specifies a controller initialization routine by invoking the `DPT_STORE` macro to place its procedure value into the interrupt transfer vector block (`CRB$L_INTD+VECSL_INITIAL`). An OpenVMS AXP device driver specifies the routine in the **`ctrlinit`** argument of the `DDTAB` macro.
- You must replace direct CSR access (for instance, by means of a `MOVL` instruction) with CSR access by means of a `CRAM`.
- The controller initialization routine of an OpenVMS VAX device driver receives the addresses of the device CSR in R4 and the IDB in R5. An OpenVMS AXP device driver's controller initialization routine is not passed the address of the CSR. It may access the controller's register by means of the controller register access mailbox (`CRAM`), the address of which is provided in `IDB$PS_CRAM`.
- A controller initialization routine that must initialize the programmable interrupt vectors of a device does so by referring to the vector offset placed in `IDB$L_VECTOR` by the driver-loading procedure. For a device with multiple interrupt vectors, `IDB$L_VECTOR` contains the address of a vector list extension (`VLE`) which contains a list of vector offsets.
- An OpenVMS AXP controller initialization routine must return success or failure status to its caller.

Driver Channel Grant Fork Routine Entry

Enabled via the IOC_STDSREQCHANx or IOC\$REQCHANx routines if the CRB is not immediately available. The procedure value of the grant routine is contained in ucb->ucb\$!_fpc. The grant routine is invoked by IOC_STDSRELCHAN which has been enhanced to support both the JSB interface and the new standard call interface. The above also applies to IOC\$RELCHAN which is now simply a JSB-to-CALL interface jacket routine around IOC_STDSRELCHAN.

Description

The JSB interface for the channel grant routine is:

JSB driver_channel_grant_routine

Inputs:

R3	contains a pointer to the IRP,
R4	contains a pointer to the IDB,
R5	contains a pointer to the UCB.

Outputs:

R0-R5	may be scratched by the routine.
-------	----------------------------------

The standard call interface for the channel grant routine is:

```
void driver_channel_grant_routine (IRP *irp, IDB *idb, UCB *ucb);
```

Inputs:

irp	is a pointer to the IRP,
idb	is a pointer to the IDB,
ucb	is a pointer to the UCB.

Driver Device Timeout Routine Entry

Enabled by the WFIKPCH or WFIRLCH macros and invoked by the EXE\$TIMEOUT routine. The EXE\$TIMEOUT routine supports both timeout routines using the JSB interface and the standard call interface.

Description

The JSB interface for the interrupt timeout routine is:

JSB driver_timeout_routine

Inputs:

R3	contains a pointer to the IRP from UCB\$Q_FR3(R5),
R4	contains the 64-bit value from UCB\$Q_FR4(R5),
R5	contains a pointer to the UCB.

Outputs:

R0-R4	may be scratched by the routine.
-------	----------------------------------

The standard call interface for the interrupt timeout routine is:

```
void driver_timeout_routine (IRP *irp, int64 fr4, UCB *ucb);
```

Inputs:

irp	is a pointer to the IRP from ucb->ucb\$q_fr3,
fr4	is the 64-bit value from ucb->fkb\$q_fr4,
ucb	is a pointer to the UCB,

The procedure value of the driver interrupt timeout routine is found in ucb->ucb\$ps_toutroun.

Note

By default the WFIKPCH macro and the IOC\$PRIMITIVE_WFIKPCH JSB interface routine set the ucb\$ps_toutroun cell to contain the same value as ucb\$l_fpc.

Driver Resume from Interrupt Routine Entry

The driver resume from interrupt routine is setup by the WFIKPCH macro and is invoked by the driver's interrupt service routine.

Description

The JSB interface for the driver interrupt resume routine is:

JSB driver_resume_routine

Inputs:

R3 contains a pointer to the IRP from UCB\$Q_FR3(R5),
R4 contains the 64-bit value from UCB\$Q_FR4(R5),
R5 contains a pointer to the UCB.

Outputs:

R0-R4 may be scratched by the routine.

The recommended standard call interface for the driver resume from interrupt routine is:

```
void driver_resume_routine (IRP *irp, int64 fr4, UCB *ucb);
```

Inputs:

irp is a pointer to the IRP from ucb->ucb\$q_fr3,
fr4 is the 64-bit value from ucb->fkb\$q_fr4,
ucb is a pointer to the UCB,

Note

The resume from interrupt routine interface must conform exactly to the calling convention used in the interrupt service routine in that driver. This differs from other routines, for example the interrupt timeout routine, which could be written to use either the traditional or the new interface.

It may be possible to eliminate the driver resume from interrupt routine by moving some processing directly into the interrupt service routine and by resuming the driver in a fork routine.

OpenVMS AXP Device Driver Entry Points
Start I/O Routine (Simple Fork, JSB Environment)

Start I/O Routine (Simple Fork, JSB Environment)

Activates a device to process a requested I/O function.

Driver Unloading Routine

*****Not supported in OpenVMS AXP drivers*****

OpenVMS AXP Device Driver Entry Points

FDT Upper-Level Action Routine

FDT Upper-Level Action Routine

Performs any device-dependent activities needed to prepare the I/O database to process an I/O request.

Format

status = driver_FDT_routine (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism
irp	IRP	input	reference
pcb	PCB	input	reference
ucb	UCB	input	reference
ccb	CCB	input	reference

irp

I/O request packet for the current I/O request. An FDT routine may read the following IRP fields:

Field	Contents
IRP\$ <i>L</i> _FUNC	I/O function code supplied in the \$QIO request
IRP\$ <i>L</i> _QIO_P <i>n</i>	Function-specific \$QIO system service arguments (p1 through p6); <i>n</i> corresponds to an integer from 1 to 6.

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb

Channel control block that describes the process-I/O channel.

Essentials

Identifying the Routine

Use the FDT_ACT macro to insert the procedure value of an upper-level FDT action routine into the FDT action routine vector slot that corresponds to a specified I/O function code.

Declaring the Entry Point

Use:

```
$DRIVER_FDT_ENTRY
```

```
[preserve=<R2,R3,R4,R5,R6,R7,R8,R9,R10,R11,R12,R13,R14,R15>] [,fetch=YES]
```

OpenVMS AXP Device Driver Entry Points FDT Upper-Level Action Routine

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO compiler) across the call to the upper-level FDT action routine.

fetch=YES, the default, loads the addresses of the IRP, PCB, UCB, and CCB into R3, R4, R5, and R6, respectively; **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver upper-level FDT action routine that uses this macro can access any of its arguments by using a symbolic name of the form **FDTARG\$_argument-name**.

Called by

The \$QIO system service calls a driver's upper-level FDT action routine from the module SYSQIOREQ. An upper-level FDT action routine can call any number of FDT support routines, as long as each routine returns control and status to the upper-level routine.

Context

An FDT routine is called at IPL\$_ASTDEL and must exit at IPL\$_ASTDEL. An FDT routine must not lower IPL below IPL\$_ASTDEL. If it raises IPL, it must lower it to IPL\$_ASTDEL before passing control to any other code. Similarly, before exiting, it must release any spin locks it may have acquired in an OpenVMS multiprocessing environment.

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.

Exit mechanism

An FDT routine must return control and status to its caller. An upper-level FDT action routine returns SSS_FDT_COMPL status to the \$QIO system service and passes the return status to be delivered to the caller of \$QIO in the FDT_CONTEXT structure.

Description

An upper-level FDT routine (and any FDT support routine it may call) validates the function-dependent arguments to a \$QIO system service request and prepares the I/O database to service the request. For each function that a device supports, an upper-level FDT action routine must provide preprocessing of requests for that function. FDT processing may complete a function that does not involve an I/O transfer. Otherwise FDT processing can abort the request or deliver it to the driver.

An OpenVMS AXP upper-level FDT action routine can invoke the \$FDTARGDEF macro, defined in SYSSLIBRARY:LIB.MLB, to provide symbolic names for the standard AP offsets of the four parameters provided as input (IRP, PCB, UCB, and CCB) to all upper-level FDT action routines. A routine that does so can use names of the form **FDTARG\$_xxx**, where *xxx* is the 3-letter structure acronym, to access the input parameters.

OpenVMS AXP Device Driver Entry Points FDT Upper-Level Action Routine

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate the entry point of each upper-level FDT action routine with a `$DRIVER_FDT_ENTRY` macro, indicating which registers must be saved and restored across routine execution.
- You should examine an FDT routine's use of R0, R7, and R8.

The FDT routine of an OpenVMS VAX device driver may obtain the address of FDT routine being called from R0, the number of the bit that specifies the code for the requested I/O function from R7, and the address of the entry in the function decision table that dispatched control to this FDT routine.

An OpenVMS AXP driver can obtain the user-supplied function code from `IRPSL_FUNC`. It can obtain the address of the start of the FDT from `DDT$PS_FDT2`. The DDT address is available from `UCB$SL_DDT`.

- An FDT routine of an OpenVMS VAX device driver accesses values of the function-dependent arguments specified in the `$QIO` request as offsets from the value of the AP; an OpenVMS AXP device driver obtains them from the IRP (at symbolic offsets `IRPSL_QIO_P1` through `IRPSL_QIO_P6`).

FDT Error-Handling Callback Routine

Processes error conditions that occur during EXE_STD\$READLOCK, EXE_STD\$WRITELOCK, and EXE_STD\$MODIFYLOCK processing.

Format

status = error_callback (irp, pcb, ucb, ccb, status)

Arguments

Argument	Type	Access	Mechanism
irp	IRP	input	reference
pcb	PCB	input	reference
ucb	UCB	input	reference
ccb	CCB	input	reference
status	integer	input	value

irp

I/O request packet for the current I/O request.

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb

Channel control block that describes the process-I/O channel.

status

Error status returned by buffer accessibility check (SS\$_ACCVIO or SS\$_BADPARAM) or buffer locking operation (SS\$_ACCVIO, SS\$_INSFWSL, or page fault status).

Essentials

Identifying the Routine

Use the **errtn** argument in a call to EXE_STD\$MODIFYLOCK, EXE_STD\$READLOCK, or EXE_STD\$WRITELOCK.

Declaring the Entry Point

Use:

```
$DRIVER_ERRRTN_ENTRY [preserve=<>] [,fetch=YES]
```

OpenVMS AXP Device Driver Entry Points FDT Error-Handling Callback Routine

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO compiler) across the call to the upper-level FDT action routine.

fetch=YES, the default, loads the addresses of the IRP, PCB, UCB, CCB, and status into R3, R4, R5, R6, and R0, respectively; **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver error-handling callback routine that uses this macro can access any of its arguments by using a symbolic name of the form `ERRARG$_argument-name`.

Called by

`EXE_STD$MODIFYLOCK`, `EXE_STD$READLOCK`, and `EXE_STD$WRITELOCK` call the driver's error-handling callback routine to process errors incurred by a buffer accessibility check or buffer locking operation.

Context

An error-handling callback routine is called at `IPL$_ASTDEL` and must exit at `IPL$_ASTDEL`. An error-handling callback routine must not lower IPL below `IPL$_ASTDEL`. If it raises IPL, it must lower it to `IPL$_ASTDEL` before passing control to any other code. Similarly, before exiting, it must release any spin locks it may have acquired in an OpenVMS multiprocessing environment.

Error-handling callback routines execute in the context of the process that requested the I/O activity. If a routine alters the stack, it must restore the stack before returning control to the caller of the routine.

Exit mechanism

An error-handling callback routine must return control to its caller and preserve the contents of R0 and R1.

Description

An error-handling callback routine processes any errors incurred by a call to `EXE_STD$MODIFYLOCK`, `EXE_STD$READLOCK`, or `EXE_STD$WRITELOCK`.

A driver typically requires an error-handling callback routine if 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. The routine performs such operations as locating the addresses of the previously allocated buffers (typically stored in the IRP) and calling `MMG_STD$UNLOCK` to release them.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate the entry point of each FDT error-handling callback routine with a `$DRIVER_ERRRTN_ENTRY` macro, indicating which registers must be saved and restored across routine execution.
- You should examine an FDT routine's use of R0, R7, and R8.

The FDT routine of an OpenVMS VAX device driver may obtain the address of FDT routine being called from R0, the number of the bit that specifies the code for the requested I/O function from R7, and the address of the entry in the function decision table that dispatched control to this FDT routine.

OpenVMS AXP Device Driver Entry Points FDT Error-Handling Callback Routine

An OpenVMS AXP driver can obtain the user-supplied function code from `IRP$FUNC`. It can obtain the address of the start of the FDT from `DDT$PS_FDT2`. The DDT address is available from `UCB$DDT`.

- An FDT routine of an OpenVMS VAX device driver accesses values of the function-dependent arguments specified in the `$QIO` request as offsets from the value of the AP; an OpenVMS AXP device driver obtains them from the IRP (at symbolic offsets `IRP$QIO_P1` through `IRP$QIO_P6`).

OpenVMS AXP Device Driver Entry Points

Interrupt Service Routine

Interrupt Service Routine

Processes interrupts generated by a device. The Interrupt Service routine is called by the system interrupt dispatcher.

Format

DRIVER_INTERRUPT (idb, scb_offset)

Arguments

Argument	Type	Access	Mechanism
idb	IDB	input	reference
scb_offset	integer	input	value

idb

Interrupt dispatch block.

Essentials

Identifying the Routine

Devices require an interrupt service routine for each interrupt vector. Use the `DPT_STORE_ISR` macro to store the ISR procedure descriptor and entry point address in the interrupt transfer vector block (VEC) at `CRB$L_INTD`. You can find the second and third VECs at `CRB$L_INTD2` and `CRB$L_INTD+2*VEC$K_LENGTH`, respectively.

Declaring the Entry Point

Indicate the entry point of an OpenVMS AXP interrupt service routine with a `.CALL_ENTRY` MACRO-32 compiler directive to indicate which registers are provided as input or used as output and which must be saved and restored. If the interrupt service routine forks, transferring control to a fork routine, it must declare, at its `.CALL_ENTRY` point, R3, R4, and R5 as **input** registers.

Called by

The interrupt service routine is called either by the OpenVMS interrupt dispatcher (for direct-vectorized adapters) or by an adapter interrupt service routine (for non-direct-vector adapters).

Context

An OpenVMS AXP driver's interrupt service routine conforms to the OpenVMS calling standard.

An interrupt service routine is called, executes, and returns at device IPL. It 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.

At the execution of a driver's interrupt service routine, the processor is running in interrupt mode on the kernel stack. As a result, an interrupt service routine can reference only those virtual addresses that reside in system (S0) space.

OpenVMS AXP Device Driver Entry Points Interrupt Service Routine

Resuming the Suspended Driver Thread

The method that an interrupt service routine should use to invoke the driver's resume from interrupt routine depends on how the driver suspended its execution.

If the driver is using the simple fork mechanism with a JSB-based environment then the driver resume from interrupt routine is invoked by the following:

```
MOVX   UCB$Q_FR3(R5),R3   ;R3 = FR3 (64-bits)
MOVX   UCB$Q_FR4(R5),R4   ;R4 = FR4 (64-bits)
JSB    @UCB$L_FPC(R5)
```

If the driver is using the simple fork mechanism with a CALL-based environment then the driver resume from interrupt routine is invoked in C by the following:

```
(ucb->ucb$l_fpc)( ucb->ucb$q_fr3, ucb->ucb$q_fr4, ucb);
```

or in MACRO-32 by the following:

```
PUSHL  R5                ;Param3 = UCB address
PUSHL  UCB$Q_FR4(R5)     ;Param2 = FR4 value
PUSHL  UCB$Q_FR3(R5)     ;Param1 = FR3 value
CALLS  #3,@UCB$L_FPC(R5)
```

If the driver is using the kernel process mechanism then the suspended kernel process can be resumed in C by the following:

```
exe$kp_restart( kpb );
```

or:

```
(ucb->ucb$l_fpc)( ucb->ucb$q_fr3, ucb->ucb$q_fr4, ucb);
```

or in MACRO-32 by the following:

```
PUSHL  UCB$Q_FR4(R5)     ;Param1 = KPB address
CALLS  #1,EXE$KP_RESTART
```

or:

```
PUSHL  R5                ;Param3 = UCB address
PUSHL  UCB$Q_FR4(R5)     ;Param2 = FR4 value
PUSHL  UCB$Q_FR3(R5)     ;Param1 = FR3 value
CALLS  #3,@UCB$L_FPC(R5)
```

Exit mechanism

The interrupt service routine returns to the interrupt dispatcher with a RET 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.

OpenVMS AXP Device Driver Entry Points Interrupt Service Routine

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- An OpenVMS VAX device driver declares an interrupt service routine by issuing the `DPT_STORE` macro to store its address in an interrupt transfer vector block. Because the OpenVMS AXP interrupt dispatcher requires the addresses of both the code entry point and the procedure descriptor of an interrupt service routine, you must use the new `DPT_STORE_ISR` macro (which generates both) to declare the routine.
- The OpenVMS VAX interrupt dispatcher issues a `JSB` instruction to pass control to an OpenVMS VAX driver's interrupt service routine; the OpenVMS AXP interrupt dispatcher issues a standard call to a driver's interrupt service routine. This results in some substantial differences:
 - You must indicate the entry point of an OpenVMS AXP interrupt service routine with a `.CALL_ENTRY MACRO-32` compiler directive to indicate which registers are provided as input or used as output and which must be saved and restored.
 - An OpenVMS VAX driver's interrupt service routine must preserve any of the non-scratch registers R2 through R15 if it uses them.
 - An OpenVMS VAX driver's interrupt service routine is passed various information on the stack, including the address of the IDB, the contents of R0 through R5, the PC, and PSL at the time of the interrupt.

The only parameter passed to an OpenVMS AXP driver's interrupt service routine is the address of the IDB (that is, the contents of `VECSL_IDB`). The routine cannot reference data on the stack.
 - Before exiting, an OpenVMS VAX driver's interrupt service routine removes the address of the pointer to the IDB from the top of the stack and restores the registers OpenVMS saved when dispatching the interrupt.

An OpenVMS AXP driver's interrupt service routine does not perform these actions.
 - An OpenVMS VAX driver's interrupt service routine exits with an `REI` instruction.

An OpenVMS AXP driver's interrupt service routine exits by returning control with a `RET` instruction.
- You must replace direct CSR access (for instance, by means of a `MOVL` instruction) with CSR access by means of a `CRAM`.
- If you alter the driver's suspension mechanism such that it uses the OpenVMS kernel process services, you must change the mechanism by which the interrupt service routine reactivates lower IPL execution threads by replacing the `IOFORK` macro with the `KP_STALL_IOFORK` macro.

Mount Verification Routine

Performs device-specific mount verification.

Format

MNTVER (irp, ucb)

Arguments

Argument	Type	Access	Mechanism
irp	IRP	input	reference
ucb	UCB	input	reference

irp

I/O request packet, or zero to complete mount verification.

ucb

Unit control block.

Essentials

Identifying the Routine

Supply the address of the mount verification routine in the **mntver** argument of the DDTAB macro. The macro places the procedure value of this routine into DDT. The default value of this argument, IOC_STDSMNTVER, is the only value allowed for device drivers not supplied by Digital.

Declaring the Entry Point

Use:

```
$DRIVER_MNTVER_ENTRY [preserve=<>] [,fetch=YES]
```

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO-32 compiler) across the call to the cancel selective routine.

fetch=YES, the default, loads the addresses of IRP and UCB into R3 and R5, respectively, **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver mount verification routine that uses this macro can access any of its arguments by using a symbolic name of the form **MNTARGS_argument-name**.

Called by

Routine DRIVER_CODE in module MOUNTVER calls a driver's mount verification routine.

Context

A mount verification routine is called at fork IPL with the corresponding fork lock held in a multiprocessing system.

OpenVMS AXP Device Driver Entry Points Mount Verification Routine

Exit mechanism

The mount verification routine returns to its caller.

Description

Reserved to Digital.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note that you must indicate the entry point of a mount verification routine with a `$DRIVER_MNTVER_ENTRY` indicating which registers must be saved and restored across routine execution.

Register Dumping Routine

Copies the contents of a device's registers to an error message buffer or a diagnostic buffer.

Format

status = REGDMP (buffer, arg_2, ucb)

Arguments

Argument	Type	Access	Mechanism
buffer	address	input	reference
arg_2	unspecified	input	reference
ucb	UCB	input	reference

buffer

Address of buffer into which a register dumping routine copies the contents of device registers.

arg_2

Device-specific argument, usually a controller register access mailbox (CRAM).

ucb

Unit control block.

Essentials

Identifying the Routine

Specify the name of the register dumping routine in the **regdmp** argument of the DDTAB macro. This macro places the procedure value of the routine into the DDT.

Declaring the Entry Point

Use:

```
$DRIVER_REGDUMP_ENTRY [preserve=<R2>] [,fetch=YES]
```

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO compiler) across the call to the register dumping routine.

fetch=YES, the default, loads the addresses of the buffer, the driver-specific argument, and the UCB into R0, R4, and R5, respectively; **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver register dumping routine that uses this macro can access any of its arguments by using a symbolic name of the form REGARG\$_**argument-name**.

OpenVMS AXP Device Driver Entry Points

Register Dumping Routine

Called by

The system error-logging routines (ERL_STD\$DEVICERR, ERL_STD\$DEVICTMO, and ERL_STD\$DEVICEATTN) and diagnostic buffer filling routine (IOC_STD\$DIAGBUFILL) call the register dumping routine.

Context

OpenVMS calls a register dumping routine at the same interrupt service routine (IPL) at which the driver called the OpenVMS AXP system routine ERL_STD\$DEVICERR, ERL_STD\$DEVICTMO, ERL_STD\$DEVICEATTN, or IOC_STD\$DIAGBUFILL. A register dumping routine must not change IPL.

A register dumping routine executes within the context of an IPL 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. 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.

Exit mechanism

The register dumping routine returns to its caller.

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

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate the entry point of the routine with a \$DRIVER_REGDUMP_ENTRY macro, indicating which registers must be saved and restored across routine execution.
- An OpenVMS VAX device driver's register dumping routine is passed the address of the device's CSR in R4 (if the driver invoked the WFIKPCH macro to wait for an interrupt or timeout).

An OpenVMS AXP device driver's register dumping routine is not passed the address of the CSR. It may access the controller's register by means of the controller register access mailbox (CRAM), the address of which is usually passed in **arg_2**.

- You must replace direct CSR access (for instance, by means of a MOVL instruction) with CSR access by means of a CRAM.

Start-I/O Routine (Simple Fork, Call Environment)

Activates a device to process a requested I/O function.

Format

START (irp, ucb)

Arguments

Argument	Type	Access	Mechanism
irp	IRP	input	reference
ucb	UCB	input	reference

irp
I/O request packet.

ucb
Unit control block. The start-I/O routine uses information from the following UCB fields to calculate the size and location of a transfer:

Field	Description
UCB\$\$_BCNT	Number of bytes to be transferred, copied from the low-order word of IRP\$\$_BCNT
UCB\$\$_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\$\$_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

Essentials

Identifying the Routine

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 the DDT.

Declaring the Entry Point

Use:

```
$DRIVER_START_ENTRY [preserve=<R2,R4>] [,fetch=YES]
```

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO compiler) across the call to the start-I/O routine

OpenVMS AXP Device Driver Entry Points

Start-I/O Routine (Simple Fork, Call Environment)

fetch=YES, the default, loads the addresses of the IRP and UCB into R3 and R5, respectively; **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver start-I/O routine that uses this macro can access any of its arguments by using a symbolic name of the form **STARTARG\$_argument-name**.

Called by

A traditional start-I/O routine is called as the result of a standard call issued by **IOC_STDS\$INITIATE** and **IOC_STDS\$REQCOM** in module **IOSUBNPAG**.

Context

A start-I/O routine is placed into execution at fork IPL, holding the associated fork lock. 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. An invocation of the **WFIKPCH** or **WFIRLCH** macro (or **KP_STALL_WFIKPCH** or **KP_STALL_WFIRLCH**) to wait for a device interrupt releases this device lock.

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. 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.

Exit mechanism

A traditional start-I/O routine suspends itself whenever it must wait for a required resource, such as a controller data channel. To do so, it invokes an OpenVMS macro (such as **REQPCHAN**) 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 **WFIKPCH** 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 an **IOFORK** macro, the routine returns 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 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, which is the OpenVMS fork dispatcher.

OpenVMS AXP Device Driver Entry Points Start-I/O Routine (Simple Fork, Call Environment)

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.

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.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate the entry point of the start-I/O routine with a \$DRIVER_START_ENTRY macro, indicating which registers must be saved and restored across routine execution.
- You must replace direct CSR access (for instance, by means of a MOVL instruction) with CSR access by means of a CRAM.
- You should examine the routine's use of suspension mechanisms (for instance, its forking, wait-for-interrupt, and resource-wait semantics) to determine whether it needs to be adapted to use the kernel process services. Typically a driver that makes subroutine calls before suspending itself (and relies on the previous context of these subroutines remaining intact on the stack), must be adapted to use the kernel process services.

OpenVMS AXP Device Driver Entry Points

Start-I/O Routine (Kernel Process)

Start-I/O Routine (Kernel Process)

Activates a device to process a requested I/O function.

Format

START (kpb)

Arguments

Argument	Type	Access	Mechanism
kpb	KPB	input	reference

kpb
Kernel process block.

Essentials

Identifying the Routine

Specify the name of the kernel process start-I/O routine (EXE_STD\$KP_STARTIO) in the **start** argument of the DDTAB macro, and the name of the driver's start-I/O routine in the **kp_startio** argument.

Declaring the Entry Point

Indicate the entry point of a kernel process start-I/O routine with a .CALL_ENTRY MACRO-32 compiler directive to indicate which registers are provided as input or used as output and which registers must be saved and restored.

Called by

A kernel-process start-I/O routine is called by EXE_STD\$KP_STARTIO in module KERNEL_PROCESS.

Context

A kernel process start-I/O routine is placed into execution at fork IPL, holding the associated fork lock. The kernel process start-I/O routine 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. An invocation of the KP_STALL_WFIKPCH or KP_STALL_WFIRLCH macro to wait for a device interrupt releases this device lock.

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.

Neither the start-I/O routine that initiates a kernel process nor the kernel process thread can depend on inheriting the synchronization capabilities (such as spin locks and IPL) of the other when control is exchanged between them. If they must share data or perform other operations that require synchronization, they must explicitly establish a synchronization mechanism.

OpenVMS AXP Device Driver Entry Points Start-I/O Routine (Kernel Process)

The kernel process cannot assume that its initiator is not running in parallel, nor can the initiator of the kernel process assume that the kernel process has already executed when EXESKP_START returns control.

Exit mechanism

A kernel process start-I/O routine suspends itself whenever it must wait for a required resource, such as a controller data channel. To do so, the kernel process start-I/O routine invokes an OpenVMS macro (such as KP_STALL_REQCHAN) 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 KP_STALL_WFIKPCH or KP_STALL_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 KP_STALL_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 a KP_STALL_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 KP_REQCOM macro. In addition to initiating device-independent postprocessing of the current request, the KP_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 kernel process start-I/O routine, EXESKP_STARTIO.

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.

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.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- If the routine need not be converted to a kernel process, you must replace any calls to EXESFORK, EXESFORK_WAIT, EXESIOFORK, IOC\$WFIKPCH, IOC\$WFIRLCH, IOC\$REQCHANH, and IOC\$REQCHANL with invocations of the appropriate suspension macro or with calls to EXE_STD\$PRIMITIVE_FORK, IOC_STD\$PRIMITIVE_WFIKPCH, IOC_STD\$PRIMITIVE_WFIRLCH, IOC_STD\$PRIMITIVE_REQCHANH, or IOC_STD\$PRIMITIVE_REQCHANL.
- You must indicate the entry point of a kernel process start-I/O routine with a .CALL_ENTRY MACRO-32 compiler directive to indicate which registers are provided as input or used as output and which registers must be saved and restored. A kernel process start-I/O routine invokes the KP_REQCOM macro (in place of the REQCOM macro) to return control properly to its caller.

OpenVMS AXP Device Driver Entry Points Timeout Handling Code (Traditional)

Timeout Handling Code (Traditional)

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.

Format

BNEQ timeout-code-address

Input

Location	Contents
R3	Contents of R3 when the last invocation of WFIKPCH or WFIRLCH occurred (usually the address of the IRP)
R4	Contents of R4 when the last invocation of WFIKPCH or WFIRLCH occurred (usually the address of the IDB)
R5	Address of UCB of the device
UCB\$SL_STS	UCB\$V_INT and UCB\$V_TIM clear; UCB\$V_TIMOUT set

Essentials

Identifying the Timeout Handler

Specify the address of timeout code in the **except** argument to the WFIKPCH or WFIRLCH macro.

Branched to

The WFIKPCH and WFIRLCH macros use this entry point, but only when the label of timeout code 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.

The OpenVMS AXP software timer interrupt service routine restarts a stalled driver fork procedure, passing to it a status (UCB\$V_TIMOUT in UCB\$SL_STS) which is inspected by one of two instructions left at the top of the fork procedure by the WFIKPCH or WFIRLCH macro. If UCB\$V_TIMOUT is set, the second instruction branches to the timeout code.

Context

Timeout code receives control at device IPL and must exit at device IPL. At the time the timeout code executes, the processor holds both the fork lock and the device lock associated with the device.

After taking whatever device-specific action is necessary at device IPL, timeout code can lower IPL to fork IPL to perform less critical activities. Because its caller restores IPL to fork IPL (and releases the device lock), if a timeout handler lowers IPL, it can do so only by forking or by performing the following steps:

1. Issue a DEVICEUNLOCK macro to lower to fork level
2. Perform timeout handling activities possible at the lower IPL

OpenVMS AXP Device Driver Entry Points Timeout Handling Code (Traditional)

3. Issue a DEVICELOCK macro to again obtain the device lock and raise to device IPL

Timeout code can access only those virtual addresses that refer to system (S0) space.

Traditional timeout code can use R0, R1, and R2 freely, but must preserve the contents of all other registers. If it 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.

Exit mechanism

Traditional timeout code issues an RSB instruction to return to the software timer interrupt service routine, restarts the I/O request, or invokes the REQCOM macro to complete the I/O request that encountered the timeout.

Description

There are no outputs required from timeout code but, depending on the characteristics of the device, timeout code might cancel or retry the current I/O request, send a message to the operator, or take some other action.

Before timeout code executes, the system has placed the device in a state in which no interrupt is expected (by clearing the bit UCB\$V_INT in field UCB\$SL_STS). If the requested interrupt occurs while this routine executes, it will appear to be an unsolicited interrupt. Many drivers handle this situation by disabling interrupts while timeout code executes.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- On OpenVMS VAX systems, the software timer interrupt service routine issues a JSB instruction to a timeout handling routine within a driver when it detects that a device has timed out. On OpenVMS AXP systems, the OpenVMS AXP suspension macros provide a mechanism by which the driver fork routine, when resumed by a timeout, tests the timeout bit in the UCB and branches, if the bit is set, to the address of the timeout code.
- You must replace direct control and status register (CSR) access (for instance, by means of a MOVL instruction) with CSR access using one of the OpenVMS AXP CSR access methods (CRAMs, platform independent access routines, or direct mapping).

Timeout Handling Code (Kernel Process)

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.

Format

BLBC timeout-code-address

Arguments

None.

Essentials

Identifying the Routine

Specify the address of the timeout code in the **except** argument to the `KP_STALL_WFIKPCH` or `KP_STALL_WFIRLCH` macro.

Branched to

The `KP_STALL_WFIKPCH`, and `KP_STALL_WFIRLCH` macros use this entry point, but only when the label of timeout code 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.

The OpenVMS AXP software timer interrupt service routine restarts a stalled driver kernel process fork procedure, passing a status (`UCBSV_TIMEOUT` in `UCBSL_STS`) to it, which is inspected by one of two instructions left at the top of the fork procedure by the `KP_STALL_WFIKPCH` or `KP_STALL_WFIRLCH` macro. If `UCBSV_TIMEOUT` is set, the second instruction branches to the timeout code.

Context

The timeout code receives control at device IPL and must exit at device IPL. At the time the timeout code executes, the processor holds both the fork lock and device lock associated with the device.

After taking whatever device-specific action is necessary at device IPL, timeout code can lower IPL to fork IPL to perform less critical activities. Because its caller restores IPL to fork IPL (and releases the device lock), if timeout code lowers IPL, it can do so only by forking or by performing the following steps:

1. Issue a `DEVICEUNLOCK` macro to lower to fork level
2. Perform timeout handling activities possible at the lower IPL
3. Issue a `DEVICELock` macro to again obtain the device lock and raise to device IPL

Timeout code can access only those virtual addresses that refer to system (S0) space.

Kernel process timeout code executes in the context of the kernel process thread that invoked the `KP_STALL_WFIKPCH` or `KP_STALL_WFIRLCH` macro.

OpenVMS AXP Device Driver Entry Points Timeout Handling Code (Kernel Process)

Exit mechanism

Kernel process timeout code executes as part of the kernel process thread that invoked WFIKPCH or WFIRLCH macro and therefore has no special exit mechanism.

Description

There are no outputs required from timeout code but, depending on the characteristics of the device, timeout code might cancel or retry the current I/O request, send a message to the operator, or take some other action.

Before timeout code executes, OpenVMS has placed 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 while this routine executes, it will appear to be an unsolicited interrupt. Many drivers handle this situation by disabling interrupts while timeout code executes.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- On OpenVMS VAX systems, the software timer interrupt service routine issues a JSB instruction to a timeout handling routine within a driver when it detects that a device has timed out. On OpenVMS AXP systems, the OpenVMS AXP suspension macros provide a mechanism by which the driver fork routine, when resumed by a timeout, tests the timeout bit in the UCB and branches, if the bit is set, to the address of the timeout code.
- You must replace direct CSR access (for instance, by means of a MOVL instruction) with CSR access using one of the OpenVMS AXP CSR access methods (CRAMs, platform independent access routines, or direct mapping).

OpenVMS AXP Device Driver Entry Points

Unit Delivery Routine

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.

Format

status = DELIVER (ddb, idb, unit_number, scratch_area, adp)

Arguments

Argument	Type	Access	Mechanism
ddb	DDB	input	reference
idb	IDB	input	reference
unit_number	integer	input	value
scratch_area	address	input	reference
adp	ADP	input	reference

ddb

Device data block.

idb

Interrupt dispatch block; 0 if none exists.

unit_number

Number of unit that the unit delivery routine must decide to configure or not to configure.

scratch_area

Address of quadword scratch area.

adp

Adapter control block.

Essentials

Identifying the Routine

Specify the name of the unit delivery routine in the **deliver** argument to the DPTAB macro. The macro puts the procedure value address of this routine in the DPT.

Declaring the Entry Point

Use:

```
$DRIVER_DELIVER_ENTRY [preserve=<R2>] [,fetch=YES]
```

OpenVMS AXP Device Driver Entry Points Unit Delivery Routine

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO compiler) across the call to the unit delivery routine.

fetch=YES, the default, loads the address of the IDB into R3 and R4, the unit number into R5, the address of the scratch area into R7, and the address of the ADP into R8; **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver unit delivery routine that uses this macro can access any of its arguments by using a symbolic name of the form **DLVRARG\$_argument-name**.

Called by

The System Management (SYSMAN) utility's IO 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.

Context

The unit delivery routine is called at IPL\$_POWER. It must not lower IPL. The unit delivery routine executes in the context of the process within which the autoconfiguration facility executes.

Exit mechanism

A unit delivery routine returns success or failure status to the autoconfiguration facility. If the routine returns error status, the unit is not configured.

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.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate the entry point of the routine with a \$DRIVER_DELIVER_ENTRY macro to indicate which registers must be saved and restored across routine execution.
- You must replace direct CSR access (for instance, by means of a MOVL instruction) with CSR access using one of the OpenVMS AXP CSR access methods (CRAMs, platform independent access routines, or direct mapping).
- The unit delivery routine of an OpenVMS VAX device driver receives the addresses of the device CSR in R4 and the IDB in R5. An OpenVMS AXP device driver's unit delivery routine is not passed the address of the CSR. It may access the controller's register by means of the controller register access mailbox (CRAM), the address of which is provided in IDB\$PS_CRAM.
- An OpenVMS AXP unit delivery routine is passed the address of the device data block and the address of a quadword scratch area.

OpenVMS AXP 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.

Format

status = UNITINIT (idb, ucb)

Arguments

Argument	Type	Access	Mechanism
idb	IDB	input	reference
ucb	UCB	input	reference

idb

Interrupt dispatch block associated with the controller.

ucb

Unit control block.

Essentials

Identifying the Routine

Specify the address of the unit initialization routine **unitinit** argument of the DDTAB macro. This macro places the procedure value of the routine into the DDT.

Declaring the Entry Point

Use:

```
$DRIVER_UNITINIT_ENTRY [preserve=<R2>] [,fetch=YES]
```

where:

preserve indicates the registers to be preserved (in addition to those automatically preserved by the MACRO-32 compiler) across the call to the unit initialization routine.

fetch=YES, the default, loads \$\$\$_NORMAL status into R0, and the addresses of the IDB and UCB into R4 and R5, respectively; **fetch=NO** disables register loading. Regardless of the value of the **fetch** argument, a driver unit initialization routine that uses this macro can access any of its arguments by using a symbolic name of the form UNITARG\$_**argument-name**.

Called by

The driver-loading procedure calls a driver's unit initialization routine when processing a CONNECT command. OpenVMS 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.

OpenVMS AXP Device Driver Entry Points Unit Initialization Routine

Context

OpenVMS 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 the driver-loading procedure calls the unit initialization routine immediately after the controller initialization returns control to it, the driver's initialization routines must synchronize their activities.

The portion of the unit initialization routine that services power failure cannot acquire any spin locks. As a result, the routine cannot fork to perform power failure servicing.

Because OpenVMS calls it in system context, a unit initialization routine can only refer to those virtual addresses that reside in system (S0) space. R0, and preserve the contents of all registers except R0, R1, and R2.

Exit mechanism

A unit initialization routine returns success or failure status to its caller.

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\$_POWER in UCB\$_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\$_ONLINE in UCB\$_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\$_OWNER.
7. For dedicated controllers, initializes controller and device hardware.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate the entry point of the routine with a `.JSB_ENTRY` MACRO-32 compiler directive to indicate which registers are provided as input or used as output and which must be saved and restored.
- An OpenVMS VAX device driver can specify a controller initialization routine by invoking the `DPT_STORE` macro to place its address into the interrupt transfer vector block (`CRB$_INTD+VEC$_UNITINIT`). An OpenVMS AXP device driver specifies the routine in the **unitinit** argument of the `DDTAB` macro.
- You must replace direct CSR access (for instance, by means of a `MOVL` instruction) with CSR access using one of the OpenVMS AXP CSR access methods (CRAMs, platform independent access routines, or direct mapping).

OpenVMS AXP Device Driver Entry Points Unit Initialization Routine

- The unit initialization routine of an OpenVMS VAX device driver receives the addresses of the primary and secondary device CSRs in R3 and R4, respectively. An OpenVMS AXP device driver's unit initialization routine is not passed the addresses of the CSRs. It may access the controller registers by means of the controller register access mailbox (CRAM), the address of which is provided in IDB\$PS_CRAM.
- An OpenVMS AXP unit initialization routine must return success or failure status to its caller.

System Routines

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

- Most routines reside in modules within the [SYS] facility of the operating system. 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, the operating system 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 4 describes how a driver can use these macros.
- System routines generally return a status value in R0 (for instance, SSS_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 *OpenVMS System Services Reference Manual* and the *OpenVMS System Messages and Recovery Procedures Reference Manual*.

Table 2-1 highlights some of the differences between OpenVMS VAX and OpenVMS AXP system routines.

Table 2-1 New, Changed, and Unsupported OpenVMS System Routines

System Routine	Description	Notes
EXESBUS_DELAY	Allows a system-specific bus delay within a timed wait	New
EXESDELAY	Provides a short-term simple delay	New
ERL\$DEVICERR, ERL\$DEVICTMO, ERL\$DEVICEATTN	Allocate an error message buffer and record in it information concerning the error	Changed
EXESFORK	Creates a fork process on the current processor	Replaced by EXESPRIMITIVE_ FORK and EXE_ STD\$PRIMITIVE_FORK
EXESFORK_WAIT	Inserts a fork block on the fork-and-wait queue	Replaced by EXESPRIMITIVE_ FORK_WAIT and EXE_ STD\$PRIMITIVE_FORK_ WAIT
EXESINSERT_IRP	Inserts an IRP into the specified queue of IRPs according to the base priority of the process that issued the I/O request	New

(continued on next page)

System Routines

Table 2–1 (Cont.) New, Changed, and Unsupported OpenVMS System Routines

System Routine	Description	Notes
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	Replaced by EXE\$INSERT_IRP
EXE\$IOFORK	Creates a fork process on the current processor for a device driver, disabling timeouts from the associated device	Replaced by EXE\$PRIMITIVE_FORK and EXE_\$STD\$PRIMITIVE_FORK
EXE\$KP_ALLOCATE_KPB	Creates a KPB and a kernel process stack, as required by the kernel process services	New
EXE\$KP_DEALLOCATE_KPB	Deallocates a KPB and its associated kernel process stack	New
EXE\$KP_END	Terminates the execution of a kernel process	New
EXE\$KP_FORK	Stalls a kernel process in such a manner that it can be resumed by the fork dispatcher	New
EXE\$KP_FORK_WAIT	Stalls a kernel process in such a manner that it can be resumed by the software timer interrupt service routine's examination of the fork-and-wait queue	New
EXE\$KP_RESTART	Resumes the execution of a kernel process	New
EXE\$KP_STALL_GENERAL	Stalls the execution of a kernel process	New
EXE\$KP_START	Starts the execution of a kernel process	New
EXE_\$STD\$KP_STARTIO	Sets up and starts a kernel process to be used by a device driver	New
EXE\$MODIFYLOCK	Validate and prepare a user buffer for a direct-I/O, DMA read/write operation.	Replaced by EXE_\$STD\$MODIFYLOCK and CALL_MODIFYLOCK macro
EXE\$MODIFYLOCKR	Validates and prepares a user buffer for a direct-I/O, DMA modify operation.	Replaced by EXE_\$STD\$MODIFYLOCK and CALL_MODIFYLOCK_ERR macro
EXE\$PRIMITIVE_FORK, EXE_\$STD\$PRIMITIVE_FORK	Creates a simple fork process on the current processor	New
EXE\$PRIMITIVE_FORK_WAIT, EXE_\$STD\$PRIMITIVE_FORK_WAIT	Inserts a fork block on the fork-and-wait queue	New
EXE\$READLOCK	Validate and prepare a user buffer for a direct-I/O, DMA read operation.	Replaced by EXE_\$STD\$READLOCK and CALL_READLOCK macro
EXE\$READLOCKR	Validates and prepares a user buffer for a direct-I/O, DMA read operation	Replaced by EXE_\$STD\$READLOCK and CALL_READLOCK_ERR macro

(continued on next page)

Table 2–1 (Cont.) New, Changed, and Unsupported OpenVMS System Routines

System Routine	Description	Notes
EXESTIMEDWAIT_COMPLETE	Determines whether the time interval of a timed wait has conclude	New
EXESTIMEDWAIT_SETUP, EXESTIMEDWAIT_SETUP_ 10US	Calculate and return the end-value used by EXESTIMEDWAIT_COMPLETE to determine when a timed wait has completed	New
EXESWRITELOCK	Validate and prepare a user buffer for a direct-I/O, DMA write operation.	Replaced by EXE_ STD\$WRITELOCK and CALL_WRITELOCK macro
EXESWRITELOCKR	Validates and prepares a user buffer for a direct-I/O, DMA write operation	Replaced by EXE_ STD\$WRITELOCK and CALL_WRITELOCK_ERR macro
IOCSALOALTMAP, IOCSALOALTMAPN, IOCSALOALTMAPSP	Allocate a set of Q22-bus alternate map registers	Not supported. See the description of IOCSALLOC_CNT_RES.
IOCSALOUBAMAP, IOCSALOUBAMAPN	Allocate a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers	Not supported. See the description of IOCSALLOC_CNT_RES.
IOCSALLOC_CNT_RES	Allocates the requested number of items of a counted resource	New
IOCSALLOC_CRAB	Allocates and initializes a counted resource allocation block (CRAB)	New
IOCSALLOC_CRCTX	Allocates and initializes a counted resource context block (CRCTX)	New
IOCSALLOCATE_CRAM	Allocates a controller register access mailbox	New
IOCSCANCEL_CNT_RES	Cancels a thread that has been stalled waiting for a counted resource	New
IOCSGRAM_CMD	Generates values for the command, mask, and remote I/O interconnect address fields of the hardware I/O mailbox that are specific to the interconnect that is the target of the mailbox operation, inserting these values into the indicated mailbox, buffer, or both	New
IOCSGRAM_IO	Queues the hardware I/O mailbox defined within a controller register access mailbox (GRAM) to the mailbox pointer register (MBPR) and awaits the completion of the mailbox transaction	New
IOCSGRAM_QUEUE	Queues the hardware I/O mailbox defined within a controller register access mailbox (GRAM) to the mailbox pointer register (MBPR)	New
IOCSGRAM_WAIT	Awaits the completion of a hardware I/O mailbox transaction to a tightly coupled I/O interconnect	New

(continued on next page)

System Routines

Table 2–1 (Cont.) New, Changed, and Unsupported OpenVMS System Routines

System Routine	Description	Notes
IOC\$DEALLOC_CNT_RES	Deallocates the requested number of items of a counted resource	New
IOC\$DEALLOC_CRAB	Deallocates a counted resource allocation block (CRAB)	New
IOC\$DEALLOC_CRCTX	Deallocates a counted resource context block (CRCTX)	New
IOC\$DEALLOCATE_CRAM	Deallocates a controller register access mailbox	New
IOC\$DIAGBUFILL	Fills a diagnostic buffer if the original \$QIO request specified such a buffer	Changed
IOC\$SKP_REQCHAN	Stalls a kernel process in such a manner that it can be resumed by the granting of a device controller channel	New
IOC\$SKP_WFIKPCH, IOC\$SKP_WFIRLCH	Stall a kernel process in such a manner that it can be resumed by device interrupt processing	New
IOC\$LOAD_MAP	Loads a set of adapter-specific map registers	New
IOC\$LOADALTMAP	Loads a set of alternate Q22-bus map registers	Not supported; see IOC\$LOAD_MAP
IOC\$LOADMBAMAP	Loads MASSBUS map registers	Not supported; see IOC\$LOAD_MAP
IOC\$LOADUBAMAP, IOC\$LOADUBAMAPA	Load a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers	Not supported; see IOC\$LOAD_MAP
IOC\$MAP_IO	Maps I/O bus physical address space into an address region accessible by the processor	New
IOC\$NODE_FUNCTION	Performs node-specific functions on behalf of a driver, such as enabling or disabling interrupts from a bus slot	New
IOC_STD\$PRIMITIVE_REQCHANH, IOC_STD\$PRIMITIVE_REQCHANL	Request a controller's data channel and, if unavailable, place process in channel wait queue	New
IOC_STD\$PRIMITIVE_WFIKPCH, IOC_STD\$PRIMITIVE_WFIRLCH	Suspend a driver fork thread and fold its context into a fork block in anticipation of a device interrupt or timeout	New
IOC\$READ_IO	Reads a value from a previously mapped location in I/O address space	New
IOC\$RELALTMAP	Releases a set of Q22-bus alternate map registers	Not supported; see IOC\$DEALLOC_CNT_RES
IOC\$RELDATAP	Releases a UNIBUS adapter's buffered data path.	Not supported
IOC\$RELMAPREG	Releases a set of UNIBUS map registers or a set of the first 496 Q22-bus map registers	Not supported; see IOC\$DEALLOC_CNT_RES

(continued on next page)

Table 2–1 (Cont.) New, Changed, and Unsupported OpenVMS System Routines

System Routine	Description	Notes
IOCSREQALTMAP	Allocates sufficient Q22-bus alternate map registers to accommodate a DMA transfer	Not supported; see IOCSALLOC_CNT_RES
IOCSREQDATAP, IOCSREQDATAPNW	Request a UNIBUS adapter's buffered data path and, optionally, if no path is available, place process in a data-path wait queue	Not supported
IOCSREQMAPREG	Allocates sufficient UNIBUS map registers or a sufficient number of the first 496 Q22-bus map registers to accommodate a DMA transfer	Not supported; see IOCSALLOC_CNT_RES
IOCSREQPCHANH, IOCSREQPCHANL, IOCSREQSCHANH, IOCSREQSCHANL	Request a controller's primary or secondary data channel and, if unavailable, place process in channel wait queue	Not supported
IOCSWFIKPCH, IOCSWFIRLCH	Suspend a driver fork thread and fold its context into a fork block in anticipation of a device interrupt or timeout	Replaced by IOC_STD\$PRIMITIVE_WFIKPCH and IOC_STD\$PRIMITIVE_WFIRLCH
IOCSWRITE_IO	Writes a value to a previously mapped location in I/O address space	New
IOCSUNMAP_IO	Unmaps a previously mapped I/O address space	New

System Routines

ACP_STD\$ACCESS

ACP_STD\$ACCESS

Accesses and creates ACP function processing.

Module

SYSACPFDT

Format

status = ACP_STD\$ACCESS (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp
I/O request packet.

pcb
Process control block of the current process.

ucb
Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb
Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$ACCVIO Access violation.
SS\$DEVNOTMOUNT Device not mounted.
SS\$DEVFOREIGN Device is mounted as foreign.
SS\$EXQUOTA File quota exceeded.
SS\$FILALRACC File already accessed.

SS\$_IVCHNLSEC	Invalid section channel.
SS\$_NORMAL	The I/O request has been successfully queued to the appropriate ACP or XQP.

Context

FDT dispatching code in the \$QIO system service calls ACP_STD\$ACCESS as an upper-level FDT action routine at IPL\$_ASTDEL.

Description

For Digital internal use only.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine ACP\$ACCESS (used by OpenVMS VAX and Step 1 OpenVMS AXP device drivers) expects, as input, a bit number indicating the requested I/O function in R7, and the address of the FDT entry from which it received control in R8.
R0, R7, and R8 are not provided as input to ACP_STD\$ACCESS.
- ACP\$ACCESS returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. ACP_STD\$ACCESS returns to its caller, passing it SS\$_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

System Routines ACP_STD\$ACCESSNET

ACP_STD\$ACCESSNET

Connects to network function processing.

Module

SYSACPFDT

Format

status = ACP_STD\$ACCESSNET (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp
I/O request packet.

pcb
Process control block of the current process.

ucb
Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb
Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$ACCVIO Access violation.
SS\$NORMAL The I/O request has been successfully queued to the appropriate ACP or XQP.
SS\$EXQUOTA File quota exceeded.
SS\$FILALRACC File already accessed.
SS\$IVCHNLSEC Invalid section channel.

Context

FDT dispatching code in the \$QIO system service calls ACP_STD\$ACCESSNET as an upper-level FDT action routine at IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine ACP\$ACCESSNET (used by OpenVMS VAX and Step 1 OpenVMS AXP device drivers) expects, as input, a bit number indicating the requested I/O function in R7, and the address of the FDT entry from which it received control in R8.

R0, R7, and R8 are not provided as input to ACP_STD\$ACCESSNET.

- ACP\$ACCESSNET returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. ACP_STD\$ACCESSNET returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

ACP_STD\$DEACCESS

Deaccesses ACP function processing.

Module

SYSACPFDT

Format

status = ACP_STD\$DEACCESS (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp
I/O request packet.

pcb
Process control block of the current process.

ucb
Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb
Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$FILNOTACC File not accessed.
SS\$IVCHNLSEC Invalid section channel.
SS\$NORMAL Normal, successful completion.

Context

FDT dispatching code in the \$QIO system service calls ACP_STD\$DEACCESS as an upper-level FDT action routine at IPL\$_ASTDEL.

Description

For Digital internal use only.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine ACP\$DEACCESS (used by OpenVMS VAX and Step 1 OpenVMS AXP device drivers) expects, as input, a bit number indicating the requested I/O function in R7, and the address of the FDT entry from which it received control in R8.

R0, R7, and R8 are not provided as input to ACP_STD\$DEACCESS.

- ACP\$DEACCESS returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. ACP_STD\$DEACCESS returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

System Routines

ACP_STD\$MODIFY

ACP_STD\$MODIFY

Deletes and modifies ACP function processing.

Module

SYSACPFDT

Format

status = ACP_STD\$MODIFY (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp
I/O request packet.

pcb
Process control block of the current process.

ucb
Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb
Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$ACCVIO Access violation.
SS\$DEVNOTMOUNT Device not mounted.
SS\$DEVFOREIGN Device is mounted as foreign.
SS\$EXQUOTA File quota exceeded.
SS\$NORMAL The I/O request has been successfully queued to the appropriate ACP or XQP.

Context

FDT dispatching code in the \$QIO system service calls ACP_STD\$MODIFY as an upper-level FDT action routine at IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine ACP\$MODIFY (used by OpenVMS VAX and Step 1 OpenVMS AXP device drivers) expects, as input, a bit number indicating the requested I/O function in R7, and the address of the FDT entry from which it received control in R8.

R0, R7, and R8 are not provided as input to ACP_STD\$MODIFY.

- ACP\$MODIFY returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. ACP_STD\$MODIFY returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

System Routines ACP_STD\$MOUNT

ACP_STD\$MOUNT

Initiates ACP mount function processing.

Module

SYSACPFDT

Format

status = ACP_STD\$MOUNT (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp
I/O request packet.

pcb
Process control block of the current process.

ucb
Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb
Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$ACCVIO Access violation.
SS\$DEVNOTMOUNT Device not mounted.
SS\$NOPRIV Process has insufficient privileges.
SS\$NORMAL The I/O request has been successfully queued to the appropriate ACP or XQP.

Context

FDT dispatching code in the \$QIO system service calls ACP_STD\$MOUNT as an upper-level FDT action routine at IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine ACP\$MOUNT (used by OpenVMS VAX and Step 1 OpenVMS AXP device drivers) expects, as input, a bit number indicating the requested I/O function in R7, and the address of the FDT entry from which it received control in R8.

R0, R7, and R8 are not provided as input to ACP_STD\$MOUNT.

- ACP\$MOUNT returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. ACP_STD\$MOUNT returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

ACP_STD\$READBLK

Processes a read block ACP function.

Module

SYSACPFDT

Format

status = ACP_STD\$READBLK (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp
I/O request packet.

pcb
Process control block of the current process.

ucb
Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb
Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$ACCVIO Access violation.
SS\$ENDOFFILE End of file reached.
SS\$FILNOTACC File not accessed on channel.
SS\$NOPRIV Process has insufficient privileges.
SS\$ILLIOFUNC Illegal I/O function.

SS\$_ILLBLKNUM	Illegal block number.
SS\$_NORMAL	Normal, successful completion.
SS\$_INSFWSL	Insufficient working set limit.

Context

FDT dispatching code in the \$QIO system service calls ACP_STD\$READBLK as an upper-level FDT action routine at IPL\$_ASTDEL.

Description

For Digital internal use only.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine ACP\$READBLK (used by OpenVMS VAX and Step 1 OpenVMS AXP device drivers) expects, as input, a bit number indicating the requested I/O function in R7, and the address of the FDT entry from which it received control in R8.
R0, R7, and R8 are not provided as input to ACP_STD\$READBLK.
- ACP\$READBLK returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. ACP_STD\$READBLK returns to its caller, passing it SS\$_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

System Routines ACP_STD\$WRITEBLK

ACP_STD\$WRITEBLK

Processes a write block ACP function.

Module

SYSACPFDT

Format

status = ACP_STD\$WRITEBLK (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp
I/O request packet.

pcb
Process control block of the current process.

ucb
Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb
Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$ACCVIO Access violation.
SS\$BADPARAM Record size is too small for magtape function processing.
SS\$ENDOFFILE End of file reached.
SS\$FILNOTACC File not accessed on channel.
SS\$NOPRIV Process has insufficient privileges.

SS\$_ILLIOFUNC	Illegal I/O function.
SS\$_ILLBLKNUM	Illegal block number.
SS\$_INSFMEM	Insufficient memory to perform erase function.
SS\$_INSFSPTS	Insufficient system page table entries to perform erase function.
SS\$_INSFWSL	Insufficient working set limit.
SS\$_NORMAL	Normal, successful completion.
SS\$_WRITLCK	Device software is write locked.

Context

FDT dispatching code in the \$QIO system service calls ACP_STD\$WRITEBLK as an upper-level FDT action routine at IPL\$_ASTDEL.

Description

For Digital internal use only.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine ACP\$WRITEBLK (used by OpenVMS VAX and Step 1 drivers) expects, as input, a bit number indicating the requested I/O function in R7, and the address of the FDT entry from which it received control in R8. R0, R7, and R8 are not provided as input to ACP_STD\$WRITEBLK.
- ACP\$WRITEBLK returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. ACP_STD\$WRITEBLK returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

System Routines

COM_STD\$DELATTNAST

COM_STD\$DELATTNAST

Delivers all attention ASTs linked in the specified list.

Module

COMDRVSUB

Format

COM_STD\$DELATTNAST (acb_lh, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
acb_lh	address	input	reference	required
ucb	UCB	input	reference	required

ast_lh

Listhead of AST control blocks

ucb

Unit control block.

Context

COM_STD\$DELATTNAST executes and exits at the caller's IPL, and acquires no spin locks. However, the caller must be executing at IPL\$_RESCHED or higher to avoid certain race conditions.

Description

COM_STD\$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_STD\$DELATTNAST dequeues each ACB from the head of the list, thus removing them in the reverse order of their declaration by COM_STD\$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.

Macro

CALL_DELATTNAST [save_r0r1]

where:

save_r0r1 indicates that the macro should preserve registers R0 and R1 across the call to COM_STD\$DELATTNAST. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

System Routines COM_STD\$DELATTNAST

In a Step 2 driver, CALL_DELATTNAST simulates a JSB to COM\$DELATTNAST. It calls COM_STD\$DELATTNAST using the current contents of R4 and R5 as the **listhead** and **ucb** arguments, respectively. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note that COM_STD\$DELATTNAST replaces COM\$DELATTNAST (used by Step 1 drivers and OpenVMS VAX drivers). Unlike COM\$DELATTNAST, COM_STD\$DELATTNAST does not preserve the contents of R0 and R1.

COM_STD\$DELATTNASTP

Delivers all attention ASTs linked in the specified list for a given process.

Module

COMDRVSUB

Format

COM_STD\$DELATTNASTP (acb_lh, ucb, ipid)

Arguments

Argument	Type	Access	Mechanism	Status
acb_lh	listhead	input	reference	required
ucb	UCB	input	reference	required
ipid	integer	input	value	required

acb_lh

Listhead of AST control blocks

ucb

Unit control block.

ipid

Internal process ID (IPID) for the target process.

Context

COM_STD\$DELATTNASTP executes and exits at the caller's IPL, and acquires no spin locks. However, the caller must be executing at IPL\$_RESCHED or higher to avoid certain race conditions.

Description

For Digital internal use only.

Macro

CALL_DELATTNASTP [save_r0r1]

where:

save_r0r1 indicates that the macro should preserve registers R0 and R1 across the call to COM_STD\$DELATTNASTP. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

In a Step 2 driver, CALL_DELATTNASTP simulates a JSB to COM\$DELATTNASTP. It calls COM_STD\$DELATTNASTP using the current contents of R4, R5 and R6 as the **listhead**, **ucb**, and **ipid** arguments, respectively. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- COM_STD\$DELATTNASTP replaces COM\$DELATTNASTP (used by Step 1 drivers and OpenVMS VAX drivers). Unlike COM\$DELATTNASTP, COM_STD\$DELATTNASTP does not preserve the contents of R0 and R1.

System Routines

COM_STD\$DELCTRLAST

COM_STD\$DELCTRLAST

Delivers all control ASTs, linked in the specified list, that match a given condition.

Module

COMDRVSUB

Format

COM_STD\$DELCTRLAST (acb_lh, ucb, matchchar, inclchar_p)

Arguments

Argument	Type	Access	Mechanism	Status
acb_lh	listhead	input	reference	required
ucb	UCB	input	reference	required
matchchar	integer	input	value	required
inclchar_p	pointer	output	value	required

acb_lh

Listhead of AST control blocks

ucb

Unit control block.

matchchar

Match character.

inclchar_p

Address in which COM_STD\$DELCTRLAST writes the character to include in the data stream, or NULL.

Context

COM_STD\$DELCTRLAST executes and exits at the caller's IPL, and acquires no spin locks. However, the caller must be executing at IPL\$RESCHED or higher to avoid certain race conditions.

Description

For Digital internal use only.

Macro

CALL_DELCTRLAST [save_r0r1]

where:

save_r0r1 indicates that the macro should preserve registers R0 and R1 across the call to COM_STD\$DELCTRLAST. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

In a Step 2 driver, CALL_DELCTRLAST simulates a JSB to COM\$DELCTRLAST. It calls COM_STD\$DELCTRLAST using the current contents of R4, R5, and R3 as the **listhead**, **ucb**, and **matchchar** arguments, respectively. When COM\$DELCTRLAST returns, it moves the include character into R3. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- COM_STD\$DELCTRLAST replaces COM\$DELCTRLAST (used by Step 1 drivers and OpenVMS VAX drivers). Unlike COM\$DELCTRLAST, COM_STD\$DELCTRLAST does not preserve the contents of R0 and R1.

System Routines

COM_STD\$DELCTRLASTP

COM_STD\$DELCTRLASTP

Delivers all control ASTs, linked in the specified list, that match a given condition.

Module

COMDRVSUB

Format

COM_STD\$DELCTRLASTP (acb_lh, ucb, ipid, matchchar, inclchar_p)

Arguments

Argument	Type	Access	Mechanism	Status
acb_lh	listhead	input	reference	required
ucb	UCB	input	reference	required
ipid	integer	input	value	required
matchchar	integer	input	value	required
inclchar_p	pointer	input	value	required

acb_lh

Listhead of AST control blocks

ucb

Unit control block.

ipid

Internal process ID (IPID) for the target process.

matchchar

Match character.

inclchar_p

Address in which COM_STD\$DELCTRLASTP writes the character to include in the data stream, or NULL.

Context

COM_STD\$DELCTRLASTP executes and exits at the caller's IPL, and acquires no spin locks. However, the caller must be executing at IPL\$_RESCHED or higher to avoid certain race conditions.

Description

For Digital internal use only.

Macro

CALL_DELCTRLASTP [save_r0r1]

where:

save_r0r1 indicates that the macro should preserve registers R0 and R1 across the call to COM_STD\$DELCTRLASTP. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

In a Step 2 driver, CALL_DELCTRLASTP simulates a JSB to COM\$DELCTRLASTP. It calls COM_STD\$DELCTRLASTP using the current contents of R4, R5, R6, and R3 as the **listhead**, **ucb**, **ipid**, and **matchchar** arguments, respectively. When COM\$DELCTRLASTP returns, it moves the include character into R3. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- COM_STD\$DELCTRLASTP replaces COM\$DELCTRLASTP (used by Step 1 drivers and OpenVMS VAX drivers). Unlike COM\$DELCTRLASTP, COM_STD\$DELCTRLASTP does not preserve the contents of R0 and R1.

COM_STD\$DRVDEALMEM

Deallocates system dynamic memory.

Module

MEMORYALC_MIN or MEMORYALC_MON

Format

COM_STD\$DRVDEALMEM (block)

Arguments

Argument	Type	Access	Mechanism	Status
ptr	structure	input	reference	required

ptr

Block to be deallocated. The block must be a standard OpenVMS data structure (in which offset FKBSW_SIZE contains its size). The block size must be at least FKBSK_LENGTH (24 bytes). (The FKBS symbols are defined by the \$FKBDEF macro in SYSSLIBRARY:LIB.MLB.)

Context

A driver can call COM_STD\$DRVDEALMEM from any IPL. COM_STD\$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_STD\$DRVDEALMEM transfers control to EXE\$DEANONPAGED to deallocate the buffer specified by the **block** parameter. If COM_STD\$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 FKBS_C_LENGTH in size, or its address is not aligned on a 16-byte boundary, COM_STD\$DRVDEALMEM issues a BADDALRQSZ bugcheck.

Macro

CALL_DRVDEALMEM [save_r0r1]

where:

save_r0r1 indicates that the macro should preserve registers R0 and R1 across the call to COM_STD\$DRVDEALMEM. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro

generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

In a Step 2 driver, **CALL_DRVDEALMEM** simulates a JSB to **COM\$DRVDEALMEM**. It calls **COM_STD\$DRVDEALMEM** using the current contents of **R0** as the **ptr** argument. Unless you specify **save_r0r1=NO**, the macro preserves the quadword registers **R0** and **R1** across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- **COM_STD\$DRVDEALMEM** replaces **COM\$DRVDEALMEM** (used by Step 1 drivers and OpenVMS VAX drivers). Unlike **COM\$DRVDEALMEM**, **COM_STD\$DRVDEALMEM** does not preserve the contents of **R0** and **R1**.

System Routines

COM_STD\$FLUSHATTNS

COM_STD\$FLUSHATTNS

Removes specified ASTs from an attention AST list.

Module

COMDRVSUB

Format

status = COM_STD\$FLUSHATTNS (pcb, ucb, chan ,acb_lh)

Arguments

Argument	Type	Access	Mechanism	Status
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
chan	integer	input	value	required
acb_lh	listhead	input	reference	required

pcb

Process control block. COM_STD\$FLUSHATTNS reads the following PCB fields:

Field	Contents
PCB\$L_PID	Process ID
PCB\$L_ASTCNT	ASTs remaining in quota

COM_STD\$FLUSHATTNS increases PCB\$L_ASTCNT once for each AST control block (ACB) it flushes.

ucb

Unit control block. COM_STD\$FLUSHATTNS reads UCB\$L_DLCK to obtain the address of the device lock.

chan

Number of the assigned I/O channel.

acb_lh

Listhead of ACBs.

Return Values

SSS_NORMAL

Normal, successful completion

Context

COM_STD\$FLUSHATTNS raises IPL to device IPL, acquiring the corresponding device lock. Before returning control to its caller at the caller's IPL, COM_STD\$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_STD\$FLUSHATTNS to flush an attention AST list. A driver FDT routine calls COM_STD\$FLUSHATTNS to service a \$QIO request that specifies a set-attention-AST function and a value of 0 in the **p1** argument (IRPSL_QIO_P1).

COM_STD\$FLUSHATTNS locates all ACBs 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.

Macro

CALL_FLUSHATTNS

In a Step 2 driver, CALL_FLUSHATTNS simulates a JSB to COM\$FLUSHATTNS. It calls COM_STD\$FLUSHATTNS using the current contents of R4, R5, R6, and R7 as the **pcb**, **ucb**, **chan**, and **acb_lh** arguments, respectively. It returns status in R0.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- COM_STD\$FLUSHATTNS replaces COM\$FLUSHATTNS (used by Step 1 drivers and OpenVMS VAX drivers).

System Routines

COM_STD\$FLUSHCTRLS

COM_STD\$FLUSHCTRLS

Removes specified ASTs from a control AST list.

Module

COMDRVSUB

Format

status = COM_STD\$FLUSHCTRLS (pcb, ucb, chan ,acb_lh, mask_p)

Arguments

Argument	Type	Access	Mechanism	Status
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
chan	integer	input	value	required
acb_lh	listhead	input	reference	required
mask_p	mask_ longword	input	reference	required

pcb

Process control block. COM_STD\$FLUSHCTRLS reads the following PCB fields:

Field	Contents
PCB\$L_PID	Process ID
PCB\$L_ASTCNT	ASTs remaining in quota

COM_STD\$FLUSHCTRLS increases PCB\$L_ASTCNT once for each control AST control block (TAST) it flushes.

ucb

Unit control block. COM_STD\$FLUSHCTRLS reads UCB\$L_DLCK to obtain the address of the device lock.

chan

Number of the assigned I/O channel.

acb_lh

Listhead of ACBs.

mask_p

Summary mask of active control characters. COM_STD\$FLUSHCTRLS updates this mask.

Return Values

SS\$NORMAL	Normal, successful completion
------------	-------------------------------

Context

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

Description

For Digital internal use only.

Macro

CALL_FLUSHCTRLS

In a Step 2 driver, CALL_FLUSHCTRLS simulates a JSB to COM\$FLUSHCTRLS. It calls COM_STD\$FLUSHCTRLS using the current contents of R2, R4, R5, R6, and R7 as the **mask**, **pcb**, **ucb**, **chan**, and **acb_lh** arguments, respectively. It returns status in R0.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- COM_STD\$FLUSHCTRLS replaces COM\$FLUSHCTRLS (used by Step 1 drivers and OpenVMS VAX drivers).

System Routines

COM_STD\$POST, COM_STD\$POST_NOCNT

COM_STD\$POST, COM_STD\$POST_NOCNT

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

Module

COMDRVSUB

Format

COM_STD\$POST (irp, ucb)

COM_STD\$POST_NOCNT (irp)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

irp

I/O request block. The following IRP fields are input to I/O postprocessing.

Field	Contents
IRP\$!_MEDIA	Data to be copied to the I/O status block
IRP\$!_MEDIA+4	Data to be copied to the I/O status block

ucb

Unit control block (COM_STD\$POST only). COM_STD\$POST increases the unit operation count (UCB\$!_OPCNT).

Context

Drivers call COM_STD\$POST at or above fork IPL. Drivers call COM_STD\$POST_NOCNT at or above IPL\$!_ASTDEL. These routines execute at their caller's 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_STD\$POST or COM_STD\$POST_NOCNT after it has completed device-dependent I/O processing for an I/O request initiated by EXE_STD\$ALTQUEPKT. Because COM_STD\$POST_NOCNT, unlike COM_STD\$POST, does not increment the unit's operations count (UCB\$!_OPCNT), a driver uses COM_STD\$POST_NOCNT to initiate completion processing for an I/O request when the associated UCB is not available.

System Routines COM_STD\$POST, COM_STD\$POST_NOCNT

COM_STD\$POST and COM_STD\$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_STD\$REQCOM, these routines do not attempt to dequeue any IRP waiting for the device or change the busy status of the device.

Macro

```
CALL_POST [save_r1]
CALL_POST_NOCNT [save_r1]
```

where:

save_r1 indicates that the macro should preserve register R1 across the call to COM_STD\$POST or COM_STD\$POST_NOCNT. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

In a Step 2 driver, CALL_POST simulates a JSB to COM\$POST. It calls COM_STD\$POST using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively. CALL_POST_NOCNT simulates a JSB to COM\$POST_NOCNT. It calls COM_STD\$POST_NOCNT using the current contents of R3 as the **irp** argument. Unless you specify **save_r1=NO**, the macro preserves the quadword register R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- COM_STD\$POST replaces COM\$POST (used by Step 1 drivers and OpenVMS VAX drivers); COM_STD\$POST_NOCNT replaces COM\$POST_NOCNT. Unlike the Step 1 routines, the Step 2 routines do not preserve R1 across the call.

System Routines

COM_STD\$SETATTNAST

COM_STD\$SETATTNAST

Enables or disables attention ASTs.

Module

COMDRVSUB

Format

status = COM_STD\$SETATTNAST (irp, pcb, ucb, ccb, acb_lh)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required
acb_lh	listhead	input	reference	required

irp

I/O request packet for the current I/O request.

COM_STD\$SETATTNAST reads the following IRP fields:

Field	Contents
IRP\$L_QIO_P1	\$QIO system service p1 argument, containing the address of the AST routine, or zero to flush the AST queue.
IRP\$L_QIO_P2	\$QIO system service p2 argument, containing the AST parameter.
IRP\$L_QIO_P3	\$QIO system service p3 argument, containing the access mode of the AST request.
IRP\$L_CHAN	I/O request channel index number.

pcb

Process control block of the current process.

COM_STD\$SETATTNAST reads the following PCB fields:

Field	Contents
PCB\$L_ASTCNT	Number of ASTs remaining in process quota
PCB\$L_PID	Process ID

COM_STD\$SETATTNAST decreases PCB\$L_ASTCNT if it successfully queues the AST.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

COM_STD\$SETATTNAST reads UCB\$L_DLCK.

ccb

Channel control block that describes the process-I/O channel

acb_lh

Address of listhead of AST control blocks.

Return Values

SS\$_FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.
SS\$_NORMAL	Normal, successful completion

Status in FDT_CONTEXT

SS\$_EXQUOTA	Process AST quota exceeded.
SS\$_INSFMEM	No memory available to allocate the expanded ACB.

Context

The FDT support routine COM_STD\$SETATTNAST must be called from code executing at IPL\$_ASTDEL. COM_STD\$SETATTNAST raises IPL and acquires the corresponding device lock, to insert the AST into the AST queue. It returns control to its caller at IPL\$_ASTDEL.

Description

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

If the **p1** argument of the request contains a zero, COM_STD\$SETATTNAST transfers control to COM_STD\$FLUSHATTNS, which disables all ASTs indicated by the PID and I/O channel number (IRP\$L_CHAN). COM_STD\$FLUSHATTNS searches through the AST control block (ACB) list, extracts each identified ACB, deallocates it, and returns SS\$_NORMAL status in R0 to COM_STD\$SETATTNAST. COM_STD\$SETATTNAST returns this status to its caller.

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

- Spin lock index SPL\$_QUEUEAST
- Address of the AST routine (as specified in **p1**)
- AST parameter (as specified in **p2**)
- Access mode (the maximum, or least privileged, access mode between the access mode specified in **p3** and the current process's access mode). Bit ACB\$_QUOTA is set in this value to indicate that the AST was requested by a process, not by the system.

System Routines

COM_STD\$SETATTNAST

- Number of the assigned I/O channel
- PID of the requesting process

COM_STD\$SETATTNAST links the ACB to the start of the specified linked list of ACBs located in a UCB extension area. 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 its AST quota, or if there is no memory available to allocate the expanded ACB, COM_STD\$SETATTNAST restores PCB\$_ASTCNT to its original value and calls EXE_STD\$ABORTIO, passing it a **qio_sts** of SSS_BADPARAM. When it regains control, COM_STD\$SETATTNAST returns to its caller with this status in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0.

The caller of COM_STD\$SETATTNAST must examine the status in R0:

- If the status is SSS_NORMAL, the attention AST has been enabled (or the AST has been flushed), as requested.
- If the status is SSS_FDT_COMPL, an error has occurred that has caused the operation to be aborted. You can determine the reason for the failure from FDT_CONTEXT\$_QIO_STATUS.

Macro

CALL_SETATTNAST

In a Step 2 driver, CALL_SETATTNAST simulates a JSB to COM\$SETATTNAST. It calls COM_STD\$SETATTNAST using the current contents of R3, R4, R5, R6, and R7, as the **irp**, **pcb**, **ucb**, **ccb**, and **acb_lh** arguments, respectively. It returns status in R0 and in the FDT_CONTEXT structure.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- COM_STD\$SETATTNAST replaces COM\$SETATTNAST (used by Step 1 drivers and OpenVMS VAX drivers). COM\$SETATTNAST returns to its caller only upon success; COM_STD\$SETATTNAST returns to its caller whether it has been successful or not. It returns SSS_NORMAL or SSS_FDT_COMPL status in R0. When it returns SSS_FDT_COMPL status, the FDT_CONTEXT structure contains additional status (SS\$_EXQUOTA or SSS_INSMEM) to explain why the request has been aborted.
- COM\$SETATTNAST preserves the addresses of the IRP and UCB in R3 and R5 across the call. COM_STD\$SETATTNAST does not.

COM_STD\$SETCTRLAST

Enables or disables control ASTs.

Module

COMDRVSUB

Format

status = COM_STD\$SETCTRLAST (irp, pcb, ucb, acb_lh, mask, tast_p)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
acb_lh	listhead	input	reference	required
mask	mask_ longword	input	value	required
tast_p	TAST	output	value	required

irp

I/O request packet for the current I/O request.

COM_STD\$SETCTRLAST reads the following IRP fields:

Field	Contents
IRP\$QIO_P1	\$QIO system service p1 argument, containing the address of the AST routine to call when an out-of-band character is typed, or zero to flush the queue.
IRP\$QIO_P2	\$QIO system service p2 argument, containing the address of the short-form terminator mask, indicating which out-of-band characters precipitate AST delivery. This address is passed as an AST parameter when the AST is delivered.
IRP\$QIO_P3	\$QIO system service p3 argument, containing the access mode of the AST request.
IRP\$CHAN	I/O request channel index number

pcb

Process control block of the current process.

System Routines

COM_STD\$SETCTRLAST

COM_STD\$SETCTRLAST reads the following PCB fields:

Field	Contents
PCB\$L_ASTCNT	Number of ASTs remaining in process quota
PCB\$L_PID	Process ID

COM_STD\$SETCTRLAST decreases PCB\$L_ASTCNT if it successfully queues the AST.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

COM_STD\$SETCTRLAST reads UCB\$L_DLCK.

acb_lh

Address of listhead of AST control blocks.

mask

Summary mask of active control characters. COM_STD\$SETCTRLAST updates the summary mask to be the inclusive-OR of all masks in the control AST list.

tast_p

Address of the control AST block (TAST), returned as output from COM_STD\$SETCTRLAST.

Return Values

SS\$FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.
SS\$NORMAL	Normal, successful completion

Status in FDT_CONTEXT

SS\$ACCVIO	Specified mask is not addressable.
SS\$EXQUOTA	Process AST quota exceeded.
SS\$INSFMEM	No memory available to allocate the expanded ACB.

Context

The FDT support routine COM_STD\$SETCTRLAST must be called from code executing at IPL\$ASTDEL. COM_STD\$SETCTRLAST raises IPL and acquires the corresponding device lock, to insert the AST into the AST queue. It returns control to its caller at IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

CALL_SETCTRLAST

In a Step 2 driver, CALL_SETCTRLAST simulates a JSB to COM\$SETCTRLAST. It calls COM_STD\$SETCTRLAST using the current contents of R3, R4, R5, R7, and R2, as the **irp**, **pcb**, **ucb**, **acb_lh**, and **mask** arguments, respectively. It returns the TAST block in R2. It returns status in R0 and in the FDT_CONTEXT structure.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- COM_STD\$SETCTRLAST replaces COM\$SETCTRLAST (used by Step 1 drivers and OpenVMS VAX drivers). The order in which formal parameters are passed to COM_STD\$SETCTRLAST differs from the order in which they are provided in registers to the Step 1 routine COM\$SETCTRLAST.
- COM_STD\$SETCTRLAST does not provide the address of the TAST block as output in R2.
- COM\$SETCTRLAST returns to its caller only upon success; COM_STD\$SETCTRLAST returns to its caller whether it has been successful or not. It returns SSS_NORMAL or SSS_FDT_COMPL status in R0. When it returns SSS_FDT_COMPL status, the FDT_CONTEXT structure contains additional status (SSS_EXQUOTA or SSS_INSFMEM) to explain why the request has been aborted.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- ERL_STD\$ALLOCEMB replaces ERL\$ALLOCEMB (used by Step 1 drivers and OpenVMS VAX drivers). Unlike ERL\$ALLOCEMB, ERL_STD\$ALLOCEMB does not return the error sequence number in R1. A driver can obtain the error sequence number from the error message buffer (EMBSW_DV_ERRSEQ).

System Routines

ERL_STD\$DEVICEATTN, ERL_STD\$DEVICERR, ERL_STD\$DEVICTMO

ERL_STD\$DEVICEATTN, ERL_STD\$DEVICERR, ERL_STD\$DEVICTMO

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

Module

ERRORLOG

Format

ERL_STD\$DEVICEATTN (driver_param, ucb)

ERL_STD\$DEVICERR (driver_param, ucb)

ERL_STD\$DEVICTMO (driver_param, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
driver_param	undefined	input	reference	required
ucb	UCB	input	reference	required

driver_param

Parameter to be passed to the register dumping routine, usually a controller register access mailbox (CRAM).

ucb

Unit control block. These routines read the following UCB fields:

Field	Contents
UCB\$L_DEVCHAR	Bit DEV\$V_ELG set.
UCB\$L_FUNC	Bit IOSV_INHERLOG clear.
UCB\$L_IRP	Address of IRP currently being processed (ERL_STD\$DEVICERR and ERL_STD\$DEVICTMO only).
UCB\$L_ORB	ORB address.
UCB\$L_DDB	DDB address.
UCB\$L_DDT	DDT address. DDT\$W_ERRORBUF contains the size of the error message buffer in bytes.

These routines write the following UCB fields:

Field	Contents
UCB\$L_ERRCNT	Increased.
UCB\$L_EMB	Address of error message buffer.
UCB\$L_STS	UCB\$V_ERLOGIP set.

Context

A driver calls ERL_STD\$DEVICEATTN, ERL_STD\$DEVICERR, or ERL_STD\$DEVICTMO at or above fork IPL, holding the corresponding fork lock in an OpenVMS 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_STD\$DEVICERR and ERL_STD\$DEVICTMO log an error associated with a particular I/O request. ERL_STD\$DEVICEATTN logs an error that is not associated with an I/O request. Each of these routines performs the following steps:

1. Increases UCBSL_ERRCNT to record a device error. If the error-log-in-progress bit (UCBSV_ERLOGIP in UCBSL_STS) is set, the routine returns control to its caller.
2. 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.
3. Places the address of the error message buffer in UCBSL_EMB.
4. Sets UCBSV_ERLOGIP in UCBSL_STS.
5. 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_STD\$DEVICEATTN	Device attention (EMB\$C_DA)
ERL_STD\$DEVICERR	Device error (EMB\$C_DE)
ERL_STD\$DEVICTMO	Device timeout (EMB\$C_DT)

6. Loads fields from the UCB, the IRP, and the DDB into the buffer, including the following:

UCBSB_DEVCLASS	Device class
UCBSB_DEVTYPE	Device type
IRP\$\$_PID	Process ID of the process originating the I/O request (ERL_STD\$DEVICERR or ERL_STD\$DEVICTMO)
IRP\$\$_BOFF	Transfer parameter (ERL_STD\$DEVICERR and ERL_STD\$DEVICTMO)
IRP\$\$_BCNT	Transfer parameter (ERL_STD\$DEVICERR and ERL_STD\$DEVICTMO)
IRP\$\$_MEDIA	Disk address
UCBSW_UNIT	Unit number
UCBSL_ERRCNT	Count of device errors
UCBSL_OPCNT	Count of completed operations
ORB\$\$_OWNER	UIC of volume owner

System Routines

ERL_STD\$DEVICEATTN, ERL_STD\$DEVICERR, ERL_STD\$DEVICTMO

- | | |
|--------------|--|
| UCB\$DEVCHAR | Device characteristics |
| IRP\$FUNC | I/O function value (ERL_STD\$DEVICERR and ERL_STD\$DEVICTMO) |
| DDB\$NAME | Device name (concatenated with cluster node name if appropriate) |
7. Loads into R0 the address of the location in the buffer in which the contents of the device registers are to be stored.
 8. Calls the driver's register dumping routine, the address of which is specified in the **regdmp** argument to the DDTAB macro.

Macro

```
CALL_DEVICEATTN [save_r0r1]
CALL_DEVICERR [save_r0r1]
CALL_DEVICTMO [save_r0r1]
```

where:

save_r0r1 indicates that the macros must preserve the contents of R0 and R1 across the call to ERL_STD\$DEVICEATTN, ERL_STD\$DEVICERR, or ERL_STD\$DEVICTMO. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

In a Step 2 driver, the CALL_DEVICEATTN, CALL_DEVICERR, and CALL_DEVICTMO macros simulate JSBs to ERL\$DEVICEATTN, ERL\$DEVICERR, and ERL\$DEVICTMO, respectively. Each macro calls the corresponding routine using the current contents of R4 and R5 as the **driverpar** and **ucb** arguments, respectively. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- ERL_STD\$DEVICEATTN, ERL_STD\$DEVICERR, and ERL_STD\$DEVICTMO replace ERL\$DEVICEATTN, ERL\$DEVICERR, and ERL\$DEVICTMO (used by Step 1 drivers and OpenVMS VAX drivers). Unlike the Step 1 routines, the Step 2 routines do not preserve the contents of R0 and R1.
- Because the UCB\$MEDIA field has been removed from the UCB local disk extension, these routines write the disk address into the EMB from IRP\$MEDIA.
- Because the UCB\$SLAVE field has been removed from the UCB local disk extension, these routines do not write that field.
- OpenVMS AXP device drivers consequently do not need to define the local disk UCB extension or local tape UCB extension to use these error logging routines.
- **driver_param** is considered required input to these routines.

ERL_STD\$RELEASEMB

Releases an error message buffer to the error logging process.

Module

ERRORLOG

Format

ERL_STD\$RELEASEMB (embdv)

Arguments

Argument	Type	Access	Mechanism	Status
embdv	EMBDV	input	reference	required

embdv

Error message buffer to be released.

Context

A driver can call ERL_STD\$RELEASEMB from any IPL. ERL_STD\$RELEASEMB raises IPL to IPL\$_EMB and obtains the corresponding spin lock to release the error message buffer. It returns control to its caller at its caller's IPL. The caller retains any spin locks it held at the time of the call.

Description

For Digital internal use only.

Macro

CALL_RELEASEMB

In a Step 2 driver, CALL_RELEASEMB simulates a JSB to ERL\$RELEASEMB. It calls ERL_STD\$RELEASEMB using the current contents of R2 as the **buff** argument.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- ERL_STD\$RELEASEMB replaces ERL\$RELEASEMB (used by Step 1 drivers and OpenVMS VAX drivers).

EXE\$BUS_DELAY

Allows a system-specific bus delay within a timed wait.

Module

[.SYSLOA]TIMEDWAIT

Macro

TIMEDWAIT

Format

EXE\$BUS_DELAY adp

Context

EXE\$BUS_DELAY conforms to the OpenVMS calling standard.

Arguments

adp

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by value

Address of ADP.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$INSFARG	Not all of the required arguments were specified.

Description

The OpenVMS VAX version of the TIMEDWAIT macro generated a processor-specific delay for the bus indicated by the ADP before executing the series of instructions, specified in the macro invocation, that check for the occurrence of a specific event or condition. In OpenVMS VAX systems, the delay helps prevent flooding the bus paths with references to device interface registers in I/O space.

An implicit call to EXE\$BUS_DELAY is included in the expansion of the TIMEDWAIT macro when you specify the **bus** argument. You can explicitly call EXE\$BUS_DELAY but, if you do, you must not also employ the TIMEDWAIT macro with the **bus** argument.

Note

In OpenVMS AXP, EXE\$BUS_DELAY checks for the required argument and, if it is present, returns to its caller with SS\$_NORMAL status.

System Routines

EXE\$DELAY

EXE\$DELAY

Provides a short-term simple delay.

Module

[SYSLOA]TIMEDWAIT

Macro

TIMEDDELAY

Format

EXE\$DELAY delta

Context

EXE\$DELAY conforms to the OpenVMS calling standard.

Arguments

delta

VMS Usage: aligned quadword
type: quadword (unsigned)
access: read only
mechanism: by reference

Delay time specified in nanoseconds.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$_NORMAL	The routine completed successfully.
SS\$_INSFARG	Not all of the required arguments were specified.

Description

EXE\$DELAY implements a simple delay by looping for at least the requested time interval. System events such as interrupt processing may have some impact on the actual time delay.

EXE\$KP_ALLOCATE_KPB

Creates a KPB and a kernel process stack, as required by the OpenVMS kernel process services.

Module

KERNEL_PROCESS_MIN, KERNEL_PROCESS_MON

Macro

KP_ALLOCATE_KPB
DDTAB (**start**=EXE\$KP_STARTIO)

Format

EXE\$KP_ALLOCATE_KPB kpb ,stack_size ,flags ,param_size

Context

EXE\$KP_ALLOCATE_KPB conforms to the OpenVMS AXP calling standard.

Because EXE\$KP_ALLOCATE_KPB raises IPL to IPL\$_SYNCH and obtains the MMG spin lock, its caller cannot be executing above IPL\$_SYNCH or hold any higher ranked spin locks. EXE\$KP_ALLOCATE_KPB returns control to its caller at its caller's IPL. The caller retains any spin locks it held at the time of the call.

Arguments

kpb

VMS Usage: address
type: longword (unsigned)
access: write only
mechanism: by reference

Address of KPB.

stack_size

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

Requested size (in bytes) of kernel process stack.

flags

VMS Usage: mask_longword
type: longword (unsigned)
access: read only
mechanism: by value

Flags indicating the type, size, and configuration of the KPB to be created. EXE\$KP_ALLOCATE_KPB accepts only the following flags:

KPBSV_VEST KPB must be a VEST KPB. (See Chapter 3 for a description of VEST KPBs.)

System Routines

EXE\$KP_ALLOCATE_KPB

KPBSV_SPLOCK	Spinlock area must be present. (Note that EXE\$KP_ALLOCATE_KPB automatically sets this bit when KPBSV_VEST is set.)
KPBSV_DEBUG	Debug area must be present.
KPBSV_DEALLOC_AT_END	KP_END should call KP_DEALLOCATE.

param_size

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

Size in bytes of KPB parameter area, if any.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$BADPARAM	An illegal value was specified in the flags argument.
SS\$INSFARG	Not all of the required arguments were specified.
SS\$INSFMEM	KPB cannot be allocated because of a failure in the nonpaged pool allocation routine.
SS\$INSFRPGS	Kernel process stack cannot be allocated because of there are not enough free pages in the system.

Description

EXE\$KP_ALLOCATE_KPB creates the KPB and the kernel process stack needed by a kernel process. It performs the following tasks:

- Verifies the contents of the **flags** parameter. If the **flags** parameter is valid, EXE\$KP_ALLOCATE_KPB uses it as the basis for the mask it writes to KPBSIS_FLAGS. It automatically sets KPBSV_SCHED for all KPBs and, for VEST KPBs, also sets KPBSV_SPLOCK. Finally, it sets KPBSV_PARAM if a non-zero **param_size** argument is specified.
- Computes the size of the KPB to be allocated. For both VEST and non-VEST KPBs, the KPB includes the base KPB and scheduling area. VEST KPBs also, by default, include the spinlock area, which is optional for non-VEST KPBs. For VEST and non-VEST KPBs alike, the debug and parameter areas are optional. The presence of KP\$V_DEBUG in the **flags** argument causes EXE\$KP_ALLOCATE_KPB to include the KPB debug area; the presence of a non-zero **param_size** argument causes it to include the KPB parameter area (rounded up to an integral number of quadwords).

System Routines EXESKP_ALLOCATE_KPB

- Allocates a KPB of the appropriate size. If the KPB cannot be allocated, it returns SSS_INSMEM status to its caller.
- Initializes the following KPB fields:

KPB\$IB_TYPE	DYN\$C_MISC
KPB\$IB_SUBTYPE	DYN\$C_KPB
KPB\$IS_FLAGS	Computed flags value
KPB\$PS_SCH_PTR	Address of KPB scheduling area
KPB\$PS_SPL_PTR	Address of KPB spinlock area, if present
KPB\$PS_DBG_PTR	Address of KPB debug area, if present
KPB\$PS_PRM_PTR	Address of KPB parameter area, if present
KPB\$IS_PRM_LENGTH	Length of the KPB parameter area, if specified, rounded up to an integral number of quadwords
- Computes the size of the kernel process stack by rounding the value of **stack_size** up to an integral number of CPU-specific pages, converting the result to bytes, and storing it in KPB\$IS_STACK_SIZE.
- Allocates and initializes sufficient system PTEs for the stack, plus two no-access guard pages. If the sufficient PTEs are not available, EXESKP_ALLOCATE_KPB deallocates the KPB and returns SSS_INSMEM status to its caller.
- Stores in KPB\$PS_STACK_BASE the system virtual address of the start of the no-access guard page at the base of the kernel process stack. The kernel process stack grows negatively from this address.
- Inserts the address of the KPB in the location specified by the **kpb** argument.

The caller of EXESKP_ALLOCATE_KPB is responsible for providing wait and retry operations in case of allocation failures.

EXE\$KP_DEALLOCATE_KPB

Deallocates a KPB and its associated kernel process stack.

Module

KERNEL_PROCESS_MIN, KERNEL_PROCESS_MON

Macro

KP_DEALLOCATE_KPB

Format

EXE\$KP_DEALLOCATE_KPB kpb

Context

EXE\$KP_DEALLOCATE_KPB conforms to the OpenVMS AXP calling standard.

EXE\$KP_DEALLOCATE_KPB forks to perform KPB cleanup and call the routines that deallocate the KPB and the kernel process stack. As a result, drivers can call EXE\$KP_DEALLOCATE_KPB from any IPL.

Arguments

kpb
VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of KPB.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$INSFARG	The kpb argument was not specified.

Description

EXE\$KP_DEALLOCATE_KPB deallocates the KPB and the associated kernel process stack. It performs the following tasks:

- Validates the structure indicated by the **kpb** argument. If the structure is not a KPB, or if it is currently valid, active, or in the process of deletion, EXE\$KP_DEALLOCATE_KPB requests an INCONSTATE bugcheck.

System Routines EXESKP_DEALLOCATE_KPB

- Indicates that KPB deletion is in progress by setting KPB\$V_DELETING in KPB\$IS_FLAGS.
- Sets up the KPB fork block (at KPB\$PS_FQFL) so that the rest of KPB cleanup can transpire at IPL\$_QUEUEAST. EXESKP_DEALLOCATE_KPB issues a call to IOC\$PRIMITIVE_FORK to queue the fork block on the IPL\$_QUEUEAST fork queue. When IOC\$PRIMITIVE_FORK returns control, EXESKP_DEALLOCATE_KPB returns SSS_NORMAL status to its caller.
- When execution resumes at IPL\$_QUEUEAST, the EXESKP_DEALLOCATE_KPB fork routine deallocates the stack and returns the KPB to nonpaged pool.

EXE\$KP_END

Terminates the execution of a kernel process.

Module

KERNEL_PROCESS_MAGIC

Macro

KP_END

Format

EXE\$KP_END kpb

Context

EXE\$KP_END conforms to the OpenVMS AXP calling standard.

The caller of EXE\$KP_END must be executing at IPL\$_RESCHED or above.

Arguments

kpb
VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of KPB.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$_NORMAL	The routine completed successfully.
SS\$_INSFARG	The kpb argument was not specified.

Description

EXE\$KP_END performs the following tasks to terminate the execution of a kernel process:

- If the **kpb** argument is not supplied, returns SS\$_INSFARG status to its caller.
- Validates the structure indicated by the **kpb** argument. If the structure is not a KPB, or if it is currently invalid or inactive, EXE\$KP_END requests an INCONSTATE bugcheck.

System Routines EXE\$KP_END

- Restores the SP of the initiator of the kernel process thread from KPB\$PS_SAVED_SP and poisons that field.
- Restores the preserved registers (as indicated by KPB\$IS_REG_MASK) and SP of the initiator of the kernel process thread.
- Marks the kernel process as inactive and invalid by clearing KPB\$V_ACTIVE and KPB\$V_VALID in KPB\$IS_FLAGS.
- If KPB\$V_DEALLOC_AT_END in KPB\$IS_FLAGS is set (as it is in VEST KPBs), call EXE\$KP_DEALLOCATE_KPB to deallocate the KPB and its associated kernel process stack.
- Returns successfully to the initiator of the kernel process thread (that is, the caller of EXE\$START_KP or EXE\$RESTART_KP).

System Routines

EXE\$KP_FORK

EXE\$KP_FORK

Stalls a kernel process in such a manner that it can be resumed by the OpenVMS fork dispatcher.

Module

KERNEL_PROCESS_MIN, KERNEL_PROCESS_MON

Macro

KP_STALL_FORK, KP_STALL_IOFORK

Format

EXE\$KP_FORK kpb [,fkb]

Context

EXE\$KP_FORK conforms to the OpenVMS AXP calling standard. It can only be called by a kernel process.

Arguments

kpb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of the caller's KPB (which must be a VEST KPB). KPB\$PS_UCB must contain the address of a UCB and KPB\$PS_IRP must contain the address of an IRP.

fkb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of a fork block, usually in the UCB. If this argument is omitted, EXE\$KP_FORK uses the fork block within the KPB (KPB\$PS_FQFL).

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$BADPARAM	The kp argument does not specify a VEST KP.
SS\$INSFARG	Not all of the required arguments were specified.

Description

EXE\$KP_FORK performs the following tasks in stalling the kernel process:

1. Saves the **kp** argument in KPB\$PS_FKBLK. If this argument is not specified to EXE\$KP_FORK, EXE\$KP_FORK writes the address of KPB\$PS_FQFL into KPB\$PS_FKBLK.
2. Inserts the procedure descriptor of subroutine STALL_FORK in KPB\$PS_SCH_STALL_RTN, thus making it the kernel process scheduling stall routine.
3. Clears KPB\$PS_SCH_RESTART, thus indicating that there is no kernel process scheduling restart routine.
4. Calls EXE\$KP_STALL_GENERAL, passing to it the address of the KP.

Having stalled the kernel process, the STALL_FORK kernel process scheduling stall routine returns control to EXE\$KP_STALL_GENERAL, which returns to the initiator of the kernel process thread (that is, the caller of EXE\$KP_START or EXE\$KP_RESTART). When the fork dispatcher ultimately resumes the suspended routine, STALL_FORK calls EXE\$KP_RESTART which, in turn, passes control back to EXE\$KP_FORK. The kernel process forking stall routine then returns to the kernel process that called it.

System Routines

EXE\$KP_FORK_WAIT

EXE\$KP_FORK_WAIT

Stalls a kernel process in such a manner that it can be resumed by the software timer interrupt service routine's examination of the fork-and-wait queue.

Module

KERNEL_PROCESS_MIN, KERNEL_PROCESS_MON

Macro

KP_STALL_FORK_WAIT

Format

EXE\$KP_FORK_WAIT kpb [,fkb]

Context

EXE\$KP_FORK_WAIT conforms to the OpenVMS AXP calling standard and can only be called by a kernel process.

The caller of EXE\$KP_FORK_WAIT must be executing at or above IPL\$_SYNCH.

Arguments

kpb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of the caller's KPB.

fkb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of a fork block. If this argument is omitted, EXE\$KP_FORK_WAIT uses the fork block within the KPB (KPB\$PS_FKBLK).

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$INSFARG	Not all of the required arguments were specified.

Description

EXE\$KP_FORK_WAIT performs the following tasks in stalling a kernel process:

1. Saves the **fk** argument, if specified, in KPB\$PS_FKBLK. If the argument is not specified, EXE\$KP_FORK_WAIT moves the address of KPB\$PS_FQFL into KPB\$PS_FKBLK.
2. Inserts the procedure descriptor of subroutine STALL_FORK_WAIT in KPB\$PS_SCH_STALL_RTN, thus making it the kernel process scheduling stall routine.
3. Clears KPB\$PS_SCH_RESTART, thus indicating that there is no kernel process scheduling restart routine.
4. Calls EXE\$KP_STALL_GENERAL, passing to it the address of the KPB.

Note that, having stalled the kernel process, the STALL_FORK_WAIT kernel process scheduling stall routine returns control to EXE\$KP_STALL_GENERAL, which returns to the initiator of the kernel process thread (that is, the caller of EXE\$KP_START or EXE\$KP_RESTART). When the fork block is ultimately removed from the fork-and-wait-queue, STALL_FORK_WAIT calls EXE\$KP_RESTART which, in turn, passes control back to EXE\$KP_FORK_WAIT. EXE\$KP_FORK_WAIT then returns to kernel process that called it.

EXE\$KP_RESTART

Resumes the execution of a kernel process.

Module

KERNEL_PROCESS_MAGIC

Macro

KP_RESTART

Format

EXE\$KP_RESTART kpb [,thread_status]

Context

EXE\$KP_RESTART conforms to the OpenVMS AXP calling standard.

The caller of EXE\$KP_RESTART, usually a kernel process scheduling stall routine, must be executing at IPL\$_RESCHED or above.

Arguments

kpb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of KPB.

thread_status

VMS Usage: longword (unsigned)
type: read only
access: by value
mechanism:

Status value to be returned to the kernel process that is to be resumed. This is the status returned by the call to EXE\$KP_STALL_GENERAL. If the **thread_status** argument is not present, EXE\$KP_RESTART returns SSS_NORMAL status to the kernel process.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$INSFARG	The kpb argument was not specified.

Description

EXE\$KP_RESTART performs the following tasks to restart a kernel process:

1. Validates the structure indicated by the **kpb** argument. If the structure is not a KPB, or if it is currently invalid, EXE\$KP_RESTART requests an INCONSTATE bugcheck.
2. Preserves the current context by saving the current stack pointer (SP) and the registers indicated by KPB\$IS_REG_MASK on the stack (which it quadword-aligns after obtaining the current SP). It saves the new value of the SP in KPB\$PS_SAVED_SP.
3. Restores the SP of the stalled kernel process from KPB\$PS_STACK_SP.
4. Restores the preserved registers (as indicated by KPB\$IS_REG_MASK) from the top of the kernel process stack, plus the original SP of the kernel process stack.
5. Makes the KPB active by setting the corresponding bit in KPB\$IS_FLAGS.
6. Calls the kernel process scheduling restart routine, if one is specified, passing it the KPB address, the return status value, and the procedure value of the kernel process spinlock restart routine.
7. Resumes the stalled kernel process.

System Routines

EXE\$KP_STALL_GENERAL

EXE\$KP_STALL_GENERAL

Stalls the execution of a kernel process.

Module

KERNEL_PROCESS_MAGIC

Macro

KP_STALL_GENERAL
KP_STALL_FORK
KP_STALL_FORK_WAIT
KP_STALL_IOFORK
KP_STALL_REQCHAN
KP_STALL_WFIKPCH
KP_STALL_WFIRLCH

Format

EXE\$KP_STALL_GENERAL kpb

Context

EXE\$KP_STALL_GENERAL conforms to the OpenVMS AXP calling standard and can only be called by a kernel process.

Arguments

kpb
VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of the caller's KPB.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$INSFARG	Not all of the required arguments were specified.
Other values	As supplied to EXE\$KP_RESTART

Description

EXESKP_STALL_GENERAL suspends execution of the current kernel process. It performs the following tasks:

- Validates the structure indicated by the **kpb** argument. If the structure is not a KPB, or if it is currently valid, active, or in the process of deletion, EXESKP_STALL_GENERAL requests an INCONSTATE bugcheck.
- Preserves the current context by saving the current kernel process stack pointer (SP) and the registers indicated by KPB\$IS_REG_MASK on the stack (which it quadword-aligns after obtaining the current SP). It saves the new value of the kernel process SP in KPB\$PS_STACK_SP.
- Restores the SP of the initiator of the kernel process thread from KPB\$PS_SAVED_SP and poisons that field.
- Restores the preserved registers (as indicated by KPB\$IS_REG_MASK) from the top of the initiator's stack, plus the original SP of the initiator of the kernel process thread.
- Marks the kernel process as inactive by clearing KPB\$V_ACTIVE in KPB\$IS_FLAGS.
- Calls the kernel process scheduling stall routine indicated by the procedure value in KPB\$PS_SCH_STALL_RTN, passing it the KPB address and the procedure value of the spin lock stall handling routine (from KPB\$PS_SPL_STALL_ROUTINE), or zero if the KPB spin lock area is not present. If there is no kernel process scheduling stall routine, EXESKP_STALL_GENERAL requests an INCONSTATE bugcheck.

OpenVMS provides the following jacket routines for EXESKP_STALL_GENERAL that supply scheduling stall routines for basic device driver functions:

Table 2–2 Kernel Process Stall Jacket Routines and Scheduling Stall Routines

Stall Jacket Routine	Scheduling Stall Routine ¹	Action of Stall Routine
EXESKP_FORK	STALL_FORK	Calls EXESPRIMITIVE_FORK on behalf of a kernel process. When it regains control from the OpenVMS fork dispatcher, this stall routine resumes the kernel process by calling EXESKP_RESTART.
EXESKP_FORK_WAIT	STALL_FORK_WAIT	Calls EXESPRIMITIVE_FORK_WAIT on behalf of a kernel process. When it regains control from the OpenVMS software timer interrupt service routine (which resumes the entries on the fork-and-wait queue), this stall routine resumes the kernel process by calling EXESKP_RESTART.

¹These scheduling stall routines are not globally accessible.

(continued on next page)

System Routines

EXE\$KP_STALL_GENERAL

Table 2–2 (Cont.) Kernel Process Stall Jacket Routines and Scheduling Stall Routines

Stall Jacket Routine	Scheduling Stall Routine ¹	Action of Stall Routine
EXE\$KP_IOFORK	STALL_FORK	Calls EXE\$PRIMITIVE_FORK (with timeouts disabled from the device unit associated with the KPB [UCB\$PS_UCB]) on behalf of a kernel process. When it regains control from the OpenVMS fork dispatcher, this stall routine resumes the kernel process by calling EXE\$KP_RESTART.
IOC\$KP_REQCHAN	STALL_REQCHAN	Calls EXE\$PRIMITIVE_REQCHAN on behalf of a kernel process. When it regains control after the channel has been granted, this stall routine resumes the kernel process by calling EXE\$KP_RESTART.
IOC\$KP_WFIKPCH	STALL_WFIXXCH	Issues the WFIKPCH macro on behalf of a kernel process. When it regains control due to a timeout or from interrupt servicing, this stall routine resumes the kernel process by calling EXE\$KP_RESTART, returning to it SSS_NORMAL or SSS_TIMEOUT status.
IOC\$KP_WFIRLCH	STALL_WFIXXCH	Issues the WFIRLCH macro on behalf of a kernel process. When it regains control due to a timeout or from interrupt servicing, it resumes the kernel process by calling EXE\$KP_RESTART, returning to it SSS_NORMAL or SSS_TIMEOUT status.

¹These scheduling stall routines are not globally accessible.

When the kernel process scheduling stall routine returns control, EXE\$KP_STALL_GENERAL returns SSS_NORMAL status to the initiator of the kernel process thread (that is, the caller if EXE\$KP_START or EXE\$KP_RESTART).

EXE\$KP_START

Starts the execution of a kernel process.

Module

KERNEL_PROCESS_MAGIC

Macro

KP_START
DDTAB (**start**=EXE\$KP_STARTIO)

Format

EXE\$KP_START kpb ,routine [,reg-mask]

Context

EXE\$KP_START conforms to the OpenVMS AXP calling standard. Its caller must be executing at IPL\$_RESCHED or above.

Neither the initiator of the kernel process thread nor the kernel process itself can assume that there is any relationship between them unless they mutually establish one. The initiator and the kernel process must establish explicit synchronization between themselves for operations that require it.

The kernel process cannot assume that its initiator is not running in parallel. Neither can it depend on inheriting the synchronization capabilities of its caller (for instance, its spin locks and IPL). The initiator of the kernel process thread cannot assume that the kernel process has already executed when EXE\$KP_START returns control.

Arguments

kpb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of KPB.

routine

VMS Usage: procedure_value
type: longword (unsigned)
access: read only
mechanism: by reference

Procedure value of the routine to be started as the top-level routine in the kernel process.

reg-mask

VMS Usage: mask_quadword
type: quadword (unsigned)
access: read only
mechanism: by value

System Routines

EXE\$KP_START

Optional register save mask, indicating which registers must be preserved across kernel process context switches. Registers R0, R1, R16 through R25, R28, R30, and R31 (KPREG\$K_ERR_REG_MASK) are never preserved across context switches; a **reg-mask** that indicates any of these registers is illegal. Registers R12 through R15, R26, R27, and R29 (KPREG\$K_MIN_REG_MASK) are always saved and need not be specified.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$BADPARAM	An illegal reg-mask was specified.
SS\$INSFARG	Not all of the required arguments were specified.

Description

EXE\$KP_START performs the following tasks to create a kernel process and start its execution:

1. Validates the structure indicated by the **kpb** argument. If the structure is not a KPB, or if it is currently valid, active, or in the process of deletion, EXE\$KP_START requests an INCONSTATE bugcheck.
2. Constructs the register save mask from the value specified in **reg-mask**, if present, and the minimal register save mask. EXE\$KP_START writes a value into this field that reflects the register save mask specified by its caller, plus a set of registers that are always preserved across such context switches (KPB\$K_MIN_REG_MASK), including R12 through R15, R27, and R29.
If an illegal **reg-mask** is specified, EXE\$KP_START returns SS\$BADPARAM status to its caller. Otherwise, EXE\$KP_START saves the register save mask in KPB\$IS_REG_MASK.
3. Preserves the current context by saving the current stack pointer (SP) and the registers indicated by KPB\$IS_REG_MASK on the stack (which it quadword-aligns after obtaining the current SP). It saves the new value of the SP in KPB\$PS_SAVED_SP.
4. Establishes kernel process context by loading the base of the kernel process stack (KPB\$PS_STACK_BASE) into the SP and KPB\$PS_STACK_SP.
5. Makes the KPB active and valid by setting the corresponding bits in KPB\$IS_FLAGS.
6. Initializes the bottom of the kernel process stack to enable implicit kernel process termination (by means of a call to EXE\$KP_END) if the top-level kernel process routine returns to EXE\$KP_START.
7. Calls the top-level kernel process routine, as indicated by the **routine** argument, passing to it the address of the KPB.

If the initiator of the kernel process thread and the kernel process must exchange additional parameters, they can do so only by using the KPB parameter area. The KPB parameter area is optionally created in the KPB by EXE\$KP_ALLOCATE_KPB.

8. When it regains control as the result of the kernel process invoking the KP_REQCOM macro, calls EXE\$KP_END.

EXE\$KP_STARTIO

Sets up and starts a kernel process to be used by a device driver.

Module

KERNEL_PROCESS_MIN, KERNEL_PROCESS_MON

Macro

DDTAB (**start**=EXE\$KP_STARTIO)

Format

JSB G^EXE\$KP_STARTIO

Context

The caller of EXE\$KP_STARTIO (usually IOC\$INITIATE) must be executing at fork IPL and hold the fork lock indicated by UCB\$B_FLCK. EXE\$KP_STARTIO returns to its caller in fork context with no explicit output values.

Input

Location	Contents
R0	Address of DDT
R3	Address of IRP
R5	Address of UCB
UCB\$L_BCNT	Number of bytes to be transferred
UCB\$L_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 direct-I/O transfer, virtual address of first page-table entry (PTE) of I/O-transfer buffer; for buffered-I/O transfer, address of buffer is system address space
DDT\$PS_KP_STARTIO	Procedure value of the driver's start-I/O routine, which serves as the top-level routine within the kernel process thread.
DDT\$IS_STACK_BCNT	Size in bytes of the kernel process stack
DDT\$IS_REG_MASK	Kernel process register save mask

Description

EXE\$KP_STARTIO uses information stored in the DDT to set up and start a kernel process that can be used by a device driver. It performs the following tasks:

1. Establishes the size of the kernel process stack as the minimum of DDT\$IS_STACK_BCNT and KPB\$K_MIN_IO_STACK (currently 8KB).

2. Issues a standard call to EXE\$KP_ALLOCATE_KPB to create the KPB and allocate the kernel process stack, passing to it the following:
 - Zero as the size of the KPB parameter area
 - KPB flags, indicating a VEST KPB with scheduling and spinlock areas, that is deallocated when the kernel process is terminated.
 - the kernel process stack size
 - IRP\$PS_KPB as the target location of the KPB address

If there were not enough free pages in the system for the kernel process stack, and the I/O request described by the IRP has not since been cancelled, EXE\$KP_STARTIO issues a fork-and-wait request. When EXESTIMEOUT resumes EXE\$KP_STARTIO, it retries the call to EXE\$KP_ALLOCATE_KPB.

If the attempt to allocated nonpaged pool for the KPB failed, EXE\$KP_STARTIO requests an INCONSTATE bugcheck.

3. Inserts the address of the IRP in KPB\$PS_IRP and the address of the UCB in KPB\$PS_UCB
4. Establishes the kernel process register save mask as the logical-OR of the registers specified in DDT\$IS_REG_MASK and those indicated by KPREG\$K_MIN_IO_REG_MASK (R2 through R5; the VAX AP, FP, SP, and PC [registers R12 through R15]; and R26, R27, and R29), minus those indicated by KPREG\$K_ERR_REG_MASK (R0 and R1; R16 through R25; R28; R30; and R31).
5. Issues a standard call to EXE\$KP_START, passing it the register save mask, the procedure value of a kernel process start-I/O routine (DDT\$PS_KP_STARTIO), and the address of the KPB.
6. Issues an RSB instruction to its caller (usually IOC\$INITIATE, or EXESTIMEOUT if EXE\$KP_STARTIO was resumed by fork-and-wait mechanism)

System Routines

EXE\$TIMEDWAIT_COMPLETE

EXE\$TIMEDWAIT_COMPLETE

Determines whether the time interval of a timed wait has concluded.

Module

[SYSLOA]TIMEDWAIT

Macro

TIMEDWAIT

Format

EXE\$TIMEDWAIT_COMPLETE end-value

Context

EXE\$TIMEDWAIT_COMPLETE conforms to the OpenVMS AXP calling standard.

Arguments

end-value

VMS Usage: aligned quadword
type: quadword (unsigned)
access: modify
mechanism: by reference

End time calculated by a previous call to EXE\$TIMEDWAIT_SETUP or EXE\$TIMEDWAIT_SETUP_10US.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$_CONTINUE	The timed wait has not yet completed. The time interval for the timed wait may or may not have expired. This is a success status.
SS\$_INSFARG	Not all of the required arguments were specified.
SS\$_TIMEOUT	The time interval for a timed wait has expired and the timed wait is complete.

Description

EXESTIMEDWAIT_COMPLETE compares the specified **end-value** (as computed by a prior call to EXESTIMEDWAIT_SETUP or EXESTIMEDWAIT_SETUP_10US) with an internal current-value. There are three results of this comparison:

- If the **end-value** is greater than or equal to the current-value value, the timed wait has not yet completed, and EXESTIMEDWAIT_COMPLETE returns SSS_CONTINUE status.
- If the **end-value** is less than the current-value, EXESTIMEDWAIT_COMPLETE sets the **end-value** to -1 and returns SSS_CONTINUE status.

When EXESTIMEDWAIT_COMPLETE returns SSS_CONTINUE status to the TIMEDWAIT macro, the macro reexecutes a specified series of instructions that tests for a particular exit condition. Having set the **end-value** to -1 prior to returning SSS_CONTINUE status, EXESTIMEDWAIT_COMPLETE allows for the possibility that the exit condition was actually met during the timed wait time interval, but after the embedded instruction series could detect it. This could be the case, for instance, if an interrupt occurred and was serviced after the instruction sequence was executed but before the call to EXESTIMEDWAIT_COMPLETE was made. As a result of this behavior, all timed wait instruction loops execute one additional time after the timed wait time interval has concluded.

- If the **end-value** is equal to -1, the timed wait has completed and EXESTIMEDWAIT_COMPLETE returns SSS_TIMEOUT status.

System Routines

EXE\$TIMEDWAIT_SETUP, EXE\$TIMEDWAIT_SETUP_10US

EXE\$TIMEDWAIT_SETUP, EXE\$TIMEDWAIT_SETUP_10US

Calculate and return the **end-value** used by EXE\$TIMEDWAIT_COMPLETE to determine when a timed wait has completed.

Module

[SYSLOA]TIMEDWAIT

Macro

TIMEDWAIT

Format

EXE\$TIMEDWAIT_SETUP delta-time ,end-value

EXE\$TIMEDWAIT_SETUP_10US delta-time ,end-value

Context

EXE\$TIMEDWAIT_SETUP and EXE\$TIMEDWAIT_SETUP_10US conform to the OpenVMS AXP calling standard.

Arguments

delta-time

VMS Usage: aligned quadword
type: quadword (unsigned)
access: read only
mechanism: by reference

Delay time specified in nanoseconds (for EXE\$TIMEDWAIT_SETUP) or 10-microsecond units (for EXE\$TIMEDWAIT_SETUP_10US)

end-value

VMS Usage: aligned quadword
type: quadword (unsigned)
access: write only
mechanism: by reference

End time token to be supplied as input to EXE\$TIMEDWAIT_COMPLETE.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

System Routines

EXE\$TIMEDWAIT_SETUP, EXE\$TIMEDWAIT_SETUP_10US

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$INSFARG	Not all of the required arguments were specified.

Description

EXE\$TIMEDWAIT_SETUP and EXE\$TIMEDWAIT_SETUP_10US compute the **end-value** that is supplied as an input argument to a subsequent call to EXE\$TIMEDWAIT_COMPLETE. EXE\$TIMEDWAIT_COMPLETE uses the **end-value** to determine whether the timed wait time interval has concluded.

EXE\$TIMEDWAIT_SETUP and EXE\$TIMEDWAIT_SETUP_10US generate a system-specific **end-value** from the sum of the specified **delta-time** and the current time, converted to a value that can be directly compared to an internal current-value. EXE\$TIMEDWAIT_SETUP_10US performs the additional step of converting the input **delta-time** to a number of nanoseconds.

System Routines

EXE_STD\$ABORTIO

EXE_STD\$ABORTIO

Completes the servicing of an I/O request without returning status to the I/O status block specified in the request.

Module

SYSQIOREQ

Format

status = EXE_STD\$ABORTIO (irp, pcb, ucb, qio_sts)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
qio_sts	integer	input	value	required

irp

I/O request packet. EXE_STD\$ABORTIO copies the **qio_sts** parameter to IRP\$SL_IOST1 and clears IRP\$PS_FDT_CONTEXT. The caller of EXE_STD\$ABORTIO should not access the IRP after the routine returns SS\$_FDT_COMPL status.

pcb

PCB of current process

ucb

Unit control block

qio_sts

Final status to be returned by the \$QIO system service to its caller. EXE_STD\$ABORTIO places this status in FDT_CONTEXT\$SL_QIO_STATUS. If you intend to access the FDT context structure after EXE_STD\$ABORTIO returns, you must obtain its address from IRP\$PS_FDT_CONTEXT and store it before making the call.

Return Values

SS\$_FDT_COMPL

Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

Contents of **qio_sts**
argument

Context

EXE_STD\$ABORTIO executes at its caller's IPL and raises to fork IPL, acquiring the associated fork lock in a multiprocessing environment. As a result, its caller cannot be executing above fork IPL. A driver usually transfers control to EXE_STD\$ABORTIO at IPL\$ASTDEL.

EXE_STD\$ABORTIO returns to its caller at the caller's IPL.

Description

The FDT completion routine EXE_STD\$ABORTIO terminates the servicing of an I/O request without returning status to the I/O status block specified in the original call to the \$QIO system service.

EXE_STD\$ABORTIO performs the following actions:

1. Examines the **qio_sts** argument. If the argument contains SSS_FDT_COMPL, EXE_STD\$ABORTIO returns to its caller. This check prevents an I/O request from being aborted more than once.
2. Places the status to be returned to the caller of the \$QIO system service in IRP\$L_IOST1 and in the FDT_CONTEXT structure.
3. Clears the pointer to the FDT_CONTEXT structure in IRP\$PS_FDT_CONTEXT.
4. Requests the fork lock, raising IPL to fork IPL, to perform the following tasks:
 - a. Clear IRP\$L_IOSB so that no status is returned by I/O postprocessing
 - b. Clear ACB\$V_QUOTA in IRP\$B_RMOD to prevent the delivery of any AST to the process specified in the I/O request
 - c. Update the count of available AST entries at PCB\$L_ASTCNT, if necessary
 - d. Insert the IRP in the local processor's I/O postprocessing queue. If the queue is empty, request a software interrupt from the local processor at IPL\$IOPOST.
5. Releases the fork lock, restoring the caller's IPL. The pending IPL\$IOPOST interrupt causes I/O postprocessing to occur before the remaining instructions in EXE_STD\$ABORTIO are executed.

When all I/O postprocessing has been completed, EXE_STD\$ABORTIO regains control and returns SSS_FDT_COMPL status to its caller.

Any ASTs specified when the I/O request was issued will not be delivered, and any event flags requested will not be set.

System Routines

EXE_STD\$ABORTIO

Macro

CALL_ABORTIO [do_ret=YES]

where:

do_ret indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

In a Step 2 driver, the CALL_ABORTIO macro simulates the JMP to EXE\$ABORTIO in the FDT routine of a Step 1 driver. It initializes the **irp**, **pcb**, **ucb**, and **qio_sts** parameters from the contents of R3, R4, R5, and R0, respectively, and calls EXE_STD\$ABORTIO. When EXE_STD\$ABORTIO returns control to the code generated by a default invocation of \$ABORTIO, a RET instruction returns control to the caller of \$ABORTIO's invoker. Status is returned in R0 and in the FDT_CONTEXT structure.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The order in which formal parameters are passed to EXE_STD\$ABORTIO differs from the order in which they are provided in registers to the Step 1 routine EXE\$ABORTIO.
- The contents of R0 (final \$QIO system service status in a call to EXE\$ABORTIO in a Step 1 driver) are destroyed across the call to EXE_STD\$ABORTIO. This is especially important if you use the \$ABORTIO macro on OpenVMS AXP systems and expect R0 to retain its value afterwards.
- Unlike EXE\$ABORTIO, EXE_STD\$ABORTIO does not lower IPL to 0 before exiting. EXE_STD\$ABORTIO returns to its caller at the caller's IPL.
- EXE\$ABORTIO returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. EXE_STD\$ABORTIO returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

EXE_STD\$ALLOCBUF, EXE_STD\$ALLOCIRP

Allocates a buffer from nonpaged pool for a buffered-I/O operation.

Module

MEMORYALC

Format

status = EXE_STD\$ALLOCBUF (reqsize, blocksize, blockptr)

status = EXE_STD\$ALLOCIRP (blocksize, blockptr)

Arguments

Argument	Type	Access	Mechanism	Status
reqsize	integer	input	value	required
alozsize_p	pointer	output	value	required
bufptr_p	pointer	output	value	required

reqsize

Size of requested buffer in bytes (EXE_STD\$ALLOCBUF only). This value should include the 12 bytes required to store header information.

alozsize_p

Location in which EXE_STD\$ALLOCBUF and EXE_STD\$ALLOCIRP write the size of the requested buffer in bytes.

bufptr_p

Location in which EXE_STD\$ALLOCBUF and EXE_STD\$ALLOCIRP write the address of allocated buffer. The following fields are initialized in the buffer:

Field	Contents
IRP\$W_SIZE (in allocated buffer)	Size of requested buffer in bytes (for EXE_STD\$ALLOCBUF), IRP\$C_LENGTH (for EXE_STD\$ALLOCIRP).
IRP\$B_TYPE (in allocated buffer)	DYN\$C_BUFIO (for EXE_STD\$ALLOCBUF), DYN\$C_IRP (for EXE_STD\$ALLOCIRP).

Return Values

SS\$NORMAL

Normal, successful completion.

SS\$INSFMEM

Insufficient memory to satisfy request.

System Routines

EXE_STD\$ALLOCBUF, EXE_STD\$ALLOCIRP

Context

EXE_STD\$ALLOCBUF and EXE_STD\$ALLOCIRP set IPL to IPL\$ASTDEL. As a result they cannot be called by code executing above IPL\$ASTDEL. They return control to the caller at IPL\$ASTDEL.

Description

EXE_STD\$ALLOCBUF attempts to allocate a buffer of the requested size from nonpaged pool; EXE_STD\$ALLOCIRP attempts to allocate an IRP from nonpaged pool.

If sufficient memory is not available, EXE_STD\$ALLOCBUF and EXE_STD\$ALLOCIRP examine the PCB (CTL\$GL_PCB) to determine whether the process has resource wait mode enabled. If PCB\$V_SSRWAIT in PCB\$_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\$_BYTCNT or JIB\$_BYTLM, or both) by calling EXE\$DEBIT_BYTCNT or EXE\$DEBIT_BYTCNT_BYTLM.

Note

You can perform this task and allocate a buffer of the requested size by using the routines EXE\$DEBIT_BYTCNT_ALO and EXE\$DEBIT_BYTCNT_BYTLM_ALO. These routines invoke EXE_STD\$ALLOCBUF.)

The normal buffered I/O postprocessing routine (IOC_STD\$REQCOM), initiated by the REQCOM macro, readjusts quotas and also deallocates the buffer.

Note

The value returned in the **alospace_p** argument and placed at IRP\$W_SIZE in the allocated buffer is the size of the allocated buffer. The actual size of the 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.

Macro

CALL_ALLOCBUF
CALL_ALLOCIRP

In a Step 2 driver, CALL_ALLOCBUF and CALL_ALLOCIRP simulate a JSB to EXE\$ALLOCBUF and EXE\$ALLOCIRP, respectively. CALL_ALLOCBUF calls EXE_STD\$ALLOCBUF using the current contents of R1 as the **reqsize** argument. Both CALL_ALLOCBUF and CALL_ALLOCIRP return status in R0, the address of the allocated buffer in R2 and its size in R1. If a resource wait occurred, these macros return the address of the PCB in R4.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$ALLOCBUF and EXE_STD\$ALLOCIRP replace EXE\$ALLOCBUF and EXE\$ALLOCIRP (used by Step 1 drivers and OpenVMS VAX drivers). Unlike the Step 1 routines, the Step 2 routines do not preserve the original contents of R4, or return the address of the PCB in R4 if a wait has occurred.

EXE_STD\$ALTQUEPKT

Delivers an IRP to a driver's alternate start-I/O routine without regard for the status of the device.

Module

SYSQIOREQ

Format

EXE_STD\$ALTQUEPKT (irp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

irp
I/O request packet.

ucb
Unit control block.

EXE_STD\$ALTQUEPKT reads the following UCB fields:

Field	Contents
UCB\$B_FLCK	Fork lock index
UCB\$L_DDT	Address of unit's DDT. EXE_STD\$ALTQUEPKT reads DDB\$PS_ALTSTART to obtain the procedure value of the driver's alternate start-I/O routine.
UCB\$L_ALTIOWQ	Address of the alternate start-I/O wait queue listhead.

Context

A driver FDT routine typically calls EXE_STD\$ALTQUEPKT at IPL\$ASTDEL. EXE_STD\$ALTQUEPKT raises to fork IPL (acquiring the associated fork lock) before calling the driver's alternate start-I/O routine. When the alternate start-I/O routine returns control to it, EXE_STD\$ALTQUEPKT returns control to its caller at the caller's IPL (having released its acquisition of the fork lock).

Description

EXE_STD\$ALTQUEPKT calls the driver's alternate start-I/O routine. It does not test whether the unit is busy before making the call.

Macro

CALL_ALTQUEPKT

In a Step 2 driver, the CALL_ALTQUEPKT macro simulates the JSB to EXE\$ALTQUEPKT in the FDT routine of a Step 1 driver. CALL_ALTQUEPKT calls EXE_STD\$ALTQUEPKT, using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$ALTQUEPKT replaces EXE\$ALTQUEPKT (used by Step 1 drivers and OpenVMS VAX drivers).

System Routines

EXE_STD\$CARRIAGE

EXE_STD\$CARRIAGE

Interprets the carriage control specifier in IRP\$B_CARCON and converts it to a generic prefix or suffix format.

Module

SYSQIOFDT

Format

EXE_STD\$CARRIAGE (irp)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required

irp
I/O request packet.

Context

A driver FDT routine calls EXE_STD\$CARRIAGE at IPL\$_ASTDEL. EXE_STD\$CARRIAGE returns control to the driver at that IPL.

Description

For Digital internal use only.

Macro

CALL_CARRIAGE

In a Step 2 driver, the CALL_CARRIAGE macro simulates the JSB to EXE\$CARRIAGE in the FDT routine of a Step 1 driver. CALL_CARRIAGE calls EXE_STD\$CARRIAGE, using the current contents of R3 as the **irp** arguments.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$CARRIAGE replaces EXE\$CARRIAGE (used by Step 1 drivers and OpenVMS VAX drivers).

EXE_STD\$CHKxxxACCES

Checks logical (EXE_STD\$CHKLOGACCES), physical (EXE_STD\$CHKPHYACCES), read (EXE_STD\$CHKRDACCES), write (EXE_STD\$CHKWRTACCES), execute (EXE_STD\$CHKEXEACCES), create (EXE_STD\$CHKCREACCES), or delete (EXE_STD\$CHKDELACCES) I/O function access, based on the specified protection information.

Module

EXSUBROUT

Format

status = EXE_STD\$CHKCREACCES (arb, orb, pcb, ucb)
 status = EXE_STD\$CHKDELACCES (arb, orb, pcb, ucb)
 status = EXE_STD\$CHKEXEACCES (arb, orb, pcb, ucb)
 status = EXE_STD\$CHKLOGACCES (arb, orb, pcb, ucb)
 status = EXE_STD\$CHKPHYACCES (arb, orb, pcb, ucb)
 status = EXE_STD\$CHKRDACCES (arb, orb, pcb, ucb)
 status = EXE_STD\$CHKWRTACCES (arb, orb, pcb, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
arb	ARB	input	reference	required
orb	ORB	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required

arb
Agent rights block.

orb
Object rights block.

pcb
Process control block of accessor.

ucb
Unit control block of accessed object.

System Routines

EXE_STD\$CHKxxxACCES

Return Values

SS\$NORMAL	Specified access allowed.
SS\$NOPRIV	Specified access denied.

Context

A driver FDT routine calls EXE_STD\$CHKPHYACCES, EXE_STD\$CHKLOGACCES, EXE_STD\$CHKWRTACCES, EXE_STD\$CHKRDACCES, EXE_STD\$CHKCREACCES, EXE_STD\$CHKEXEACCES, and EXE_STD\$CHKDELACCES, at IPL\$ASTDEL. These routines return control to the driver at that IPL.

Description

For Digital internal use only.

Macro

```
CALL_CHKCREACCES [save_r1]
CALL_CHKDELACCES [save_r1]
CALL_CHKEXEACCES [save_r1]
CALL_CHKLOGACCES [save_r1]
CALL_CHKPHYACCES [save_r1]
CALL_CHKRDACCES [save_r1]
CALL_CHKWRTACCES [save_r1]
```

where:

save_r1 indicates that the macro must preserve the contents of R1 across the call to EXE_STD\$CHKPHYACCES, EXE_STD\$CHKLOGACCES, EXE_STD\$CHKWRTACCES, EXE_STD\$CHKEXEACCES, EXE_STD\$CHKCREACCES, EXE_STD\$CHKDELACCES or EXE_STD\$CHKRDACCES. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

In a Step 2 driver, the CALL_CHKCREACCES, CALL_CHKDELACCES, CALL_CHKEXEACCES, CALL_CHKLOGACCES, CALL_CHKPHYACCES, CALL_CHKWRTACCES, and CALL_CHKRDACCES, macros simulate the JSB to EXE\$CHKCREACCES, EXE\$CHKDELACCES, EXE\$CHKEXEACCES, EXE\$CHKPHYACCES, EXE\$CHKLOGACCES, EXE\$CHKWRTACCES, or EXE\$CHKRDACCES in a Step 1 driver. Each macro calls the corresponding access-checking routine, using the current contents of R0, R1, R4, and R5 as the **arb**, **orb**, **pcb**, and **ucb** arguments. Unless you specify **save_r1=NO**, the macro preserves the quadword register R1 across the call. All macros return status in R0.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The OpenVMS AXP I/O function access checking routines replace their Step 1 and OpenVMS VAX counterparts, but do not does not preserve R1 across a call.

EXE_STD\$FINISHIO

Completes the servicing of an I/O request and returns status to the I/O status block specified in the original call to the \$QIO system service.

Module

SYSQIOREQ

Format

status = EXE_STD\$FINISHIO (irp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

irp

I/O request packet. EXE_STD\$FINISHIO clears IRP\$PS_FDT_CONTEXT. The caller of EXE_STD\$FINISHIO should not access the IRP after the routine returns SSS_FDT_COMPL status.

ucb

Unit control block. EXE_STD\$FINISHIO increases UCB\$L_OPCNT.

Return Values

SSS_FDT_COMPL

Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SSS_NORMAL

The routine completed successfully.

Context

EXE_STD\$FINISHIO executes at its caller's IPL and raises to fork IPL, acquiring the associated fork lock in a multiprocessing environment. As a result, its caller cannot be executing above fork IPL. A driver usually transfers control to EXE_STD\$FINISHIO at IPL\$ASTDEL.

EXE_STD\$FINISHIO returns to its caller at the caller's IPL.

System Routines

EXE_STD\$FINISHIO

Description

The FDT completion routine EXE_STD\$FINISHIO completes the servicing of an I/O request and returns status to the I/O status block specified in the original call to the \$QIO system service. It performs the following actions:

1. Clears the pointer to the FDT context structure in IRP\$PS_FDT_CONTEXT.
2. Requests the fork lock, raising IPL to fork IPL, to perform the following tasks:
 - a. Increase the number of I/O operations completed on the current device in the operation count field of the UCB (UCB\$SL_OPCNT). This task is performed at fork IPL, holding the associated fork lock in a multiprocessing environment.
 - b. Insert the IRP in the local processor's I/O postprocessing queue. If the queue is empty, request a software interrupt from the local processor at IPL\$IOPOST.
3. Releases the fork lock, restoring the caller's IPL. The pending IPL\$IOPOST interrupt causes I/O postprocessing to occur before the remaining instructions in EXE_STD\$FINISHIO are executed.

When all I/O postprocessing has been completed, EXE_STD\$FINISHIO regains control and returns SS\$_FDT_COMPL status to its caller, passing SS\$_NORMAL as the final \$QIO completion status in the FDT_CONTEXT structure.

The image that requested the I/O operation receives SS\$_NORMAL status, indicating that the I/O request has completed without device-independent error.

Macro

```
CALL_FINISHIO [do_ret=YES]  
CALL_FINISHIOC [do_ret=YES]
```

where:

do_ret indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

In a Step 2 driver, the CALL_FINISHIO macro simulates the JMP to EXE\$FINISHIO in the FDT routine of a Step 1 driver. The CALL_FINISHIOC macro simulates the JMP to EXE\$FINISHIOC. The former macro moves the current contents of R0 and R1 into IRP\$SL_IOST1 and IRP\$SL_IOST2, respectively; the latter initializes IRP\$SL_IOST1 from R0 and clears IRP\$SL_IOST2. Both macros initialize the **irp** and **ucb** parameters from the contents of R3 and R5, respectively before calling EXE_STD\$FINISHIO. When EXE_STD\$FINISHIO returns control to the code generated by a default invocation of CALL_FINISHIO or CALL_FINISHIOC, a RET instruction returns control to the caller of the macro's invoker. Status is returned in R0 and in the FDT_CONTEXT structure.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- A Step 1 driver supplies the first and second longwords of device-specific status (in R0 and R1) as input to EXE\$FINISHIO. EXE\$FINISHIO writes these longwords to IRP\$L_IOST1 and IRP\$L_IOST2, respectively, from which I/O postprocessing transfers their values to the I/O status block specified in the original \$QIO call. These status longwords are not input parameters to EXE_STD\$FINISHIO. Rather, a Step 2 driver's FDT routine must fill in IRP\$L_IOST1 and IRP\$L_IOST2 before calling EXE_STD\$FINISHIO.

Because the OpenVMS VAX routines EXE\$FINISHIO and EXE\$FINISHIOC differ only in that the latter routine clears the second longword on I/O status, there is no Step 2 equivalent of EXE\$FINISHIOC. If the driver needs to clear the second I/O status longword, it simply does so before calling EXE_STD\$FINISHIO.

- The address of the PCB, supplied as input to EXE\$FINISHIO on OpenVMS VAX systems, is not provided as input to EXE_STD\$FINISHIO.
- Unlike EXE\$FINISHIO, EXE_STD\$FINISHIO does not lower IPL to 0 before exiting. EXE_STD\$FINISHIO returns to its caller at the caller's IPL.
- EXE\$FINISHIO returns control to the system service dispatcher, passing it the final \$QIO system service status (SS\$_NORMAL) in R0. EXE_STD\$FINISHIO returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status (SS\$_NORMAL) in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

EXE\$ILLIOFUNC

Aborts I/O preprocessing for an I/O function not supported a driver.

Module

SYSQIOFDT

Format

status = EXE\$ILLIOFUNC (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp

I/O request packet for the current I/O request

pcb

Process control block of the current process

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request

ccb

Channel control block that describes the process-I/O channel

Return Values

SS\$FDT_COMPL

Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Context

FDT dispatching code in the \$QIO system service calls EXE\$ILLIOFUNC at IPL\$ASTDEL when processing an I/O function that is not supported by a driver. EXE\$ILLIOFUNC returns to the system service dispatcher at IPL\$ASTDEL.

Description

Because any slot corresponding to an unsupported function in a driver's FDT action vector contains the procedure value of EXE\$ILLIOFUNC, FDT dispatching code in the \$QIO system service calls EXE\$ILLIOFUNC to process any I/O request specifying an unsupported I/O function code.

EXE\$ILLIOFUNC calls EXE_STD\$ABORTIO to terminate the processing of the I/O request.

Macro

CALL_INSERT_IRP

In a Step 2 driver, the CALL_INSERT_IRP macro simulates a JSB to EXE\$INSERT_IRP in a Step 1 driver. CALL_INSERT_IRP calls EXE_STD\$INSERT_IRP, using the current contents of R2 and R3 as the **irp_lh** and **irp** arguments, respectively. It returns status in R0.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$INSERT_IRP replaces EXE\$INSERT_IRP (used by Step 1 drivers) and EXE\$INSERTIRP (used by OpenVMS VAX drivers).

System Routines

EXE_STD\$INSIOQ, EXE_STD\$INSIOQC

EXE_STD\$INSIOQ, EXE_STD\$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

Format

EXE_STD\$INSIOQ (irp, ucb)

EXE_STD\$INSIOQC (irp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

irp
I/O request packet.

ucb
Unit control block.

EXE_STD\$INSIOQ and EXE_STD\$INSIOQC read the following UCB fields:

Field	Contents
UCB\$B_FLCK	Fork lock index
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\$L_QLEN	Length of pending-I/O queue

EXE_STD\$INSIOQ and EXE_STD\$INSIOQC write the following UCB fields:

Field	Contents
UCB\$L_STS	UCB\$V_BSY set
UCB\$W_QLEN	Increased

Context

EXE_STD\$INSIOQ and EXE_STD\$INSIOQC immediately raise to fork IPL and, in a 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_STD\$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_STD\$INSIOQC instead.

Description

EXE_STD\$INSIOQ and EXE_STD\$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.

EXE_STD\$INSIOQ and EXE_STD\$INSIOQC increase UCB\$L_QLEN and proceed according to the status of the device (as indicated by UCB\$V_BSY in UCB\$L_STS) as follows:

- If the device is busy, call EXE_STD\$INSERT_IRP to place the IRP on the device's pending-I/O queue.
- If the device is idle, call IOC_STD\$INITIATE to begin device processing of the I/O request immediately. IOC_STD\$INITIATE transfers control to the driver's start-I/O routine.

Macro

CALL_INSIOQ
CALL_INSIOQC

In a Step 2 driver, the CALL_INSIOQ and CALL_INSIOQC macros simulate a JSB to EXE\$INSIOQ and EXE\$INSIOQC, respectively, in a Step 1 driver. \$INSIOQ calls EXE_STD\$INSIOQ, and \$INSIOQC calls EXE_STD\$INSIOQC, using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively.

Notes for Converting Step 1 Drivers

None.

System Routines

EXE_STD\$IORSNWAIT

EXE_STD\$IORSNWAIT

Places a process in a resource wait state if it has enabled resource waits.

Module

SYSQIOFDT

Format

status = EXE_STD\$IORSNWAIT (irp, pcb, ucb, ccb, qio_sts, rsn)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required
qio_sts	input	integer	value	required
rsn	input	integer	value	required

irp

I/O request packet.

pcb

Process control block.

ucb

Unit control block.

ccb

Channel control block.

qio_sts

Final status to be returned by the \$QIO system service to its caller if the caller has not enabled resource wait mode. EXE_STD\$IORSNWAIT calls EXE_STD\$ABORTIO to place this status in FDT_CONTEXT\$QIO_STATUS. If you intend to access the FDT context structure after EXE_STD\$IORSNWAIT returns, you must obtain its address from IRP\$PS_FDT_CONTEXT and store it before making the call.

rsn

Number of the resource for which the request is waiting.

Return Values

SS\$FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.
---------------	---

Status in FDT_CONTEXT

Contents of qio_sts argument	Process has not enabled resource waits.
SS\$WAIT_CALLERS_MODE	Process has been placed in a resource wait state.

Context

EXE_STD\$IORSNWAIT is called by, and returns to, a driver's FDT routine at IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

CALL_IORSNWAIT [do_ret=YES]

where:

do_ret indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

In a Step 2 driver, the CALL_IORSNWAIT macro simulates the JMP to EXE\$IORSNWAIT in the FDT routine of a Step 1 driver. It calls EXE_STD\$IORSNWAIT using the current contents of R3, R4, R5, R6, R0, and R1 as the **irp**, **pcb**, **ucb**, **ccb**, **qio_sts**, and **rsn** arguments, respectively. When EXE_STD\$IORSNWAIT returns control to the code generated by a default invocation of CALL_IORSNWAIT, a RET instruction returns control to the caller of the macro's invoker. Status is returned in R0 and in the FDT_CONTEXT structure.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$IORSNWAIT replaces EXE\$IORSNWAIT (used by Step 1 drivers and OpenVMS VAX drivers). The order in which formal parameters are passed to EXE_STD\$IORSNWAIT differs from the order in which they are provided in registers to the Step 1 routine EXE\$IORSNWAIT.
- EXE\$IORSNWAIT returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. EXE_STD\$IORSNWAIT returns to its caller, passing it SS\$FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT context structure. The \$QIO system service retrieves the status from this structure.

System Routines

EXE_STD\$LCLDSKVALID

EXE_STD\$LCLDSKVALID

Processes I/O functions that affect the online count and local valid status of a disk.

Module

SYSQIOFDT

Format

status = EXE_STD\$LCLDSKVALID (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp

I/O request packet for the current I/O request. The I/O function for the current request is available in IRP\$L_FUNC.

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

EXE_STD\$LCLDSKVALID reads the following UCB fields.

Field	Contents
UCB\$B_FLCK	Fork lock index
UCB\$L_STS	UCB\$V_LCL_VALID set if the volume is valid; clear if the drive is unloaded or available
UCB\$B_ONLCNT	Number of hosts that have set this disk on line

EXE_STD\$LCLDSKVALID writes the following UCB fields:

Field	Contents
UCB\$L_STS	UCB\$V_LCL_VALID set if the requested function is IOS_PACKACK; cleared if the requested function is IOS_UNLOAD or IOS_AVAILABLE

Field	Contents
UCB\$B_ONLCNT	Incremented if UCB\$V_LCL_VALID is not set and the requested function is IOS_PACKACK; decremented if UCB\$V_LCL_VALID is set and the requested function is IOS_UNLOAD or IOS_AVAILABLE

ccb

Channel control block that describes the process-I/O channel

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$NORMAL The routine completed successfully.

Context

FDT dispatching code calls EXE_STD\$LCLDSKVALID at IPL\$ASTDEL. EXE_STD\$LCLDSKVALID immediately raises IPL to fork IPL, requesting the associated fork lock in a multiprocessing environment. When it regains control from EXE_STD\$QIODRVPKT or EXE_STD\$FINISHIO, EXE_STD\$LCLDSKVALID lowers IPL to IPL\$ASTDEL and relinquishes the fork lock before returning to the system service dispatcher.

Description

A disk driver specifies the system-supplied upper-level FDT action routine EXE_STD\$LCLDSKVALID in an FDT_ACT macro invocation to service a request for an IOS_PACKACK, IOS_AVAILABLE, or IOS_UNLOAD function for a local disk. The actions of EXE_STD\$LCLDSKVALID depend on the I/O function indicated by R7 and the value of UCB\$V_LCL_VALID in UCB\$S_STS.

For an IOS_PACKACK function, EXE_STD\$LCLDSKVALID proceeds as follows:

- If UCB\$V_LCL_VALID is clear:
 - Sets UCB\$V_LCL_VALID.
 - Increases UCB\$B_ONLCNT.
 - If this is the first cluster pack acknowledgment on the disk (that is, if UCB\$B_ONLCNT equals 1), invokes the \$QIODRVPKT macro to deliver the IRP to the driver's start-I/O routine. EXE_STD\$LCLDSKVALID regains control with SS\$FDT_COMPL status in R0 and a final \$QIO system service status of SS\$NORMAL in the FDT_CONTEXT structure.
- If UCB\$V_LCL_VALID is set, EXE_STD\$LCLDSKVALID requests that the FDT completion routine EXE_STD\$FINISHIO complete the I/O request. EXE_STD\$FINISHIO returns to EXE_STD\$LCLDSKVALID with SS\$FDT_COMPL status in R0 and a final \$QIO system service status of SS\$NORMAL in the FDT_CONTEXT structure.

System Routines

EXE_STD\$LCLDSKVALID

For an IO\$_UNLOAD or IO\$_AVAILABLE function, EXE_STD\$LCLDSKVALID proceeds as follows:

- If UCBSV_LCL_VALID is set:
 - Clears UCBSV_LCL_VALID
 - Decreases UCB\$B_ONLCNT
 - If this is the last cluster unload or available request, invokes the \$QIODRVPKT macro to deliver the IRP to the driver's start-I/O routine. EXE_STD\$LCLDSKVALID regains control with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure.
- If UCBSV_LCL_VALID is clear, EXE_STD\$LCLDSKVALID requests that the FDT completion routine EXE_STD\$FINISHIO complete the I/O request. EXE_STD\$FINISHIO returns to EXE_STD\$LCLDSKVALID with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure.

A driver must define the local disk UCB extension to use this routine.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The upper-level FDT routine EXE\$LCLDSKVALID (used by OpenVMS VAX and Step 1 OpenVMS AXP device drivers) expects as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table (FDT) from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$LCLDSKVALID.

EXE_STD\$MNTVERSIO

Initiates a mount verification I/O request to a device.

Module

MOUNTVER

Format

EXE_STD\$MNTVERSIO (rout, irp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
rout	procedure value	input	value	required
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

rout

Procedure value of action routine to postprocess the mount verification I/O request.

irp

I/O request packet.

ucb

Unit control block.

Context

EXE_STD\$MNTVERSIO raises IPL to fork IPL, obtaining the corresponding fork lock in an OpenVMS multiprocessing system. It releases the fork lock and returns control to its caller at its caller's IPL.

Description

For Digital internal use only.

Macro

CALL_MNTVERSIO

In a Step 2 driver, the CALL_MNTVERSIO macro simulates a JSB to EXE\$MNTVERSIO in a Step 1 driver. CALL_MNTVERSIO calls EXE_STD\$MNTVERSIO, using the current contents of R0, R3, and R5 as the **rout**, **irp**, and **ucb** arguments, respectively.

System Routines

EXE_STD\$MNTVERSIO

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$MNTVERSIO replaces EXE\$MNTVERSIO (used by Step 1 and OpenVMS VAX drivers).

EXE_STD\$MODIFY

Translates a logical read/write function into a physical read/write function, transfers \$QIO system service parameters to the IRP, validates and prepares a user buffer, and aborts the request or proceeds with a direct-I/O, DMA read/write operation.

Module

SYSQIOFDT

Format

status = EXE_STD\$MODIFY (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp

I/O request packet for the current I/O request.

EXE_STD\$MODIFY reads the following IRP fields:

Field	Contents
IRP\$QIO_P1	\$QIO system service p1 argument, containing the buffer's virtual address.
IRP\$QIO_P2	\$QIO system service p2 argument, containing the number of bytes in transfer. The maximum number of bytes that EXE_STD\$MODIFY can transfer is 65,535 (128 pages minus one byte).
IRP\$QIO_P4	\$QIO system service p4 argument, containing the carriage control byte.
IRP\$FUNC	I/O function code.
IRP\$RMOD	Access mode of the caller of the \$QIO system service.

EXE_STD\$MODIFY writes the following IRP fields:

Field	Contents
IRP\$CARCON	Carriage control byte (from IRP\$QIO_P4)
IRP\$FUNC	Logical read/write function code converted to physical

System Routines

EXE_STD\$MODIFY

Field	Contents
IRP\$L_STS	IRP\$V_FUNC set to indicate read function
IRP\$L_SVAPTE	System virtual address of the PTE that maps the first page of the buffer
IRP\$L_BOFF	Byte offset to start of transfer in page
IRP\$L_OBOFF	Original byte offset into the first page of a segmented direct-I/O transfer
IRP\$L_BCNT	Size of transfer in bytes

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb

Channel control block that describes the process-I/O channel

Return Values

SS\$FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.
---------------	---

Status in FDT_CONTEXT

SS\$ACCVIO	Buffer specified in buffer parameter does not allow read access.
SS\$BADPARAM	size parameter is less than zero.
SS\$INSFWSL	Insufficient working set limit.
SS\$NORMAL	The I/O request has been successfully queued.
SS\$QIO_CROCK	Buffer page must be faulted into memory.

Context

FDT dispatching code in the \$QIO system service calls EXE_STD\$MODIFY as an upper-level FDT action routine at IPL\$ASTDEL.

Description

A driver specifies the system-supplied upper-level FDT action routine EXE_STD\$MODIFY to prepare a direct-I/O read/write request. A driver cannot specify EXE_STD\$MODIFY for buffered-I/O functions. Drivers that process functions that require an intermediate system buffer typically supply their own upper-level FDT action routines to handle them.

EXE_STD\$MODIFY performs the following functions:

- Sets IRP\$V_FUNC in IRP\$L_STS to indicate a read function
- Copies the **p4** argument of the \$QIO request from IRP\$L_QIO_P4 to IRP\$B_CARCON

- Translates a logical read/write function to a physical read/write function and stores the new function code in IRP\$L_FUNC.
- Examines the size of the transfer, as specified in the **p2** argument of the \$QIO request (IRP\$L_QIO_P2), and takes one of the following actions:
 - If the transfer byte count is zero, EXE_STD\$MODIFY invokes the \$QIODRVPKT macro to deliver the IRP to the driver's start-I/O routine. EXE_STD\$MODIFY regains control with S\$\$_FDT_COMPL status in R0 and a final \$QIO system service status of S\$\$_NORMAL in the FDT_CONTEXT structure. It returns to the \$QIO system service, passing these status values.

The driver start-I/O routine should check for zero-length buffers to avoid mapping to adapter node space. An attempted mapping can cause a system failure.
 - If the byte count is not zero, EXE_STD\$MODIFY calls EXE_STD\$MODIFYLOCK, passing 0 as the value of the **err_rout** argument.

EXE_STD\$MODIFYLOCK disables an optimization in MMG_STD\$IOLOCK and joins the code for EXE_STD\$READLOCK. EXE_STD\$MODIFYLOCK invokes the \$READCHK macro, which calls EXE_STD\$READCHK.

EXE_STD\$READCHK performs the following actions:

- Moves the transfer byte count (**size** parameter) into IRP\$L_BCNT.

If the byte count is negative, it calls EXE_STD\$ABORTIO, passing it a **qio_sts** of S\$\$_BADPARAM. When it regains control, EXE_STD\$READCHK returns to EXE_STD\$MODIFYLOCK with S\$\$_BADPARAM status in the FDT_CONTEXT structure and S\$\$_FDT_COMPL status in R0. EXE_STD\$MODIFYLOCK immediately returns to EXE_STD\$MODIFY, passing these status values. EXE_STD\$MODIFY, in turn, returns to the \$QIO system service.
- 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 returns S\$\$_NORMAL in R0 to EXE_STD\$MODIFYLOCK.
 - If the buffer does not allow write access, EXE_STD\$READCHK calls EXE_STD\$ABORTIO, passing it a **qio_sts** of S\$\$_ACCVIO. When it regains control, EXE_STD\$READCHK returns to EXE_STD\$MODIFYLOCK with S\$\$_ACCVIO status in the FDT_CONTEXT structure and S\$\$_FDT_COMPL status in R0. EXE_STD\$MODIFYLOCK immediately returns to EXE_STD\$MODIFY, passing these status values. EXE_STD\$MODIFY returns to the \$QIO system service.

If EXE_STD\$READCHK succeeds, EXE_STD\$MODIFYLOCK moves into IRP\$L_BOFF and IRP\$L_OBOFF the byte offset to the start of the buffer and calls MMG_STD\$IOLOCK.

MMG_STD\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:

- If MMG_STD\$IOLOCK succeeds, EXE_STD\$MODIFYLOCK stores in IRP\$L_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns S\$\$_NORMAL status in R0 to EXE_STD\$MODIFYLOCK. EXE_STD\$MODIFYLOCK returns immediately to EXE_STD\$MODIFY, passing to it this status value.

System Routines

EXE_STD\$MODIFY

EXE_STD\$MODIFY invokes the \$QIODRVPKT macro to deliver the IRP to the driver's start-I/O routine. EXE_STD\$MODIFY regains control with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure. It returns to the \$QIO system service, passing these status values.

- If MMG_STD\$IOLOCK fails, it returns SSS_ACCVIO, SSS_INSFWSL, or page fault status to EXE_STD\$MODIFYLOCK.

For SSS_ACCVIO and SSS_INSFWSL status, EXE_STD\$MODIFYLOCK calls EXE_STD\$ABORTIO, passing it one of these status values as a **qio_sts** argument. When it regains control, EXE_STD\$MODIFYLOCK returns EXE_STD\$MODIFY the specified status value in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$MODIFY returns to the \$QIO system service.

For page fault status, EXE_STD\$MODIFYLOCK sets the final \$QIO status in the FDT_CONTEXT structure to SSS_QIO_CROCK and initializes FDT_CONTEXT\$SL_QIO_R1_VALUE to the virtual address to be faulted. It then adjusts the 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.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine EXE\$MODIFY (used by OpenVMS VAX and Step 1 OpenVMS AXP device drivers) expects as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$MODIFY.
- EXE\$MODIFY returns control to the system service dispatcher, passing it the final \$QIO system service status (SSS_NORMAL, SSS_ACCVIO, or SSS_BADPARAM, or SSS_INSFWSL) in R0. EXE_STD\$MODIFY returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

EXE_STD\$MODIFYLOCK

Validates and prepares a user buffer for a direct-I/O, DMA read/write operation.

Module

SYSQIOFDT

Format

status = EXE_STD\$MODIFYLOCK (irp, pcb, ucb, ccb, buf, bufsiz, err_rout)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required
buf	address	input	reference	required
bufsiz	integer	input	value	required
err_rout	procedure value	input	value	required

irp

I/O request packet for the current I/O request.

EXE_STD\$MODIFYLOCK reads IRPSB_RMOD to determine the access mode of the caller of the \$QIO system service.

EXE_STD\$MODIFYLOCK writes the following IRP fields:

Field	Contents
IRP\$S_SVAPTE	System virtual address of the PTE that maps the first page of the buffer
IRP\$S_BOFF	Byte offset to start of transfer in page
IRP\$S_OBOFF	Original byte offset into the first page of a segmented direct-I/O transfer
IRP\$S_BCNT	Size of transfer in bytes

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

System Routines

EXE_STD\$MODIFYLOCK

ccb

Channel control block that describes the process-I/O channel.

buf

Virtual address of buffer.

bufsiz

Number of bytes in transfer.

err_rout

Procedure value of error-handling callback routine, or 0 if the driver does not process errors.

A driver typically specifies an error-handling callback routine when the driver must lock multiple areas into memory for a single I/O request and regain control to unlock these areas, if the request is to be aborted. The routine performs those tasks required before the request is backed out of or aborted. Such operations could include calling MMG_STD\$UNLOCK to release previous buffers participating in the I/O operation. The error-handling routine must preserve R0 and R1 and return back to EXE_STD\$MODIFYLOCK.

Chapter 1 describes the error-handling callback routine interface.

Return Values

SS\$NORMAL	The buffer is read-accessible and has been locked in memory.
SS\$FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$ACCVIO	Buffer specified in buf parameter does not allow read access.
SS\$BADPARAM	bufsiz parameter is less than zero.
SS\$INSFWSL	Insufficient working set limit.
SS\$NORMAL	Nothing has occurred yet to prevent the I/O request from being successfully queued. This is the initial value of the status field in an FDT_CONTEXT structure.
SS\$INSFWSL	Insufficient working set limit.
SS\$QIO_CROCK	Buffer page must be faulted into memory.

Context

The system-supplied upper-level FDT action routine EXE_STD\$MODIFY, or a driver-specific upper-level FDT action routine, calls EXE_STD\$MODIFYLOCK at IPL\$ASTDEL.

Description

A driver FDT routine calls the system-supplied FDT support routine EXE_STD\$MODIFYLOCK to check the read accessibility of an I/O buffer supplied in a SQIO request for a read/write function, and lock the buffer in memory in preparation for a DMA read/write operation.

A driver cannot specify EXE_STD\$MODIFY for buffered-I/O functions. Drivers that process functions that require an intermediate system buffer typically supply their FDT routines to handle them.

EXE_STD\$MODIFYLOCK disables an optimization in MMG_STD\$IOLOCK and joins the code for EXE_STD\$READLOCK. EXE_STD\$MODIFYLOCK invokes the \$READCHK macro, which calls EXE_STD\$READCHK.

EXE_STD\$READCHK performs the following actions:

- Moves the transfer byte count (**bufsiz** parameter) into IRP\$LCB_CNT.
If the byte count is negative, EXE_STD\$READCHK returns SSS_BADPARAM status to EXE_STD\$MODIFYLOCK.
- 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_STD\$READCHK sets IRP\$V_FUNC in IRP\$STS and returns SSS_NORMAL in R0 to EXE_STD\$MODIFYLOCK.
 - If the buffer does not allow write access, EXE_STD\$READCHK returns SSS_ACCVIO status to EXE_STD\$MODIFYLOCK.

If error status (SSS_BADPARAM or SSS_ACCVIO) is returned, EXE_STD\$MODIFYLOCK immediately calls the specified error-handling callback routine, passing to it the IRP, PCB, UCB, CCB, and status value. The callback routine must preserve R0 and R1 and return control to EXE_STD\$MODIFYLOCK. When the callback routine returns (or if no callback routine is specified), EXE_STD\$MODIFYLOCK calls EXE_STD\$ABORTIO, passing it the error status as **qio_sts**. EXE_STD\$ABORTIO returns to EXE_STD\$MODIFYLOCK with the error status in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$MODIFYLOCK immediately returns to its caller, passing these status values.

If SSS_NORMAL status is returned, EXE_STD\$MODIFYLOCK moves into IRP\$BOFF and IRP\$OBOFF the byte offset to the start of the buffer and calls MMG_STD\$IOLOCK.

MMG_STD\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:

- If MMG_STD\$IOLOCK succeeds, EXE_STD\$MODIFYLOCK stores in IRP\$SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns SSS_NORMAL status in R0 to EXE_STD\$MODIFYLOCK. EXE_STD\$MODIFYLOCK returns immediately to its caller, passing to it this status value.
- If MMG_STD\$IOLOCK fails, it returns SSS_ACCVIO, SSS_INSFWSL, or page fault status to EXE_STD\$MODIFYLOCK. EXE_STD\$MODIFYLOCK immediately calls the specified error-handling callback routine, passing to it the IRP, PCB, UCB, CCB, and status value. The callback routine must preserve R0 and R1 and return control to EXE_STD\$MODIFYLOCK. When

System Routines

EXE_STD\$MODIFYLOCK

the callback routine returns (or if no callback routine is specified), EXE_STD\$MODIFYLOCK proceeds as follows:

- For SSS_ACCVIO and SSS_INSFWSL status, EXE_STD\$MODIFYLOCK calls EXE_STD\$ABORTIO, passing it one of these status values as a **qio_sts** argument. When it regains control, EXE_STD\$MODIFYLOCK returns to its caller the specified status value in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0.

For page fault status, EXE_STD\$MODIFYLOCK sets the final \$QIO status in the FDT_CONTEXT structure to SSS_QIO_CROCK and initializes FDT_CONTEXT\$SL_QIO_R1_VALUE to the virtual address to be faulted. It then adjusts the 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.

The caller of EXE_STD\$MODIFYLOCK must examine the status in R0:

- If the status is SSS_NORMAL, the buffer is write accessible and has been successfully locked into memory and the starting virtual address of the page table entries that map the buffer is available in IRP\$SL_SVAPTE.
- If the status is SSS_FDT_COMPL, an error has occurred that has caused the I/O request to be aborted. You can determine the reason for the failure from FDT_CONTEXT\$SL_QIO_STATUS. Ordinarily a driver specifies an error-handling callback routine to process such errors.

Note that a driver cannot access the IRP once it has received SSS_FDT_COMPL status. If you know you need access to information stored in the IRP to back out an I/O request that has been aborted, you must store that information elsewhere prior to calling EXE_STD\$MODIFYLOCK.

Macro

```
CALL_MODIFYLOCK  
CALL_MODIFYLOCK_ERR [interface_warning=YES]
```

where:

interface_warning=YES, the default, specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. **interface_warning=NO** suppresses the warning.

In a Step 2 driver, CALL_MODIFYLOCK simulates a JSB to EXE\$MODIFYLOCK and CALL_MODIFYLOCK_ERR simulates a JSB to EXE\$MODIFYLOCK_ERR. CALL_MODIFYLOCK calls EXE_STD\$MODIFYLOCK, specifying 0 as the **err_rout** argument; CALL_MODIFYLOCK_ERR also calls EXE_STD\$MODIFYLOCK, using the contents of R2 as the **err_rout** argument. Both macros supply the current contents of R3, R4, R5, R6, R0, and R1 as the **irp**, **pcb**, **ucb**, **ccb**, **buf**, and **bufsiz** arguments, respectively.

When EXE_STD\$MODIFYLOCK or EXE_STD\$MODIFYLOCK_ERR returns, code generated by the macro examines the return status:

- If success status (SS\$_NORMAL) is returned, the macro moves the contents of IRP\$\$_SVAPTE into R1 and writes a 5 into R2 to indicate a modify operation. Status is returned in R0 and in the FDT_CONTEXT structure.
- If failure status (SS\$_FDT_COMPL) is returned, the macro writes a 5 to R2 to indicate a modify operation and returns to FDT dispatching code in the \$QIO system service.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$MODIFYLOCK replaces EXE\$MODIFYLOCK and EXE\$MODIFYLOCK_ERR (used by Step 1 drivers). For compatibility with the Step 1 routines, use the \$MODIFYLOCK and \$MODIFYLOCK_ERR macros.
- EXE\$MODIFYLOCK and EXE\$MODIFYLOCK_ERR expect as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.

R0, R7, and R8 are not provided as input to EXE_STD\$MODIFYLOCK.

- The order in which formal parameters are passed to EXE_STD\$MODIFYLOCK differs from the order in which they are provided in registers to the Step 1 routines EXE\$MODIFYLOCK and EXE\$MODIFYLOCK_ERR.
- EXE\$MODIFYLOCK_ERR provides a mechanism by which a driver callback routine obtains control upon an error condition prior to the abortion of an I/O request. EXE_STD\$MODIFYLOCK accepts the address of an error-handling callback routine in the **err_rout** argument. The error-handling routine is called after an I/O request encounters a buffer access or memory allocation failure and before the request is aborted.
- The design of FDT processing for Step 2 OpenVMS AXP device drivers guarantees that the caller of EXE_STD\$MODIFYLOCK regains control whether the modify lock operation is successful or not. When a driver regains control from a call to EXE_STD\$MODIFYLOCK, return status in R0 indicates that the buffer has been successfully locked (SS\$_NORMAL) or that the operation failed and the request has been aborted (SS\$_FDT_COMPL). The driver must check the return status and take appropriate action. Final \$QIO completion status, indicating the reason the operation failed, is stored in the FDT_CONTEXT structure.

Normally, a driver services a modify lock failure by supplying the address of an error-handling callback routine to EXE_STD\$MODIFYLOCK.

- Driver code that executes after receiving failure status (SS\$_FDT_COMPL) from EXE_STD\$MODIFYLOCK cannot access information in the IRP. If the driver anticipates accessing IRP fields when EXE_STD\$MODIFYLOCK returns, it must store these fields elsewhere before calling EXE_STD\$MODIFYLOCK.

System Routines

EXE_STD\$MODIFYLOCK

- Upon successful completion, EXE\$MODIFYLOCK and EXE\$MODIFYLOCK_ERR provide as output the system virtual address of the first process PTE that maps the buffer in R1 and in IRP\$L_SVAPTE. Because EXE_STD\$MODIFYLOCK does not provide R1 as output, a driver must obtain this information from IRP\$L_SVAPTE. Similarly, the Step 1 routines set R2 to 1 to indicate a read function. EXE_STD\$MODIFYLOCK does not provide R2 as output; a driver can determine whether a function is write or read by examining IRP\$V_FUNC in IRP\$L_STS.

EXE_STD\$MOUNT_VER

During I/O postprocessing, determines whether mount verification should be initiated on a given disk or tape device on behalf of the I/O request being completed.

Module

MOUNTVER

Format

status = EXE_STD\$MOUNT_VER (iost1, iost2, irp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
iost1	integer	input	value	required
iost2	integer	input	value	required
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

iost1

First longword of I/O status.

iost2

Second longword of I/O status.

irp

I/O request packet. This argument is 0 if there is no IRP to clean up.

ucb

Unit control block.

Return Values

status

Low bit set indicates that mount verification has not been initiated and that the caller should continue; low bit clear indicates that mount verification has been initiated and that the caller should return.

Context

EXE_STD\$MOUNT_VER is typically called at or above IPL\$IOPOST.

System Routines

EXE_STD\$MOUNT_VER

Description

For Digital internal use only.

Macro

CALL_MOUNT_VER [save_r0r1]

where:

save_r0r1 indicates that the macro should preserve registers R0 and R1 across the call to EXE_STD\$MOUNT_VER. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

In a Step 2 driver, the CALL_MOUNT_VER macro simulates a JSB to EXE\$MOUNT_VER in a Step 1. CALL_MOUNT_VER calls EXE_STD\$MOUNT_VER, using the current contents of R0, R1, R3, and R5 as the **iost1**, **iost2**, **irp**, and **ucb** arguments, respectively. When EXE_STD\$MOUNT_VER returns, code generated by this macro copies return status from R0 to R2. Unless you specify **save_r0r1=NO**, the macro preserves the quadword registers R0 and R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$MOUNT_VER replaces EXE\$MOUNT_VER (used by Step 1 and OpenVMS VAX drivers). Unlike EXE\$MOUNT_VER, EXE_STD\$MOUNT_VER does not preserve R0 and R1 across the call, or provide its return status in R2.

EXE_STD\$ONEPARM

Copies a single \$QIO parameter from IRP\$\$_QIO_P1 to IRP\$\$_MEDIA and delivers the IRP to a driver's start-I/O routine.

Module

SYSQIOFDT

Format

status = EXE_STD\$ONEPARM (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp

I/O request packet for the current I/O request. EXE_STD\$ONEPARM copies the first \$QIO function-specific parameter (**p1**) from IRP\$\$_QIO_P1 to IRP\$\$_MEDIA.

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb

Channel control block that describes the process-I/O channel

Return Values

SS\$_FDT_COMPL

Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$_NORMAL

The routine completed successfully.

System Routines

EXE_STD\$ONEPARM

Context

FDT dispatching code in the \$QIO system service calls EXE_STD\$ONEPARM as an upper-level FDT action routine at IPL\$_ASTDEL.

Description

A driver specifies the system-supplied upper-level FDT action routine EXE_STD\$ONEPARM to process an I/O function code that requires only one parameter. This parameter should need no checking: for instance, for read or write accessibility.

EXE_STD\$ONEPARM copies the first \$QIO function-dependent parameter (**p1**) from IRP\$_QIO_P1 to IRP\$_MEDIA and invokes the \$QIODRVPKT macro to deliver the IRP to the driver. EXE_STD\$ONEPARM regains control with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine EXE\$ONEPARM (used by OpenVMS VAX and Step 1 drivers) obtains the value of the first function-dependent argument (**p1**) specified in the \$QIO request from 00(AP). An OpenVMS AXP FDT routine cannot obtain the argument as an offset from the AP; rather, it accesses the argument from a new IRP field, IRP\$_QIO_P1.
In order to convert an OpenVMS VAX driver to an OpenVMS AXP driver, the upper-level action routine EXE_STD\$ONEPARM exists, to move the value of (**p1**) from IRP\$_QIO_P1 to IRP\$_MEDIA and invokes the \$QIODRVPKT macro. If your driver does not access (**p1**) from IRP\$_MEDIA, but rather, uses the contents of IRP\$_QIO_P1, specifying EXE_STD\$ONEPARM as an upper-level FDT action routine may defy all logic. (If your driver ignores the contents of IRP\$_MEDIA, it is immaterial whether you specify EXE_STD\$ZEROPARM or EXE_STD\$ONEPARM as the upper-level FDT action routine that delivers the IRP to the driver.) To avoid the unnecessary copy, you can write an upper-level FDT action routine that invokes the \$QIODRVPKT macro.
- EXE\$ONEPARM expects as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table (FDT) from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$ONEPARM.
- EXE\$ONEPARM returns control to the system service dispatcher, passing it the final \$QIO system service status (SSS_NORMAL) in R0. EXE_STD\$ONEPARM returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

EXE_STD\$PRIMITIVE_FORK

Creates a simple fork process on the local processor.

Module

FORKCNTRL

Format

EXE_STD\$PRIMITIVE_FORK (fr3, fr4, fkb)

Arguments

Argument	Type	Access	Mechanism	Status
fr3	int64	input	value	required
fr4	int64	input	value	required
fkb	FKB	input	reference	required

fr3

Value to pass to the fork routine in FKBSQ_FR3.

fr4

Value to pass to the fork routine in FKBSQ_FR4.

fkb

Fork block. At input, FKBSB_FLCK must contain the fork lock index and FKBSL_FPC must contain the procedure value of the fork routine.

Context

EXE_STD\$PRIMITIVE_FORK acquires no spin locks and leaves IPL unchanged. EXE_STD\$PRIMITIVE_FORK, unlike the OpenVMS VAX system routine EXE\$FORK, returns to its caller and not to its caller's caller. It assumes that, prior to the call, its caller has placed the procedure value of the fork routine into FKBSL_FPC.

EXE_STD\$PRIMITIVE_FORK provides fork context to the fork routine in FKBSQ_FR3 (contents of **fr3**) and FKBSQ_FR4 (contents **fr4**). All other registers are destroyed. The fork routine executes at the IPL indicated by the fork lock index stored in FKBSB_FLCK.

Description

EXE_STD\$PRIMITIVE_FORK moves the contents of the **fr3** and **fr4** arguments into FKBSQ_FR3 and FKBSQ_FR4, respectively. It determines the fork IPL by using the value of FKBSB_FLCK as an index into the spin lock IPL vector (SMP\$AL_IPLVEC). EXE_STD\$PRIMITIVE_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_STD\$PRIMITIVE_FORK issues a SOFTINT macro, requesting a software interrupt from the local processor at that fork IPL.

System Routines

EXE_STD\$PRIMITIVE_FORK

A driver that calls EXE_STD\$PRIMITIVE_FORK explicitly (that is, instead of invoking the IOFORK macro) must ensure that UCB\$V_TIM in the UCB\$L_STS field is clear before making the call.

Macro

FORK
IOFORK

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$PRIMITIVE_FORK is a call-based routine that performs the same operation as the JSB-based routine EXE\$PRIMITIVE_FORK on OpenVMS AXP systems. The OpenVMS VAX routines EXE\$FORK and EXE\$IOFORK are not provided on OpenVMS AXP systems.

EXE_STD\$PRIMITIVE_FORK_WAIT

Inserts a fork block on the fork-and-wait queue.

Module

FORKCNTRL

Format

EXE_STD\$PRIMITIVE_FORK_WAIT (fr3, fr4, fkb)

Arguments

Argument	Type	Access	Mechanism	Status
fr3	int64	input	value	required
fr4	int64	input	value	required
fkb	FKB	input	reference	required

fr3

Value to pass to the fork routine in FKBSQ_FR3.

fr4

Value to pass to the fork routine in FKBSQ_FR4.

fkb

Fork block. At input, FKBSB_FLCK must contain the fork lock index and FKBSL_FPC must contain the procedure value of the fork routine.

Context

The caller of EXE_STD\$PRIMITIVE_FORK_WAIT must be executing at or above IPL\$ SYNCH. EXE_STD\$PRIMITIVE_FORK_WAIT acquires the MEGA (SPL\$C_MEGA) spin lock, raising IPL to IPL\$ MEGA in the process, to access the fork-and-wait queue (EXE\$AR_FORK_WAIT_QUEUE). It releases the spin lock, restoring the previous IPL, prior to returning to its caller.

EXE_STD\$PRIMITIVE_FORK_WAIT, unlike the OpenVMS VAX system routine EXE\$FORK_WAIT, returns to its caller and not to its caller's caller. It assumes that, prior to the call, its caller has placed the procedure value of the fork routine into FKBSL_FPC.

EXE_STD\$PRIMITIVE_FORK_WAIT provides fork context to the fork routine in FKBSQ_FR3 (contents of **fr3**) and FKBSQ_FR4 (contents of **fr4**). All other registers are destroyed. The fork routine executes at the IPL indicated by the fork lock index stored in FKBSB_FLCK.

System Routines

EXE_STD\$PRIMITIVE_FORK_WAIT

Description

EXE_STD\$PRIMITIVE_FORK_WAIT moves the contents of **fr3** and **fr4** into FKBSQ_FR3 and FKBSQ_FR4 respectively. Having obtained the MEGA spin lock, it inserts the fork block indicated by **fkf** at end of the fork-and-wait queue (EXESGL_FKWAITBL) and releases the spin lock.

Up to one second later, the software timer interrupt service routine will remove this and all other entries from the fork-and-wait queue and resume their respective fork routines.

Macro

FORK_WAIT

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$PRIMITIVE_FORK_WAIT is a call-based routine that performs the same operation as the JSB-based routine EXE\$PRIMITIVE_FORK_WAIT on OpenVMS AXP systems. The OpenVMS VAX routines EXE\$FORK and EXE\$IOFORK are not provided on OpenVMS AXP systems.

EXE_STD\$QIOACPPKT

Delivers an IRP to the appropriate ACP or XQP.

Module

SYSQIOREQ

Format

status = EXE_STD\$QIOACPPKT (irp, pcb, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required

irp
I/O request packet.

pcb
Process control block.

ucb
Unit control block.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$NORMAL The routine completed successfully.

Context

EXE_STD\$QIOACPPKT is called by, and returns to, a driver's FDT routine at IPL\$ASTDEL.

Description

For Digital internal use only.

System Routines

EXE_STD\$QIOACPPKT

Macro

CALL_QIOACPPKT [do_ret=YES]

where:

do_ret indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

In a Step 2 driver, the CALL_QIOACPPKT macro simulates the JMP to EXE\$QIOACPPKT in the FDT routine of a Step 1 driver. It calls EXE_STD\$QIOACPPKT using the current contents of R3, R4, and R5 as the **irp**, **pcb**, and **ucb** arguments, respectively. When EXE_STD\$QIOACPPKT returns control to the code generated by a default invocation of CALL_QIOACPPKT, a RET instruction returns control to the caller of the macro's invoker. Status is returned in R0 and in the FDT_CONTEXT structure.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE\$QIOACPPKT returns control to the system service dispatcher, passing it the final \$QIO system service status (SS\$_NORMAL) in R0. EXE_STD\$QIOACPPKT returns to its caller, passing it SS\$_FDT_COMPL status in R0 and storing the final \$QIO system service status (SS\$_NORMAL) in the FDT context structure. The \$QIO system service retrieves the status from this structure.

EXE_STD\$QIODRVPKT

Delivers an IRP to a driver's start-I/O routine or pending-I/O queue.

Module

SYSQIOREQ

Format

status = EXE_STD\$QIODRVPKT (irp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

irp

I/O request packet. EXE_STD\$QIODRVPKT clears IRP\$PS_FDT_CONTEXT. The caller of EXE_STD\$QIODRVPKT should not access the IRP after the routine returns SS\$_FDT_COMPL status.

ucb

Unit control block.

EXE_STD\$QIODRVPKT (by means of the call to EXE_STD\$INSIOQ) reads the following UCB fields:

Field	Contents
UCB\$_FLCK	Fork lock index
UCB\$_STS	UCB\$_BSY set if device is busy, clear if device is idle
UCB\$_IOQFL	Address of pending-I/O queue listhead
UCB\$_QLEN	Length of pending-I/O queue

EXE_STD\$QIODRVPKT (by means of the call to EXE_STD\$INSIOQ) writes the following UCB fields:

Field	Contents
UCB\$_STS	UCB\$_BSY set
UCB\$_QLEN	Increased

System Routines

EXE_STD\$QIODRVPKT

Return Values

SS\$FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.
---------------	---

Status in FDT_CONTEXT

SS\$NORMAL	The routine completed successfully.
------------	-------------------------------------

Context

EXE_STD\$QIODRVPKT is called by, and returns to, a driver's FDT routine at IPL\$ASTDEL.

Description

The FDT completion routine EXE_STD\$QIODRVPKT delivers an IRP to the driver's start-I/O routine or pending-I/O queue.

EXE_STD\$QIODRVPKT clears the pointer to the FDT context structure in IRP\$PS_FDT_CONTEXT and calls EXE_STD\$INSIOQ. EXE_STD\$INSIOQ checks the status of the device and calls either EXE_STD\$INSERT_IRP or IOC_STD\$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_STD\$INSIOQ returns, EXE_STD\$QIODRVPKT returns SS\$FDT_COMPL status to its caller, passing SS\$NORMAL as the final \$QIO completion status in the FDT context structure.

The image that requested the I/O operation receives SS\$NORMAL status, indicating that the I/O request has completed without device-independent error.

Macro

CALL_QIODRVPKT [do_ret=YES]

where:

do_ret indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

In a Step 2 driver, the CALL_QIODRVPKT macro simulates the JMP to EXE\$QIODRVPKT in the FDT routine of a Step 1 driver. CALL_QIODRVPKT clears IRP\$PS_FDT_CONTEXT and calls EXE_STD\$INSIOQ, using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively. When EXE_STD\$INSIOQ returns control to the code generated by a default invocation of \$QIODRVPKT, a RET instruction returns control to the caller of the macro's invoker. Status is returned in R0.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The address of the PCB, supplied as input to EXE\$QIODRVPKT on OpenVMS VAX systems, is not provided as input to EXE_STD\$QIODRVPKT.
- Unlike EXE\$QIODRVPKT, EXE_STD\$QIODRVPKT does not lower IPL to 0 before exiting. EXE_STD\$QIODRVPKT returns to its caller at the caller's IPL.
- EXE\$QIODRVPKT returns control to the system service dispatcher, passing it the final \$QIO system service status (SS\$_NORMAL) in R0. EXE_STD\$QIODRVPKT returns to its caller, passing it SS\$_FDT_COMPL status in R0 and storing the final \$QIO system service status (SS\$_NORMAL) in the FDT context structure. The \$QIO system service retrieves the status from this structure.

System Routines
EXE_STD\$QUEUE_FORK

EXE_STD\$QUEUE_FORK

TBS

EXE_STD\$QXQPPKT

Inserts an I/O request packet on the end of the XQP work queue and initiates its processing if it is the only request on the queue.

Module

SYSQIOREQ

Format

status = EXE_STD\$QXQPPKT (pcb, acb)

Arguments

Argument	Type	Access	Mechanism	Status
pcb	PCB	input	reference	required
acb	ACB	input	reference	required

pcb
Process control block.

acb
AST control block within the IRP.

Return Values

SS\$FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.
---------------	---

Status in FDT_CONTEXT

SS\$NORMAL	The routine completed successfully.
------------	-------------------------------------

Context

EXE_STD\$QXQPPKT is called by, and returns to, a driver's FDT routine at IPL\$ASTDEL.

Description

For Digital internal use only.

System Routines

EXE_STD\$QXQPPKT

Macro

CALL_QXQPPKT

In a Step 2 driver, the CALL_QXQPPKT macro simulates the JMP to EXE\$QXQPPKT in the FDT routine of a Step 1 driver. It calls EXE_STD\$QXQPPKT using the current contents of R4 and R5 as the **pcb** and **acb** arguments, respectively. Status is returned in R0 and in the FDT_CONTEXT structure.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE\$QXQPPKT returns control to the system service dispatcher, passing it the final \$QIO system service status (SS\$_NORMAL) in R0. EXE_STD\$QXQPPKT returns to its caller, passing it SS\$_FDT_COMPL status in R0 and storing the final \$QIO system service status (SS\$_NORMAL) in the FDT context structure. The \$QIO system service retrieves the status from this structure.

EXE_STD\$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 aborts the request or proceeds with a direct-I/O, DMA write operation.

Module

SYSQIOFDT

Format

status = EXE_STD\$READ (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp

I/O request packet for the current I/O request.

EXE_STD\$READ reads the following IRP fields:

Field	Contents
IRP\$QIO_P1	\$QIO system service p1 argument, containing the buffer's virtual address.
IRP\$QIO_P2	\$QIO system service p2 argument, containing the number of bytes in transfer. The maximum number of bytes that EXE_STD\$READ can transfer is 65,535 (128 pages minus one byte).
IRP\$QIO_P4	\$QIO system service p4 argument, containing the carriage control byte.
IRP\$FUNC	I/O function code.
IRP\$B_RMOD	Access mode of the caller of the \$QIO system service.

EXE_STD\$READ writes the following IRP fields:

Field	Contents
IRP\$B_CARCON	Carriage control byte (from IRP\$QIO_P4)
IRP\$FUNC	Logical read function code converted to physical
IRP\$STS	IRP\$V_FUNC set to indicate read function

System Routines

EXE_STD\$READ

Field	Contents
IRP\$ <u>L</u> _SVAPTE	System virtual address of the PTE that maps the first page of the buffer
IRP\$ <u>L</u> _BOFF	Byte offset to start of transfer in page
IRP\$ <u>L</u> _OBOFF	Original byte offset into the first page of a segmented direct-I/O transfer
IRP\$ <u>L</u> _BCNT	Size of transfer in bytes

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb

Channel control block that describes the process-I/O channel

Return Values

SS\$ <u>FDT</u> _COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.
------------------------	---

Status in FDT_CONTEXT

SS\$ <u>ACCVIO</u>	Buffer specified in buf parameter does not allow write access.
SS\$ <u>BADPARAM</u>	bufsiz parameter is less than zero.
SS\$ <u>INSFWSL</u>	Insufficient working set limit.
SS\$ <u>NORMAL</u>	The I/O request has been successfully queued.
SS\$ <u>QIO_CROCK</u>	Buffer page must be faulted into memory.

Context

FDT dispatching code in the \$QIO system service calls EXE_STD\$READ as an upper-level FDT action routine at IPL\$ASTDEL.

Description

A driver specifies the system-supplied upper-level FDT action routine EXE_STD\$READ to prepare a direct-I/O read request. A driver cannot specify EXE_STD\$READ for buffered-I/O functions. Drivers that process functions that require an intermediate system buffer typically supply their own upper-level FDT action routines to handle them.

EXE_STD\$READ performs the following functions:

- Sets IRP\$V_FUNC in IRP\$L_STS to indicate a read function
- Copies the **p4** argument of the \$QIO request from IRP\$L_QIO_P4 to IRP\$B_CARCON

- Translates a logical read function to a physical read function and stores the new function code in IRP\$L_FUNC.
- Examines the size of the transfer, as specified in the **p2** argument of the \$QIO request (IRP\$L_QIO_P2), and takes one of the following actions:
 - If the transfer byte count is zero, EXE_STD\$READ invokes the \$QIODRVPKT macro to deliver the IRP to the driver's start-I/O routine. EXE_STD\$READ regains control with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure. It returns to the \$QIO system service, passing these status values.
The driver start-I/O routine should check for zero-length buffers to avoid mapping to adapter node space. An attempted mapping can cause a system failure.
 - If the byte count is not zero, EXE_STD\$READ calls EXE_STD\$READLOCK, specifying 0 as the **err_rout** argument.

EXE_STD\$READLOCK invokes the \$READCHK macro, which calls EXE_STD\$READCHK.

EXE_STD\$READCHK performs the following actions:

- Moves the transfer byte count (**bufsiz** parameter) into IRP\$L_BCNT.
If the byte count is negative, it calls EXE_STD\$ABORTIO, passing it a **qio_sts** of SSS_BADPARAM. When it regains control, EXE_STD\$READCHK returns to EXE_STD\$READLOCK with SSS_BADPARAM status in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$READLOCK immediately returns to EXE_STD\$READ, passing these status values. EXE_STD\$READ, in turn, returns to the \$QIO system service.
- 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_STD\$READCHK sets IRP\$V_FUNC in IRP\$L_STS and returns SSS_NORMAL in R0 to EXE_STD\$READLOCK.
 - If the buffer does not allow write access, EXE_STD\$READCHK calls EXE_STD\$ABORTIO, passing it a **qio_sts** of SSS_ACCVIO. When it regains control, EXE_STD\$READCHK returns to EXE_STD\$READLOCK with SSS_ACCVIO status in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$READLOCK immediately returns to EXE_STD\$READ, passing these status values. EXE_STD\$READ returns to the \$QIO system service.

If EXE_STD\$READCHK succeeds, EXE_STD\$READLOCK moves into IRP\$L_BOFF and IRP\$L_OBOFF the byte offset to the start of the buffer and calls MMG_STD\$IOLOCK.

MMG_STD\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:

- If MMG_STD\$IOLOCK succeeds, EXE_STD\$READLOCK stores in IRP\$L_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns SSS_NORMAL status in R0 to EXE_STD\$READLOCK. EXE_STD\$READLOCK returns immediately to EXE_STD\$READ, passing to it this status value.

System Routines

EXE_STD\$READ

EXE_STD\$READ invokes the \$QIODRVPKT macro to deliver the IRP to the driver's start-I/O routine. EXE_STD\$READ regains control with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure. It returns to the \$QIO system service, passing these status values.

- If MMG_STD\$IOLOCK fails, it returns SSS_ACCVIO, SSS_INSFWSL, or page fault status to EXE_STD\$READLOCK.

For SSS_ACCVIO and SSS_INSFWSL status, EXE_STD\$READLOCK calls EXE_STD\$ABORTIO, passing it one of these status values as a **qio_sts** argument. When it regains control, EXE_STD\$READLOCK returns EXE_STD\$READ the specified status value in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$READ returns to the \$QIO system service.

For page fault status, EXE_STD\$READLOCK sets the final \$QIO status in the FDT_CONTEXT structure to SSS_QIO_CROCK and initializes FDT_CONTEXT\$SL_QIO_R1_VALUE to the virtual address to be faulted. It then adjusts the 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.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine EXE\$READ (used by OpenVMS VAX and Step 1 device drivers) expects as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$READ.
- EXE\$READ returns control to the system service dispatcher, passing it the final \$QIO system service status (SSS_NORMAL, SSS_ACCVIO, or SSS_BADPARAM, or SSS_INSFWSL) in R0. EXE_STD\$READ returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

EXE_STD\$READCHK

Verifies that a process has write access to the pages in the buffer specified in a \$QIO request.

Module

SYSQIOFDT

Format

status = EXE_STD\$READCHK (irp, pcb, ucb, buf, bufsiz)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
buf	address	input	reference	required
bufsiz	integer	input	value	required

irp

I/O request packet for the current I/O request.

EXE_STD\$READCHK reads IRP\$B_RMOD to determine the access mode of the caller of the \$QIO system service.

EXE_STD\$READCHK writes the following IRP fields:

Field	Contents
IRP\$L_STS	IRP\$V_FUNC set, indicating a read function
IRP\$L_BCNT	Size of transfer in bytes

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

buf

Virtual address of buffer.

bufsiz

Number of bytes in transfer.

System Routines

EXE_STD\$READCHK

Return Values

SS\$_NORMAL	The buffer is write-accessible.
SS\$_FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$_ACCVIO	Buffer specified in buf parameter does not allow write access.
SS\$_BADPARAM	bufsiz parameter is less than zero.
SS\$_INSFWSL	Insufficient working set limit.
SS\$_NORMAL	Nothing has occurred yet to prevent the I/O request from being successfully queued. This is the initial value of the status field in an FDT_CONTEXT structure.

Context

The FDT support routine EXE_STD\$READLOCK, or a driver-specific FDT routine, calls EXE_STD\$READCHK at IPL\$_ASTDEL.

Description

A driver FDT routine calls the system-supplied FDT support routine EXE_STD\$READCHK to check the write accessibility of an I/O buffer supplied in a \$QIO request for a read function.

EXE_STD\$READCHK performs the following actions:

- Moves the transfer byte count (**bufsiz** parameter) into IRP\$\$_BCNT. If the byte count is negative, it calls EXE_STD\$ABORTIO, passing it a **qio_sts** of SS\$_BADPARAM. When it regains control, EXE_STD\$READCHK returns to its caller with SS\$_BADPARAM status in the FDT_CONTEXT structure and SS\$_FDT_COMPL status in R0.
- 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_STD\$READCHK sets IRP\$\$_FUNC in IRP\$\$_STS and returns SS\$_NORMAL in R0 to its caller.
 - If the buffer does not allow write access, EXE_STD\$READCHK calls EXE_STD\$ABORTIO, passing it a **qio_sts** of SS\$_ACCVIO. When it regains control, EXE_STD\$READCHK returns to its caller with SS\$_ACCVIO status in the FDT_CONTEXT structure and SS\$_FDT_COMPL status in R0.

The caller of EXE_STD\$READCHK must examine the status in R0:

- If the status is SS\$_NORMAL, the buffer is write-accessible.
- If the status is SS\$_FDT_COMPL, an error has occurred that has caused the I/O request to be aborted. You can determine the reason for the failure from FDT_CONTEXT\$\$_QIO_STATUS.

Certain drivers must perform additional processing to back out an I/O request after it has aborted. For instance, if the driver has locked multiple buffers into memory for a single I/O request, it must unlock them once the request has been aborted. A driver cannot access the IRP once it has received SSS_FDT_COMPL status. If you know you need access to information stored in the IRP to back out an I/O request that has been aborted, you must store that information elsewhere prior to calling EXE_STD\$READCHK.

Macro

CALL_READCHK
CALL_READCHKR

In a Step 2 driver, CALL_READCHK simulates a JSB to EXE\$READCHK and CALL_READCHKR simulates a JSB to EXE\$READCHKR. Both macros call EXE_STD\$READCHK using the current contents of R3, R4, R5, R0, and R1 as the **irp**, **pcb**, **ucb**, **buf**, and **bufsiz** arguments, respectively.

When EXE_STD\$READCHK returns, \$READCHK and \$READCHKR move 1 into R2 to indicate a read operation and examines the return status:

- If success status (SS\$_NORMAL) is returned, CALL_READCHK and CALL_READCHKR copy the contents of IRP\$L_BCNT into R1. CALL_READCHK writes the starting address of the I/O buffer in R0; CALL_READCHKR preserves the return status value in R0.
- If failure status (SS\$_FDT_COMPL) is returned, CALL_READCHK returns to FDT dispatching code in the \$QIO system service. CALL_READCHKR does not return control to \$QIO.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$READCHK replaces EXE\$READCHK and EXE\$READCHKR (used by Step 1 drivers). For compatibility with the Step 1 routines, use the CALL_READCHK and CALL_READCHKR macros.
- EXE\$READCHK and EXE\$READCHKR expect as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$READCHK.
- The order in which formal parameters are passed to EXE_STD\$READCHK differs from the order in which they are provided in registers to the Step 1 routines EXE\$READCHK and EXE\$READCHKR.
- EXE\$READCHK and EXE\$READCHKR provide a mechanism by which a driver callback routine or coroutine obtains control upon an error condition prior to the abortion of an I/O request. The design of FDT processing for Step 2 device drivers guarantees that the caller of EXE_STD\$READCHK regains control whether the read check operation is successful. The caller must examine the return status in R0 (SS\$_NORMAL indicates the buffer is write accessible, SS\$_FDT_COMPL indicates that the operation failed and the request has been aborted) and take appropriate action. Final \$QIO

System Routines

EXE_STD\$READCHK

completion status, indicating the reason the operation failed, is stored in the `FDT_CONTEXT` structure.

- Driver code that services failure status (`SS$FDT_COMPL`) from `EXE_STD$READLOCK` (for instance a callback routine formerly specified to `EXE$READLOCK_ERR`) cannot access information in the IRP. If the driver anticipates handling failure status by using the contents of IRP fields, it must store these fields elsewhere before calling `EXE_STD$READLOCK`.

This is especially important for driver code that expects `EXE_STD$READCHK` to access the transfer size in `R1` after the call. Unlike `EXE$READCHK` and `EXE$READCHKR`, `EXE_STD$READCHK` does not preserve the contents of `R1` and `R3` across the call. If you must repeat a `CALL_READCHK` macro invocation, you must be sure to reload `R0`, `R1`, and `R3` with the virtual address of the buffer, the transfer size, and the address of the IRP, respectively, before each subsequent invocation.

- Upon successful completion, `EXE$READCHK` and `EXE$READCHKR` set `R2` to 1 for a read function. `EXE_STD$READCHK` does not provide `R2` as output; a driver can determine whether a function is read or write by examining `IRPSV_FUNC` in `IRP$SL_STS`.

EXE_STD\$READLOCK

Validates and prepares a user buffer for a direct-I/O, DMA write operation.

Module

SYSQIOFDT

Format

status = EXE_STD\$READLOCK (irp, pcb, ucb, ccb, buf, bufsiz, err_rout)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required
buf	address	input	reference	required
bufsiz	integer	input	value	required
err_rout	procedure value	input	value	required

irp

I/O request packet for the current I/O request.

EXE_STD\$READLOCK reads IRP\$B_RMOD to determine the access mode of the caller of the \$QIO system service.

EXE_STD\$READLOCK writes the following IRP fields:

Field	Contents
IRP\$L_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\$L_BOFF	Byte offset to start of transfer in page
IRP\$L_OBOFF	Original byte offset into the first page of a segmented direct-I/O transfer
IRP\$L_BCNT	Size of transfer in bytes

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

System Routines

EXE_STD\$READLOCK

ccb

Channel control block that describes the process-I/O channel.

buf

Virtual address of buffer.

bufsiz

Number of bytes in transfer

err_rout

Procedure value of error-handling callback routine, or 0 if the driver does not process errors.

A driver typically specifies an error-handling callback routine when the driver must lock multiple areas into memory for a single I/O request and regain control to unlock these areas, if the request is to be aborted. The routine performs those tasks required before the request is backed out of or aborted. Such operations could include calling MMG_STD\$UNLOCK to release previous buffers participating in the I/O operation. The error-handling routine must preserve R0 and R1 and return back to EXE_STD\$READLOCK.

Chapter 1 describes the error-handling callback routine interface.

Return Values

SS\$NORMAL

The buffer is write-accessible and has been locked in memory.

SS\$_FDT_COMPL

Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$_ACCVIO

Buffer specified in **buf** parameter does not allow write access.

SS\$_BADPARAM

bufsiz parameter is less than zero.

SS\$_INSFWSL

Insufficient working set limit.

SS\$_NORMAL

Nothing has occurred yet to prevent the I/O request from being successfully queued. This is the initial value of the status field in an FDT_CONTEXT structure.

SS\$_QIO_CROCK

Buffer page must be faulted into memory.

Context

The system-supplied upper-level FDT action routine EXE_STD\$READ, or a driver-specific upper-level FDT action routine, calls EXE_STD\$READLOCK at IPL\$ASTDEL.

Description

A driver FDT routine calls the system-supplied FDT support routine EXE_STD\$READLOCK to check the write accessibility of an I/O buffer supplied in a \$QIO request for a read function, and lock the buffer in memory in preparation for a DMA write operation.

A driver cannot specify EXE_STD\$READ for buffered-I/O functions. Drivers that process functions that require an intermediate system buffer typically supply their own FDT routines to handle them.

EXE_STD\$READLOCK invokes the \$READCHK macro, which calls EXE_STD\$READCHK.

EXE_STD\$READCHK performs the following actions:

- Moves the transfer byte count (**bufsiz** parameter) into IRP\$LCBNT. If the byte count is negative, EXE_STD\$READCHK returns SSS_BADPARAM status to EXE_STD\$READLOCK.
- 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_STD\$READCHK sets IRP\$V_FUNC in IRP\$STS and returns SSS_NORMAL in R0 to EXE_STD\$READLOCK.
 - If the buffer does not allow write access, EXE_STD\$READCHK returns SSS_ACCVIO status to EXE_STD\$READLOCK.

If error status (SSS_BADPARAM or SSS_ACCVIO) is returned, EXE_STD\$READLOCK immediately calls the specified error-handling callback routine, passing to it the IRP, PCB, UCB, CCB, and status value. The callback routine must preserve R0 and R1 and return control to EXE_STD\$READLOCK. When the callback routine returns (or if no callback routine is specified), EXE_STD\$READLOCK calls EXE_STD\$ABORTIO, passing it the error status as **qio_sts**. EXE_STD\$ABORTIO returns to EXE_STD\$READLOCK with the error status in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$READLOCK immediately returns to its caller, passing these status values.

If SSS_NORMAL status is returned, EXE_STD\$READLOCK moves into IRP\$LCBOFF and IRP\$LCBOFF the byte offset to the start of the buffer and calls MMG_STD\$IOLOCK.

MMG_STD\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:

- If MMG_STD\$IOLOCK succeeds, EXE_STD\$READLOCK stores in IRP\$LCB_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns SSS_NORMAL status in R0 to EXE_STD\$READLOCK. EXE_STD\$READLOCK returns immediately to its caller, passing to it this status value.
- If MMG_STD\$IOLOCK fails, it returns SSS_ACCVIO, SSS_INSFWSL, or page fault status to EXE_STD\$READLOCK. EXE_STD\$READLOCK immediately calls the specified error-handling callback routine, passing to it the IRP, PCB, UCB, CCB, and status value. The callback routine must preserve R0 and R1 and return control to EXE_STD\$READLOCK. When

System Routines

EXE_STD\$READLOCK

the callback routine returns (or if no callback routine is specified), EXE_STD\$READLOCK proceeds as follows:

- For SSS_ACCVIO and SSS_INSFWSL status, EXE_STD\$READLOCK calls EXE_STD\$ABORTIO, passing it one of these status values as a **qio_sts** argument. When it regains control, EXE_STD\$READLOCK returns to its caller the specified status value in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0.
- For page fault status, EXE_STD\$READLOCK sets the final \$QIO status in the FDT_CONTEXT structure to SSS_QIO_CROCK and initializes FDT_CONTEXT\$SL_QIO_R1_VALUE to the virtual address to be faulted. It then adjusts the 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.

The caller of EXE_STD\$READLOCK must examine the status in R0:

- If the status is SSS_NORMAL, the buffer is write accessible and has been successfully locked into memory and the starting virtual address of the page table entries that map the buffer is available in IRP\$SL_SVAPTE.
- If the status is SSS_FDT_COMPL, an error has occurred that has caused the I/O request to be aborted. You can determine the reason for the failure from FDT_CONTEXT\$SL_QIO_STATUS. Ordinarily a driver specifies an error-handling callback routine to process such errors.

Note that a driver cannot access the IRP once it has received SSS_FDT_COMPL status. If you know you need access to information stored in the IRP to back out an I/O request that has been aborted, you must store that information elsewhere prior to calling EXE_STD\$READLOCK.

Macro

```
CALL_READLOCK  
CALL_READLOCK_ERR [interface_warning=YES]
```

where:

interface_warning=YES, the default, specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. **interface_warning=NO** suppresses the warning.

In a Step 2 driver, the CALL_READLOCK simulates a JSB to EXE\$READLOCK and CALL_READLOCK_ERR simulates a JSB to EXE\$READLOCK_ERR. CALL_READLOCK calls EXE_STD\$READLOCK, specifying 0 as the **err_rout** argument; CALL_READLOCK_ERR also calls EXE_STD\$READLOCK, using the contents of R2 as the **err_rout** argument. Both macros supply the current contents of R3, R4, R5, R6, R0, and R1 as the **irp**, **pcb**, **ucb**, **ccb**, **buf**, and **bufsiz** arguments, respectively.

When EXE_STD\$READLOCK or EXE_STD\$READLOCK_ERR returns, code generated by the macro examines the return status:

- If success status (SS\$NORMAL) is returned, the macro copies the contents of IRP\$LVAPTE into R1 and writes a 1 to R2 to indicate a read operation. Status is returned in R0 and in the FDT_CONTEXT structure.
- If failure status (SS\$_FDT_COMPL) is returned, the macro writes a 1 to R2 to indicate a read operation and returns to FDT dispatching code in the \$QIO system service.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$READLOCK replaces EXES\$READLOCK and EXES\$READLOCK_ERR (used by Step 1 drivers). For compatibility with the Step 1 routines, use the CALL_READLOCK and CALL_READLOCK_ERR macros.
- EXES\$READLOCK and EXES\$READLOCK_ERR expect as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$READLOCK.
- The order in which formal parameters are passed to EXE_STD\$READLOCK differs from the order in which they are provided in registers to the Step 1 routines EXES\$READLOCK and EXES\$READLOCK_ERR.
- EXES\$READLOCK_ERR provides a mechanism by which a driver callback routine obtains control upon an error condition prior to the abortion of an I/O request. EXE_STD\$READLOCK accepts the address of an error-handling callback routine in the **err_rout** argument. The error-handling routine is called after an I/O request encounters a buffer access or memory allocation failure and before the request is aborted.
- The design of FDT processing for Step 2 drivers guarantees that the caller of EXE_STD\$READLOCK regains control whether the read lock operation is successful. When a driver regains control from a call to EXE_STD\$READLOCK, return status in R0 indicates that the buffer has been successfully locked (SS\$NORMAL) or that the operation failed and the request has been aborted (SS\$_FDT_COMPL). The driver must check the return status and take appropriate action. Final \$QIO completion status, indicating the reason the operation failed, is stored in the FDT_CONTEXT structure.

Normally, a driver services a read lock failure by supplying the address of an error-handling callback routine to EXE_STD\$READLOCK.

- Driver code that executes after receiving failure status (SS\$_FDT_COMPL) from EXE_STD\$READLOCK cannot access information in the IRP. If the driver anticipates accessing IRP fields when EXE_STD\$READLOCK returns, it must store these fields elsewhere before calling EXE_STD\$READLOCK.
- Upon successful completion, EXES\$READLOCK and EXES\$READLOCK_ERR provide as output the system virtual address of the first process PTE that maps the buffer in R1 and in IRP\$LVAPTE. Because EXE_STD\$READLOCK does not provide R1 as output, a driver must obtain this information from IRP\$LVAPTE. Similarly, the Step 1 routines set R2 to 1

System Routines

EXE_STD\$READLOCK

for a read function and clear it otherwise. EXE_STD\$READLOCK does not provide R2 as output; a driver can determine whether a function is read or write by examining IRPSV_FUNC in IRP\$L_STS.

EXE_STD\$SENSEMODE

Copies device-dependent characteristics from the device's UCB into the second longword of the I/O status block (IOSB) specified in a \$QIO system service call, and completes the I/O operation successfully.

Module

SYSQIOFDT

Format

status = EXE_STD\$SENSEMODE (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp

I/O request packet for the current I/O request.

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request. EXE_STD\$SENSEMODE reads the device-dependent status stored in UCB\$SL_DEVDEPEND.

ccb

Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL

Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$NORMAL

The routine completed successfully.

System Routines

EXE_STD\$SENSEMODE

Context

FDT dispatching code in the \$QIO system service calls EXE_STD\$SENSEMODE as an upper-level FDT action routine at IPL\$ASTDEL.

Description

A driver specifies the system-supplied upper-level FDT action routine EXE_STD\$SENSEMODE to process the sense-device-mode (IO\$_SENSEMODE) and sense-device-characteristics (IO\$_SENSECHAR) I/O functions.

EXE_STD\$SENSEMODE loads the contents of UCB\$_DEVDEPEND into the second longword of the I/O status block (IOSB) specified in the original \$QIO system service call. It then places SSS\$_NORMAL status into the FDT_CONTEXT structure and transfers control to EXE_STD\$FINISHIO to insert the IRP in the local processor's I/O postprocessing queue.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine EXE\$SENSEMODE (used by OpenVMS VAX and Step 1 device drivers) expects as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$SENSEMODE.
- EXE\$SENSEMODE returns control to the system service dispatcher, passing it the final \$QIO system service status (SSS\$_NORMAL) in R0. EXE_STD\$SENSEMODE returns to its caller, passing it SSS\$_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

EXE_STD\$SETCHAR, EXE_STD\$SETMODE

Write device-specific status and control information into the device's UCB and complete the I/O request (EXE_STD\$SETCHAR); or write the information into the IRP and deliver the IRP to the driver's start-I/O routine (EXE_STD\$SETMODE).

Module

SYSQIOFDT

Format

status = EXE_STD\$SETCHAR (irp, pcb, ucb, ccb)

status = EXE_STD\$SETMODE (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp

I/O request packet for the current I/O request.

EXE_STD\$SETCHAR and EXE_STD\$SETMODE read the following IRP fields:

Field	Contents
IRP\$L_FUNC	I/O function code supplied in the \$QIO request
IRP\$B_RMOD	Mode of the \$QIO caller
IRP\$L_QIO_P1	\$QIO system service p1 argument, containing the device characteristics quadword.

EXE_STD\$SETMODE writes the following IRP fields:

Field	Contents
IRP\$L_MEDIA	First longword of device characteristics
IRP\$L_MEDIA+4	Second longword of device characteristics

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

System Routines

EXE_STD\$SETCHAR, EXE_STD\$SETMODE

EXE_STD\$SETCHAR writes the following UCB fields:

Field	Contents
UCB\$B_DEVCLASS	Byte 0 of device characteristics quadword
UCB\$B_DEVTYPE	Byte 1 of device characteristics quadword
UCB\$W_DEVBUSIZ	Bytes 2 and 3 of device characteristics quadword
UCB\$L_DEVDEPEND	Bytes 4 through 7 of device characteristics quadword

ccb

Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$NORMAL The routine completed successfully.
SS\$ACCVIO Process calling the \$QIO system service with the IO\$SETMODE or IO\$SETCHAR function does not have read access to the quadword containing the new device characteristics.
SS\$ILLIOFUNC IO\$SETMODE and IO\$SETCHAR functions are not legal for disk devices.

Context

FDT dispatching code in the \$QIO system service calls EXE_STD\$SETCHAR and EXE_STD\$SETMODE as upper-level FDT action routines at IPL\$ASTDEL.

Description

A driver specifies the system-supplied upper-level FDT action routine EXE_STD\$SETCHAR or EXE_STD\$SETMODE to process the set-device-mode (IO\$SETMODE) and set-device-characteristics (IO\$SETCHAR) functions, respectively. If setting device characteristics requires device activity or synchronization with fork processing, the driver's FDT_ACT macro invocation *must* specify EXE_STD\$SETMODE. Otherwise, it can specify EXE_STD\$SETCHAR.

EXE_STD\$SETCHAR and EXE_STD\$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 the FDT_CONTEXT structure and transfer control to EXE_STD\$ABORTIO to terminate the request.

EXE_STD\$SETCHAR and EXE_STD\$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 the FDT_CONTEXT structure and transfer control to EXE_STD\$ABORTIO to terminate the request.

System Routines EXE_STD\$SETCHAR, EXE_STD\$SETMODE

If the request passes these checks, EXE_STD\$SETCHAR and EXE_STD\$SETMODE proceed as follows:

- EXE_STD\$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_STD\$SETCHAR writes the second word into the default-buffer-size field (UCB\$W_DEVBUFSIZ) and the third and fourth words into the device-dependent-characteristics field (UCB\$Q_DEVDEPEND).

Finally, EXE_STD\$SETCHAR stores normal completion status (SS\$_NORMAL) in the FDT_CONTEXT structure and transfers control to the FDT completion routine EXE_STD\$FINISHIO to insert the IRP in the local processor's I/O postprocessing queue. EXE_STD\$FINISHIO returns to EXE_STD\$SETCHAR with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure.

- EXE_STD\$SETMODE stores the specified quadword of characteristics in IRP\$L_MEDIA, places normal completion status (SS\$_NORMAL) in the FDT_CONTEXT structure, and transfers control to FDT completion routine EXE_STD\$QIODRVPKT to deliver the IRP to the driver's start-I/O routine. EXE_STD\$QIODRVPKT returns to EXE_STD\$SETMODE with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure.

The driver's start-I/O routine copies data from IRP\$L_MEDIA and the following longword into UCB\$W_DEVBUFSIZ, UCB\$L_DEVDEPEND, and, if the I/O function is IO\$_SETCHAR, UCB\$B_DEVCLASS and UCB\$B_DEVTYPE as well.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routines EXE\$SETCHAR and EXE\$SETMODE (used by OpenVMS VAX and Step 1 device drivers) expect as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.

R0, R7, and R8 are not provided as input to EXE_STD\$SETCHAR and EXE_STD\$SETMODE.

- EXE\$SETCHAR and EXE\$SETMODE return control to the system service dispatcher, passing it the final \$QIO system service status (SS\$_NORMAL, SSS_ACCVIO, or SSS_ILLIOFUNC) in R0. EXE_STD\$SETCHAR or EXE_STD\$SETMODE returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

System Routines

EXE_STD\$SNDEVMSG

EXE_STD\$SNDEVMSG

Builds and sends a device-specific message to the mailbox of a system process, such as the job controller or OPCOM.

Module

MBDRIVER

Format

status = EXE_STD\$SNDEVMSG (mb_ucb, msgtyp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
mb_ucb	MB_UCB	input	reference	required
msgtyp	integer	input	value	required
ucb	UCB	input	reference	required

mb_ucb

Mailbox UCB. (SYSSAR_JOBCTLMB contains the address of the job controller's mailbox; SYSSAR_OPRMBX contains the address of OPCOM's mailbox.)

msgtyp

Message type. OPCOM message types have the prefix OPC\$_ and are defined by the \$OPCMMSG macro in SYSS\$LIBRARY:STARLET.MLB.

ucb

Device UCB. EXE_STD\$SNDEVMSG reads the following UCB fields:

UCB\$W_UNIT	Device unit number.
UCB\$L_DDB	Address of device DDB. EXE_STD\$SNDEVMSG constructs the device controller name from DDB\$_NAME and mailbox UCB fields.

Return Values

SS\$_DEVNOTMBX	mb_ucb does not specify a mailbox UCB.
SS\$_INSFMEM	The system is unable to allocate memory for the message.
SS\$_MBFULL	The message mailbox is full of messages.
SS\$_MBTOOSML	The message is too large for the mailbox.
SS\$_NOPRIV	The caller lacks privilege to write to the mailbox.
SS\$_NORMAL	Normal, successful completion.

Context

Because EXE_STD\$\$SNDEVMSG raises IPL to IPL\$_MAILBOX and obtains the MAILBOX spin lock in a multiprocessing environment, its caller cannot be executing above IPL\$_MAILBOX. EXE_STD\$\$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_STD\$\$SNDEVMSG builds a 32-byte message on the stack that includes the following information:

Bytes	Contents
0 and 1	Low word of msgtyp parameter
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_STD\$\$SNDEVMSG then calls EXE_STD\$WRTMAILBOX to send the message to a mailbox.

Macro

CALL_SNDEVMSG [save_r1]

where:

save_r1 indicates that the macro should preserve register R1 across the call to COM_STD\$POST. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

In a Step 2 driver, the CALL_SNDEVMSG macro simulates a JSB to EXE\$\$SNDEVMSG in a Step 1 driver. \$\$SNDEVMSG calls EXE_STD\$\$SNDEVMSG, using the current contents of R3, R4, and R5 as the **mb_ucb**, **msgtyp**, and **ucb** arguments, respectively. It returns status in R0. Unless you specify **save_r1=NO**, the macro preserves R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$\$SNDEVMSG replaces EXE\$\$SNDEVMSG (used by Step 1 and OpenVMS VAX drivers). Unlike EXE\$\$SNDEVMSG, EXE_STD\$\$SNDEVMSG does not preserve R1 across the call.

System Routines

EXE_STD\$WRITE

EXE_STD\$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 aborts the request or proceeds with a direct-I/O, DMA read operation.

Module

SYSQIOFDT

Format

status = EXE_STD\$WRITE (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp

I/O request packet for the current I/O request.

EXE_STD\$WRITE reads the following IRP fields:

Field	Contents
IRP\$L_QIO_P1	\$QIO system service p1 argument, containing the buffer's virtual address.
IRP\$L_QIO_P2	\$QIO system service p2 argument, containing the number of bytes in transfer. The maximum number of bytes that EXE_STD\$WRITE can transfer is 65,535 (128 pages minus one byte).
IRP\$L_QIO_P4	\$QIO system service p4 argument, containing the carriage control byte.
IRP\$L_FUNC	I/O function code.
IRP\$B_RMOD	Access mode of the caller of the \$QIO system service.

EXE_STD\$WRITE writes the following IRP fields:

Field	Contents
IRP\$B_CARCON	Carriage control byte (from IRP\$L_QIO_P4)
IRP\$L_FUNC	Logical write function code converted to physical

Field	Contents
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\$\$_OBOFF	Original byte offset into the first page of a segmented direct-I/O transfer
IRP\$\$_BCNT	Size of transfer in bytes

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb

Channel control block that describes the process-I/O channel

Return Values

SS\$_FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.
----------------	---

Status in FDT_CONTEXT

SS\$_ACCVIO	Buffer specified in buf parameter does not allow read access.
SS\$_BADPARAM	bufsiz parameter is less than zero.
SS\$_INSFWSL	Insufficient working set limit.
SS\$_NORMAL	The I/O request has been successfully queued.
SS\$_QIO_CROCK	Buffer page must be faulted into memory.

Context

FDT dispatching code in the \$QIO system service calls EXE_STD\$WRITE as an upper-level FDT action routine at IPL\$_ASTDEL.

Description

A driver specifies the system-supplied upper-level FDT action routine EXE_STD\$WRITE to prepare a direct-I/O write request. A driver cannot specify EXE_STD\$WRITE for buffered-I/O functions. Drivers that process functions that require an intermediate system buffer typically supply their own upper-level FDT action routines to handle them.

EXE_STD\$WRITE performs the following functions:

- Copies the **p4** argument of the \$QIO request from IRP\$_QIO_P4 to IRP\$_CARCON
- Translates a logical write function to a physical write function and stores the new function code in IRP\$_FUNC.

System Routines

EXE_STD\$WRITE

- Examines the size of the transfer, as specified in the **p2** argument of the \$QIO request (IRP\$L_QIO_P2), and takes one of the following actions:
 - If the transfer byte count is zero, EXE_STD\$WRITE invokes the \$QIODRVPKT macro to deliver the IRP to the driver's start-I/O routine. EXE_STD\$WRITE regains control with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure. It returns to the \$QIO system service, passing these status values.

The driver start-I/O routine should check for zero-length buffers to avoid mapping to adapter node space. An attempted mapping can cause a system failure.
 - If the byte count is not zero, EXE_STD\$WRITE calls EXE_STD\$WRITELOCK, passing 0 as the value of the **err_rout** argument.

EXE_STD\$WRITELOCK invokes the \$WRITECHK macro, which calls EXE_STD\$WRITECHK.

EXE_STD\$WRITECHK performs the following actions:

- Moves the transfer byte count (**bufsiz** parameter) into IRP\$L_BCNT.

If the byte count is negative, it calls EXE_STD\$ABORTIO, passing it a **qio_sts** of SSS_BADPARAM. When it regains control, EXE_STD\$WRITECHK returns to EXE_STD\$WRITELOCK with SSS_BADPARAM status in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$WRITELOCK immediately returns to EXE_STD\$WRITE, passing these status values. EXE_STD\$WRITE, in turn, returns to the \$QIO system service.
- 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 returns SSS_NORMAL in R0 to EXE_STD\$WRITELOCK.
 - If the buffer does not allow read access, EXE_STD\$WRITECHK calls EXE_STD\$ABORTIO, passing it a **qio_sts** of SSS_ACCVIO. When it regains control, EXE_STD\$WRITECHK returns to EXE_STD\$WRITELOCK with SSS_ACCVIO status in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$WRITELOCK immediately returns to EXE_STD\$WRITE, passing these status values. EXE_STD\$WRITE returns to the \$QIO system service.

If EXE_STD\$WRITECHK succeeds, EXE_STD\$WRITELOCK moves into IRP\$L_BOFF and IRP\$L_OBOFF the byte offset to the start of the buffer and calls MMG_STD\$IOLOCK.

MMG_STD\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:

- If MMG_STD\$IOLOCK succeeds, EXE_STD\$WRITELOCK stores in IRP\$L_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns SSS_NORMAL status in R0 to EXE_STD\$WRITELOCK. EXE_STD\$WRITELOCK returns immediately to EXE_STD\$WRITE, passing to it this status value.

EXE_STD\$WRITE invokes the \$QIODRVPKT macro to deliver the IRP to the driver's start-I/O routine. EXE_STD\$WRITE regains control with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure. It returns to the \$QIO system service, passing these status values.

- If MMG_STD\$IOLOCK fails, it returns SSS_ACCVIO, SSS_INSFWSL, or page fault status to EXE_STD\$WRITELOCK.

For SSS_ACCVIO and SSS_INSFWSL status, EXE_STD\$WRITELOCK calls EXE_STD\$ABORTIO, passing it one of these status values as a **qio_sts** argument. When it regains control, EXE_STD\$WRITELOCK returns EXE_STD\$WRITE the specified status value in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$WRITE returns to the \$QIO system service.

For page fault status, EXE_STD\$WRITELOCK sets the final \$QIO status in the FDT_CONTEXT structure to SSS_QIO_CROCK and initializes FDT_CONTEXT\$SL_QIO_R1_VALUE to the virtual address to be faulted. It then adjusts the 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.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine EXE\$WRITE (used by OpenVMS VAX and Step 1 device drivers) expects as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$WRITE.
- EXE\$WRITE returns control to the system service dispatcher, passing it the final \$QIO system service status (SSS_NORMAL, SSS_ACCVIO, or SSS_BADPARAM, or SSS_INSFWSL) in R0. EXE_STD\$WRITE returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

System Routines

EXE_STD\$WRITECHK

EXE_STD\$WRITECHK

Verifies that a process has read access to the pages in the buffer specified in a \$QIO request.

Module

SYSQIOFDT

Format

status = EXE_STD\$WRITECHK (irp, pcb, ucb, buf, bufsiz)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
buf	address	input	reference	required
bufsiz	integer	input	value	required

irp

I/O request packet for the current I/O request.

EXE_STD\$WRITECHK reads IRP\$B_RMOD to determine the access mode of the caller of the \$QIO system service.

EXE_STD\$WRITECHK writes the size of the transfer in bytes to IRP\$SL_BCNT.

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

buf

Virtual address of buffer.

bufsiz

Number of bytes in transfer.

Return Values

SS\$_NORMAL	The buffer is read-accessible.
SS\$_FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$_ACCVIO	Buffer specified in buf parameter does not allow read access.
SS\$_BADPARAM	bufsiz parameter is less than zero.
SS\$_INSFWSL	Insufficient working set limit.
SS\$_NORMAL	Nothing has occurred yet to prevent the I/O request from being successfully queued. This is the initial value of the status field in an FDT_CONTEXT structure.

Context

The FDT support routine EXE_STD\$WRITELOCK, or a driver-specific FDT routine, calls EXE_STD\$WRITECHK at IPL\$ASTDEL.

Description

A driver FDT routine calls the system-supplied FDT support routine EXE_STD\$WRITECHK to check the read accessibility of an I/O buffer supplied in a \$QIO request for a write function.

EXE_STD\$WRITECHK performs the following actions:

- Moves the transfer byte count (**bufsiz** parameter) into IRP\$L_BCNT. If the byte count is negative, it calls EXE_STD\$ABORTIO, passing it a **qio_sts** of SS\$_BADPARAM. When it regains control, EXE_STD\$WRITECHK returns to its caller with SS\$_BADPARAM status in the FDT_CONTEXT structure and SS\$_FDT_COMPL status in R0.
- 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_STD\$WRITECHK returns SS\$_NORMAL in R0 to its caller.
 - If the buffer does not allow read access, EXE_STD\$WRITECHK calls EXE_STD\$ABORTIO, passing it a **qio_sts** of SS\$_ACCVIO. When it regains control, EXE_STD\$WRITECHK returns to its caller with SS\$_ACCVIO status in the FDT_CONTEXT structure and SS\$_FDT_COMPL status in R0.

The caller of EXE_STD\$WRITECHK must examine the status in R0:

- If the status is SS\$_NORMAL, the buffer is read-accessible.
- If the status is SS\$_FDT_COMPL, an error has occurred that has caused the I/O request to be aborted. You can determine the reason for the failure from FDT_CONTEXT\$L_QIO_STATUS.

System Routines

EXE_STD\$WRITECHK

Certain drivers must perform additional processing to back out an I/O request after it has aborted. For instance, if the driver has locked multiple buffers into memory for a single I/O request, it must unlock them once the request has been aborted. Note that a driver cannot access the IRP once it has received `SS$_FDT_COMPL` status. If you know you need access to information stored in the IRP to back out an I/O request that has been aborted, you must store that information elsewhere prior to calling `EXE_STD$WRITELOCK`.

Macro

`CALL_WRITECHK`
`CALL_WRITECHKR`

In a Step 2 driver, `CALL_WRITECHK` simulates a JSB to `EXE$WRITECHK` and `CALL_READCHKR` simulates a JSB to `EXE$READCHKR`. Both macros call `EXE_STD$READCHK` using the current contents of R3, R4, R5, R0, and R1 as the **irp**, **pcb**, **ucb**, **buf**, and **bufsiz** arguments, respectively.

When `EXE_STD$WRITECHK` returns, `CALL_WRITECHK` and `CALL_WRITECHKR` clear R2 to indicate a write operation and examines the return status:

- If success status (`SS$_NORMAL`) is returned, `CALL_WRITECHK` and `CALL_WRITECHKR` copy the contents of `IRP$SL_BCNT` into R1. `CALL_WRITECHK` writes the starting address of the I/O buffer in R0; `CALL_WRITECHKR` preserves the return status value in R0.
- If failure status (`SS$_FDT_COMPL`) is returned, `CALL_WRITECHK` returns to FDT dispatching code in the `$QIO` system service. `CALL_WRITECHKR` does not return control to `$QIO`.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- `EXE_STD$WRITECHK` replaces `EXE$WRITECHK` and `EXE$WRITECHKR` (used by Step 1 drivers). For compatibility with the Step 1 routines, use the `CALL_WRITECHK` and `CALL_WRITECHKR` macros.
- `EXE$WRITECHK` and `EXE$WRITECHKR` expect as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to `EXE_STD$WRITECHK`.
- The order in which formal parameters are passed to `EXE_STD$WRITECHK` differs from the order in which they are provided in registers to the Step 1 routines `EXE$WRITECHK` and `EXE$WRITECHKR`.
- `EXE$WRITECHK` and `EXE$WRITECHKR` provide a mechanism by which a driver callback routine or coroutine obtains control upon an error condition prior to the abortion of an I/O request. The design of FDT processing for Step 2 device drivers guarantees that the caller of `EXE_STD$WRITECHK` regains control whether the write check operation is successful. The caller must examine the return status in R0 (`SS$_NORMAL` indicates the buffer is read accessible, `SS$_FDT_COMPL` indicates that the operation failed and the request has been aborted) and take appropriate action. Final `$QIO`

completion status, indicating the reason the operation failed, is stored in the `FDT_CONTEXT` structure.

- Driver code that services failure status (`SS$FDT_COMPL`) from `EXE_STD$WRITELOCK` (for instance, a callback routine formerly specified to `EXE$WRITELOCK_ERR`) cannot access information in the IRP. If the driver anticipates handling failure status by using the contents of IRP fields, it must store these fields elsewhere before calling `EXE_STD$WRITELOCK`.

This is especially important for driver code that expects `EXE_STD$WRITECHK` to access the transfer size in `R1` after the call. Unlike `EXE$WRITECHK` and `EXE$WRITECHKR`, `EXE_STD$WRITECHK` does not preserve the contents of `R1` and `R3` across the call. If you must repeat a `CALL_WRITECHK` macro invocation, be sure to reload `R0`, `R1`, and `R3` with the virtual address of the buffer, the transfer size, and the address of the IRP, respectively, before each subsequent invocation.

- Upon successful completion, `EXE$WRITECHK` and `EXE$WRITECHKR` clear `R2` to indicate a write function. `EXE_STD$WRITECHK` does not provide `R2` as output; a driver can determine whether a function is write or read by examining `IRPSV_FUNC` in `IRPSL_STS`.

EXE_STD\$WRITELOCK

Validates and prepares a user buffer for a direct-I/O, DMA read operation.

Module

SYSQIOFDT

Format

status = EXE_STD\$WRITELOCK (irp, pcb, ucb, ccb, buf, bufsiz, err_rout)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required
buf	address	input	reference	required
bufsiz	integer	input	value	required
err_rout	procedure value	input	value	required

irp

I/O request packet for the current I/O request.

EXE_STD\$WRITELOCK reads IRP\$B_RMOD to determine the access mode of the caller of the \$QIO system service.

EXE_STD\$WRITELOCK writes the following IRP fields:

Field	Contents
IRP\$L_SVAPTE	System virtual address of the PTE that maps the first page of the buffer
IRP\$L_BOFF	Byte offset to start of transfer in page
IRP\$L_OBOFF	Original byte offset into the first page of a segmented direct-I/O transfer
IRP\$L_BCNT	Size of transfer in bytes

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb

Channel control block that describes the process-I/O channel.

buf

Virtual address of buffer.

bufsiz

Number of bytes in transfer.

err_rout

Procedure value of error-handling callback routine, or 0 if the driver does not process errors.

A driver typically specifies an error-handling callback routine when it must lock multiple areas into memory for a single I/O request and must regain control to unlock these areas, if the request is to be aborted. The routine performs those tasks required before the request is backed out of or aborted. Such operations could include calling MMG_STD\$UNLOCK to release previous buffers participating in the I/O operation. The error-handling routine must preserve R0 and R1 and return back to EXE_STD\$WRITELOCK.

Chapter 1 describes the error-handling callback routine interface.

Return Values

SS\$NORMAL	The buffer is read-accessible and has been locked in memory.
SS\$FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$ACCVIO	Buffer specified in buf parameter does not allow read access.
SS\$BADPARAM	bufsiz parameter is less than zero.
SS\$INSFWSL	Insufficient working set limit.
SS\$NORMAL	Nothing has occurred yet to prevent the I/O request from being successfully queued. This is the initial value of the status field in an FDT_CONTEXT structure.
SS\$INSFWSL	Insufficient working set limit.
SS\$QIO_CROCK	Buffer page must be faulted into memory.

Context

The system-supplied upper-level FDT action routine EXE_STD\$WRITE, or a driver-specific upper-level FDT action routine, calls EXE_STD\$WRITELOCK at IPL\$ASTDEL.

System Routines

EXE_STD\$WRITELOCK

Description

A driver FDT routine calls the system-supplied FDT support routine EXE_STD\$WRITELOCK to check the read accessibility of an I/O buffer supplied in a \$QIO request for a write function, and lock the buffer in memory in preparation for a DMA write operation.

A driver cannot specify EXE_STD\$WRITE for buffered-I/O functions. Drivers that process functions that require an intermediate system buffer typically supply their FDT routines to handle them.

EXE_STD\$WRITELOCK invokes the \$WRITECHK macro, which calls EXE_STD\$WRITECHK.

EXE_STD\$WRITECHK performs the following actions:

- Moves the transfer byte count (**bufsiz** parameter) into IRP\$LCB_CNT. If the byte count is negative, EXE_STD\$WRITECHK returns SSS_BADPARAM status to EXE_STD\$READLOCK.
- 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_STD\$WRITECHK returns SSS_NORMAL in R0 to EXE_STD\$WRITELOCK.
 - If the buffer does not allow write access, EXE_STD\$READCHK returns SSS_ACCVIO status to EXE_STD\$READLOCK.

If error status (SSS_BADPARAM or SSS_ACCVIO) is returned, EXE_STD\$WRITELOCK immediately calls the specified error-handling callback routine, passing to it the IRP, PCB, UCB, CCB, and status value. The callback routine must preserve R0 and R1 and return control to EXE_STD\$WRITELOCK. When the callback routine returns (or if no callback routine is specified), EXE_STD\$WRITELOCK calls EXE_STD\$ABORTIO, passing it the error status as **qio_sts**. EXE_STD\$ABORTIO returns to EXE_STD\$WRITELOCK with the error status in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0. EXE_STD\$WRITELOCK immediately returns to its caller, passing these status values.

If SSS_NORMAL status is returned, EXE_STD\$WRITELOCK moves into IRP\$LCB_BOFF and IRP\$LCB_OBOFF the byte offset to the start of the buffer and calls MMG_STD\$IOLOCK.

MMG_STD\$IOLOCK attempts to lock into memory those pages that contain the buffer, with one of the following results:

- If MMG_STD\$IOLOCK succeeds, EXE_STD\$WRITELOCK stores in IRP\$LCB_SVAPTE the system virtual address of the process PTE that maps the first page of the buffer, and returns SSS_NORMAL status in R0 to EXE_STD\$WRITELOCK. EXE_STD\$WRITELOCK returns immediately to its caller, passing to it this status value.
- If MMG_STD\$IOLOCK fails, it returns SSS_ACCVIO, SSS_INSFWSL, or page fault status to EXE_STD\$WRITELOCK. EXE_STD\$WRITELOCK immediately calls the specified error-handling callback routine, passing to it the IRP, PCB, UCB, CCB, and status value. The callback routine must preserve R0 and R1 and return control to EXE_STD\$WRITELOCK. When

the callback routine returns (or if no callback routine is specified), EXE_STD\$WRITELOCK proceeds as follows:

- For SSS_ACCVIO and SSS_INSFWSL status, EXE_STD\$WRITELOCK calls EXE_STD\$ABORTIO, passing it one of these status values as a **qio_sts** argument. When it regains control, EXE_STD\$WRITELOCK returns to its caller the specified status value in the FDT_CONTEXT structure and SSS_FDT_COMPL status in R0.
- For page fault status, EXE_STD\$WRITELOCK sets the final \$QIO status in the FDT_CONTEXT structure to SSS_QIO_CROCK and initializes FDT_CONTEXT\$SL_QIO_R1_VALUE to the virtual address to be faulted. It then adjusts the 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.

The caller of EXE_STD\$WRITELOCK must examine the status in R0:

- If the status is SSS_NORMAL, the buffer is write accessible and has been successfully locked into memory and the starting virtual address of the page table entries that map the buffer is available in IRP\$SL_SVAPTE.
- If the status is SSS_FDT_COMPL, an error has occurred that has caused the I/O request to be aborted. You can determine the reason for the failure from FDT_CONTEXT\$SL_QIO_STATUS. Ordinarily a driver specifies an error-handling callback routine to process such errors.

Note that a driver cannot access the IRP once it has received SSS_FDT_COMPL status. If you know you need access to information stored in the IRP to back out an I/O request that has been aborted, you must store that information elsewhere prior to calling EXE_STD\$WRITELOCK.

Macro

```
CALL_WRITELOCK  
CALL_WRITELOCK_ERR [interface_warning=YES]
```

where:

interface_warning=YES, the default, specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. **interface_warning=NO** suppresses the warning.

In a Step 2 driver, the CALL_WRITELOCK simulates a JSB to EXE\$WRITELOCK and CALL_WRITELOCK_ERR simulates a JSB to EXE\$WRITELOCK_ERR. CALL_WRITELOCK calls EXE_STD\$WRITELOCK, specifying 0 as the **err_rout** argument; CALL_WRITELOCK_ERR also calls EXE_STD\$WRITELOCK, using the contents of R2 as the **err_rout** argument. Both macros supply the current contents of R3, R4, R5, R6, R0, and R1 as the **irp**, **pcb**, **ucb**, **ccb**, **buf**, and **bufsiz** arguments, respectively.

System Routines

EXE_STD\$WRITELOCK

When EXE_STD\$WRITELOCK or EXE_STD\$WRITELOCK_ERR returns, code generated by the macro examines the return status:

- If success status (SS\$_NORMAL) is returned, the macro moves the contents of IRP\$\$_SVAPTE into R1 and clears R2 to indicate a write operation. Status is returned in R0 and in the FDT_CONTEXT structure.
- If failure status (SS\$_FDT_COMPL) is returned, the macro clears R2 to indicate a write operation and returns to FDT dispatching code in the \$QIO system service.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$WRITELOCK replaces EXE\$WRITELOCK and EXE\$WRITELOCK_ERR (used by Step 1 drivers). For compatibility with the Step 1 routines, use the CALL_WRITELOCK and CALL_WRITELOCK_ERR macros.
- EXE\$WRITELOCK and EXE\$WRITELOCK_ERR expect as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$WRITELOCK.
- The order in which formal parameters are passed to EXE_STD\$WRITELOCK differs from the order in which they are provided in registers to the Step 1 routines EXE\$WRITELOCK and EXE\$WRITELOCK_ERR.
- EXE\$WRITELOCK_ERR provides a mechanism by which a driver callback routine or coroutine obtains control upon an error condition prior to the abortion of an I/O request. EXE_STD\$WRITELOCK accepts the address of an error-handling callback routine in the **err_rout** argument. The error-handling routine is called after an I/O request encounters a buffer access or memory allocation failure and before the request is aborted.
- The design of FDT processing for Step 2 device drivers guarantees that the caller of EXE_STD\$WRITELOCK regains control whether the write lock operation is successful. When a driver regains control from a call to EXE_STD\$WRITELOCK, return status in R0 indicates that the buffer has been successfully locked (SS\$_NORMAL) or that the operation failed and the request has been aborted (SS\$_FDT_COMPL). The driver must check the return status and take appropriate action. Final \$QIO completion status, indicating the reason the operation failed, is stored in the FDT_CONTEXT structure.

Normally, a driver services a read lock failure by supplying the address of an error-handling callback routine to EXE_STD\$WRITELOCK.

- Driver code that executes after receiving failure status (SS\$_FDT_COMPL) from EXE_STD\$WRITELOCK cannot access information in the IRP. If the driver anticipates accessing IRP fields when EXE_STD\$WRITELOCK returns, it must store these fields elsewhere before calling EXE_STD\$WRITELOCK.

System Routines EXE_STD\$WRITELOCK

- Upon successful completion, EXE\$WRITELOCK and EXE\$WRITELOCK_ERR provide as output the system virtual address of the first process PTE that maps the buffer in R1 and in IRP\$L_SVAPTE. Because EXE_STD\$WRITELOCK does not provide R1 as output, a driver must obtain this information from IRP\$L_SVAPTE. Similarly, the Step 1 routines clear R2 for a write function. EXE_STD\$WRITELOCK does not provide R2 as output; a driver can determine whether a function is write or read by examining IRP\$V_FUNC in IRP\$L_STS.

System Routines

EXE_STD\$WRTMAILBOX

EXE_STD\$WRTMAILBOX

Sends a message to a mailbox.

Module

MBDRIVER

Format

status = EXE_STD\$WRTMAILBOX (mb_ucb, msgsiz, msg)

Arguments

Argument	Type	Access	Mechanism	Status
mb_ucb	MB_UCB	input	reference	required
msgsiz	integer	input	value	required
msg	address	input	reference	required

mb_ucb

Mailbox UCB. (SYSSAR_JOBCTLMB contains the address of the job controller's mailbox; SYSSAR_OPRMBX contains the address of OPCOM's mailbox.)

msgsiz

Message size.

msg

Address of buffer containing the message.

Return Values

SS\$INSFMEM	The system is unable to allocate memory for the message.
SS\$MBFULL	The message mailbox is full of messages.
SS\$MBTOOSML	The message is too large for the mailbox.
SS\$NOPRIV	The caller lacks privilege to write to the mailbox.
SS\$NORMAL	Normal, successful completion.

Context

Because EXE_STD\$WRTMAILBOX raises IPL to IPL\$MAILBOX and obtains the MAILBOX spin lock in a multiprocessing environment, its caller cannot be executing above IPL\$MAILBOX. EXE_STD\$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_STD\$WRTRMAILBOX checks fields in the mailbox UCB (UCB\$W_MSGQUO, UCB\$W_DEVMSGISZ) 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 EXESALONONPAGED to allocate a block of nonpaged pool to contain the message. If it fails any of these operations, EXE_STD\$WRTRMAILBOX returns error status to its caller.

If it is successful thus far, EXE_STD\$WRTRMAILBOX creates a message and delivers it to the mailbox's message queue, adjusts its UCB fields accordingly, and returns success status to its caller.

Macro

CALL_WRTMAILBOX [save_r1]

where:

save_r1 indicates that the macro should preserve register R1 across the call to COM_STD\$POST. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

In a Step 2 driver, the CALL_WRTMAILBOX macro simulates a JSB to EXE\$WRTRMAILBOX in a Step 1 driver. CALL_WRTMAILBOX calls EXE_STD\$WRTRMAILBOX, using the current contents of R5, R3, and R4 as the **mb_ucb**, **msgsiz**, and **msg** arguments, respectively. It returns status in R0. Unless you specify **save_r1=NO**, the macro preserves the R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- EXE_STD\$WRTRMAILBOX replaces EXE\$WRTRMAILBOX (used by Step 1 and OpenVMS VAX drivers). The order in which formal parameters are passed to EXE_STD\$WRTRMAILBOX differs from the order in which they are provided in registers to the Step 1 routine EXE\$WRTRMAILBOX.
- Unlike EXE\$WRTRMAILBOX, EXE_STD\$WRTRMAILBOX does not preserve R1 across the call.

System Routines

EXE_STD\$ZEROPARM

EXE_STD\$ZEROPARM

Delivers an I/O request that requires no parameters to a driver's start-I/O routine.

Module

SYSQIOFDT

Format

status = EXE_STD\$ZEROPARM (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp

I/O request packet for the current I/O request. EXE_STD\$ZEROPARM clears IRP\$L_MEDIA.

pcb

Process control block of the current process.

ucb

Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb

Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL

Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$NORMAL

The routine completed successfully.

Context

FDT dispatching code in the \$QIO system service calls EXE_STD\$ZEROPARM as an upper-level FDT action routine at IPL\$ASTDEL.

Description

A driver specifies the system-supplied upper-level FDT action routine EXE_STD\$ZEROPARM to process an I/O function code that has no required parameters.

EXE_STD\$ZEROPARM clears IRP\$L_MEDIA and invokes the \$QIODRVPKT macro to deliver the IRP to the driver. EXE_STD\$ZEROPARM regains control with SSS_FDT_COMPL status in R0 and a final \$QIO system service status of SSS_NORMAL in the FDT_CONTEXT structure.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine EXE\$ZEROPARM (used by OpenVMS VAX and Step 1 device drivers) expects as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to EXE_STD\$ZEROPARM.
- EXE\$ZEROPARM returns control to the system service dispatcher, passing it the final \$QIO system service status (SSS_NORMAL) in R0. EXE_STD\$ZEROPARM returns to its caller, passing it SSS_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

System Routines

IOC\$ALOALTMAP, IOC\$ALOALTMAPN, IOC\$ALOALTMAPSP

IOC\$ALOALTMAP, IOC\$ALOALTMAPN, IOC\$ALOALTMAPSP

Allocate a set of Q22-bus alternate map registers.

Notes for Converting Step 1 Drivers

Not supported on OpenVMS AXP systems. See the description of IOC\$ALLOC_CNT_RES.

IOC\$ALOUBAMAP, IOC\$ALOUBAMAPN

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

Notes for Converting Step 1 Drivers

Not supported on OpenVMS AXP systems. See the description of IOC\$ALLOC_CNT_RES.

System Routines

IOC\$ALLOC_CNT_RES

IOC\$ALLOC_CNT_RES

Allocates the requested number of items of a counted resource.

Module

ALLOC_CNT_RES

Format

IOC\$ALLOC_CNT_RES crab ,crctx

Context

IOC\$ALLOC_CNT_RES conforms to the OpenVMS AXP calling standard. Its caller must be executing at fork IPL, holding the corresponding fork lock.

Arguments

crab

VMS Usage: address
type: longword (signed)
access: read only
mechanism: by reference

Address of CRAB that describes the counted resource. For adapters that supply a counted resource, such as map registers, ADPSL_CRAB often contains this address.

crctx

VMS Usage: address
type: longword (signed)
access: read only
mechanism: by reference

Address of CRCTX structure that describes the request for the counted resource.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL The routine completed successfully.

SS\$_BADPARAM	Request count was greater than the total number of items managed by the CRAB or the total number of items defined by a bounded request. This status is also returned if the lower bound of the request (CRCTX\$LOW_BOUND) is greater than the upper bound (CRCTX\$UP_BOUND).
SS\$_INSFMAPREG	Insufficient resources to satisfy request, or other requests precede this one in the resource-wait queue.

Description

IOC\$ALLOC_CNT_RES allocates a requested number of items from a counted resource. The resource request is described in the CRCTX structure; the counted resource itself is described in the CRAB.

A driver typically initializes the following fields of the CRCTX before submitting it in a call to IOC\$ALLOC_CNT_RES.

Field	Description
CRCTX\$ITEM_CNT	Number of items to be allocated. When requesting map registers, this value in this field should include an extra map register to be allocated and loaded as a guard page to prevent runaway transfers.
CRCTX\$CALLBACK	Procedure value of the callback routine to be called when the deallocation of resource items allows a stalled resource request to be granted. A value of 0 in this field indicates that, on an allocation failure, control should return to the caller immediately without queueing the CRCTX to the CRAM's wait queue.

A caller can also specify the upper and lower bounds of the search for allocatable resource items by supplying values for CRCTX\$LOW_BOUND and CRCTX\$UP_BOUND.

IOC\$ALLOC_CNT_RES performs the following tasks:

- It acquires the spin lock indicated by CRAB\$SPINLOCK, raising IPL to IPL\$SYNCH in the process.
- If there are no waiters for the counted resource (that is, the resource wait queue headed by CRAB\$WQFL is empty) or if the CRCTX describes a high-priority allocation request (CRCTX\$HIGH_PRIO in CRCTX\$FLAGS is set), IOC\$ALLOC_CNT_RES attempts the allocation immediately. It scans the CRAB allocation array for a descriptor that contains as many free items as requested by the caller (in CRCTX\$ITEM_CNT).

In performing the scan, IOC\$ALLOC_CNT_RES considers any indicated range of counted resource items that are to be involved in the scan, and limits its search to those item descriptors in the allocation array that describe items within these bounds. A bounded search is indicated by nonzero values in CRCTX\$UP_BOUND and CRCTX\$LOW_BOUND. IOC\$ALLOC_CNT_RES rounds up the allocation request to the minimal allocation granularity, as indicated by CRAB\$ALLOC_GRAN_MASK.

System Routines

IOC\$ALLOC_CNT_RES

The number of the first resource item granted to the caller is placed in CRCTX\$SL_ITEM_NUM and CRCTX\$V_ITEM_VALID is set in CRCTX\$SL_FLAGS.

- If this allocation attempt fails, saves the current values of R3, R4, and R5 in the CRCTX fork block. IOC\$ALLOC_CNT_RES writes a -1 to CRCTX\$SL_ITEM_NUM, and inserts the CRCTX in the resource-wait queue (headed by CRAB\$SL_WQFL). It then returns SSS\$INSFMAPREG status to its caller.

Note

If a counted resource request does not specify a callback routine (CRCTX\$SL_CALLBACK), IOC\$ALLOC_CNT_RES does not insert its CRCTX in the resource-wait queue. Rather, it returns SSS\$INSFMAPREG status to its caller.

When a counted resource deallocation occurs, the CRCTX is removed from the wait queue and the allocation is attempted again.

When the allocation succeeds, IOC\$ALLOC_CNT_RES issues a JSB instruction to the callback routine (CRCTX\$SL_CALLBACK), passing it the following values:

Location	Contents
R0	SSS\$NORMAL
R1	Address of CRAB
R2	Address of CRCTX
R3	Contents of R3 at the time of the original allocation request (CRCTX\$Q_FR3)
R4	Contents of R4 at the time of the original allocation request (CRCTX\$Q_FR4)
R5	Contents of R5 at the time of the original allocation request (CRCTX\$Q_FR5)
Other registers	Destroyed

The callback routine checks R0 to determine whether it has been called with SSS\$NORMAL or SSS\$CANCEL status (from IOC\$CANCEL_CNT_RES). If the former, it typically proceeds to loads the map registers that have been allocated. It must preserve all registers it uses other than R0 through R5 and exit with an RSB instruction.

- It releases the spin lock indicated by CRAB\$SL_SPINLOCK (upon the condition that its caller did not already own that spin lock at the time of the call) and returns to its caller.

OpenVMS AXP allows you to indicate that a counted resource request should take precedence over any waiting request by setting the CRCTX\$V_HIGH_PRIO bit in CRCTX\$SL_FLAGS. A driver uses a high-priority counted resource request to preempt normal I/O activity and service some exception condition from the device. (For instance, during a multivolume backup, a tape driver might make a high-priority request, when it encounters the end-of-tape marker, to get a subsequent tape loaded before normal I/O activity to the tape can resume. A disk driver might issue a high-priority request to service a disk offline condition.)

System Routines IOC\$ALLOC_CNT_RES

IOC\$ALLOC_CNT_RES never stalls a high-priority counted resource request or places its CRCTX in a resource-wait queue. Rather, it attempts to allocate the requested number of resource items immediately. If IOC\$ALLOC_CNT_RES cannot grant the requested number of items, it returns SSS_INSFMAPREG status to its caller.

IOC\$ALLOC_CRAB

Allocates and initializes a counted resource allocation block (CRAB).

Module

ALLOC_CNT_RES

Format

IOC\$ALLOC_CRAB item_cnt ,req_alloc_gran ,crab_ref

Context

IOC\$ALLOC_CRAB conforms to the OpenVMS AXP calling standard. Because IOC\$ALLOC_CRAB calls EXE\$ALONONPAGED to allocate sufficient memory for a CRAB, its caller cannot be executing above IPL\$_POOL.

Arguments

item_cnt

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

Number of items associated with the resource.

req_alloc_gran

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

Requested allocation granularity associated with the resource.

crab_ref

VMS Usage: address
type: longword (unsigned)
access: write only
mechanism: by reference

Address of a cell to which IOC\$ALLOC_CRAB returns the address of the allocated CRAB.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$_BADPARAM	Specified allocation granularity is larger than the specified item count.
SS\$_NORMAL	The routine completed successfully.
SS\$_INSFMEM	Memory allocation request failed.

Description

A driver calls IOC\$ALLOC_CRAB to allocate a counted resource allocation block (CRAB) that describes a counted resource. A counted resource, such as a set of map registers, has the following attributes:

- The resource consists of an ordered set of items.
- The allocator can request one or more items. When requesting multiple items, the requester expects to receive a contiguous set of items. Thus, allocated items can be described by a starting number and a count.
- Allocation and deallocation of the resource are common operations and, thus, must be efficient and quick.
- A single deallocation may allow zero or more stalled allocation requests to proceed.

IOC\$ALLOC_CRAB computes the size of the CRAB as the sum of the fixed portion of the CRAB, plus the maximum number of descriptors required in the allocation array. It then calls EXE\$ALONONPAGED to allocate the CRAB. If the allocation request succeeds, IOC\$ALLOC_CRAB initializes the CRAB as follows and returns SS\$_NORMAL to its caller:

Field	Description
CRAB\$W_SIZE	Size of the CRAB in bytes
CRAB\$B_TYPE	DYN\$C_MISC
CRAB\$B_SUBTYPE	DYN\$C_CRAB
CRAB\$L_WQFL	CRAB\$L_WQFL
CRAB\$L_WQBL	CRAB\$L_WQFL
CRAB\$L_TOTAL_ITEMS	Contents of the item_cnt argument
CRAB\$L_ALLOC_GRAN_MASK	One less than the contents of the req_alloc_gran argument (rounded up to the next highest power of two if the value specified is not a power of two)
CRAB\$L_VALID_DESC_CNT	1
CRAB\$L_SPINLOCK	Address of dynamic spin lock used to synchronize access to this CRAB. Currently, CRAB spin locks are obtained at IPL\$IOLCK8.

IOC\$ALLOC_CRAB initializes the first descriptor in the allocation array to indicate a set of **item_cnt** items of the resource, starting at item 0.

IOC\$ALLOC_CRCTX

Allocates and initializes a counted resource context block (CRCTX).

Module

ALLOC_CNT_RES

Format

IOC\$ALLOC_CRCTX crab ,crctx_ref ,fleck_index

Context

IOC\$ALLOC_CRCTX conforms to the OpenVMS AXP calling standard. Because IOC\$ALLOC_CRCTX calls EXE\$ALONONPAGED to allocate sufficient memory for a CRCTX, its caller cannot be executing above IPL\$_POOL.

Arguments

crab

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of CRAB that describes the counted resource. For adapters that supply a counted resource, such as map registers, ADPSL_CRAB often contains this address.

crctx_ref

VMS Usage: address
type: longword (unsigned)
access: write only
mechanism: by reference

Address of a location in which IOC\$ALLOC_CRCTX places the address of the allocated CRCTX.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$_NORMAL	The routine completed successfully.
SS\$_INSMEM	Memory allocation request failed.

Description

A driver calls IOC\$ALLOC_CRCTX to allocate a CRCTX to describe a specific request for a given counted resource, such as a set of map registers. The driver subsequently uses the CRCTX as input to IOC\$ALLOC_CNT_RES to allocate a given set of the objects managed as a counted resource.

IOC\$ALLOC_CRCTX calls EXE\$ALONONPAGED to allocate the CRCTX. If the allocation request succeeds, IOC\$ALLOC_CRCTX initializes the CRCTX as follows and returns SSS_NORMAL to its caller:

Field	Description
CRCTX\$W_SIZE	Size of the CRCTX in bytes
CRCTX\$B_TYPE	DYN\$C_MISC
CRCTX\$B_SUBTYPE	DYN\$C_CRCTX
CRCTX\$L_CRAB	Address of CRAB as specified in the crab argument
CRCTX\$W_FSIZE	FKBSK_LENGTH
CRCTX\$B_FTYPE	DYN\$C_FRK
CRCTX\$B_FLCK	IPL\$IOLOCK8

IOC\$ALLOCATE_CRAM

Allocates a controller register access mailbox.

Module

CRAM-ALLOC

Macro

DPTAB (**ucb_crams** and **idb_crams** arguments) CRAM_ALLOC

Format

IOC\$ALLOCATE_CRAM cram [,idb] [,ucb] [,adp]

Context

IOC\$ALLOCATE_CRAM conforms to the OpenVMS AXP calling standard. Because IOC\$ALLOCATE_CRAM may need to allocate pages from the free page list, its caller must be executing at or below IPL\$_SYNCH and must not hold spin locks ranked higher than IO_MISC.

IOC\$ALLOCATE_CRAM acquires and releases the IO_MISC spin lock and returns to its caller at its caller's IPL.

Arguments

cram

VMS Usage: address
type: longword (unsigned)
access: write only
mechanism: by reference

Address of CRAM allocated by IOC\$ALLOCATE_CRAM

idb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of IDB for device.

ucb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of UCB for device.

adp

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of ADP for device.

Returns

VMS Usage: cond_value
 type: longword_unsigned
 access: longword (unsigned)
 mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$_NORMAL	CRAM has been successfully allocated.
SS\$_INSFARG	Insufficient arguments supplied in call

Description

IOC\$ALLOCATE_CRAM allocates a single controller register access mailbox (CRAM) and fills in the following fields:

CRAM\$W_SIZE	Size of CRAM
CRAM\$B_TYPE	Structure type (DYN\$C_MISC)
CRAM\$B_SUBTYPE	Structure type (DYN\$C_CRAM)
CRAM\$Q_RBADR	Address of remote tightly-coupled I/O interconnect (from IDB\$Q_CSR)
CRAM\$Q_HW_MBX	Physical address of hardware I/O mailbox
CRAM\$L_MBPR	Mailbox pointer register (from ADP\$PS_MBPR)
CRAM\$Q_QUEUE_TIME	Default mailbox queue timeout value (from ADP\$Q_QUEUE_TIME)
CRAM\$Q_WAIT_TIME	Default mailbox wait-for-completion timeout value (from ADP\$Q_WAIT_TIME)
CRAM\$B_HOSE	Number of remote tightly-coupled I/O interconnect (from ADP\$B_HOSE_NUM)
CRAM\$L_IDB	IDB address
CRAM\$L_UCB	UCB address

A driver may choose to allocate a CRAM on a per-controller or a per-unit basis. Typically a driver specifies values in the **idb_cramps** and **ucb_cramps** arguments of the DPTAB macro that indicate how many CRAMs should be allocated to a controller (IDB) or a unit (UCB). If these values (DPT\$W_IDB_CRAMS and DPT\$W_UCB_CRAMS) are nonzero in the DPT, the driver loading procedure automatically invokes IOC\$ALLOCATE_CRAM to allocate the specified number of CRAMs. The driver-loading procedure thereafter sets up IDB\$PS_CRAM to point to a linked list of CRAMs associated with a controller, UCB\$PS_CRAM to a linked list of CRAMs associated with a device unit.

IOC\$CANCEL_CNT_RES

Cancels a thread that has been stalled waiting for a counted resource.

Module

ALLOC_CNT_RES

Format

IOC\$CANCEL_CNT_RES crab ,crctx [,resume_flag]

Context

IOC\$CANCEL_CNT_RES conforms to the OpenVMS AXP calling standard. Its caller must be executing at fork IPL, holding the corresponding fork lock.

Arguments

crab

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of CRAB that describes the counted resource. For adapters that supply a counted resource, such as map registers, ADPSL_CRAB often contains this address.

crctx

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of CRCTX structure that describes the request for the counted resource.

[resume_flag]

VMS Usage: boolean
type: longword (unsigned)
access: read only
mechanism: by value

Indication of whether the cancelled thread should be resumed. If true, IOC\$CANCEL_CNT_RES calls the driver callback routine with SSS_CANCEL status. If not specified or false, IOC\$CANCEL_CNT_RES does not resume the cancelled thread.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$BADPARAM	The specified CRCTX was not found in the CRAB wait queue.

Description

IOC\$CANCEL_CNT_RES cancels a thread that has been stalled waiting for a counted resource. The resource request is described in the CRCTX structure; the counted resource itself is described in the CRAB.

IOC\$CANCEL_CNT_RES scans the CRAB wait queue (CRAB\$WFQL) to locate the specified CRCTX. If it cannot locate the CRCTX, it returns SS\$BADPARAM status to its caller.

If it locates the CRCTX in the CRAB wait queue and the **resume_flag** argument is not specified or is false, it removes the CRCTX from the queue and returns SS\$NORMAL status to its caller. Otherwise, after removing the CRCTX, it issues a JSB to the driver's callback routine (CRCTX\$CALLBACK), passing it the following values:

Location	Contents
R0	SS\$CANCEL
R1	Address of CRAB
R2	Address of CRCTX
R3	CRCTX\$Q_FR3
R4	CRCTX\$Q_FR4
R5	CRCTX\$Q_FR5
Other registers	Destroyed

The callback routine checks R0 to determine whether it has been called with SS\$NORMAL (from IOC\$ALLOC_CNT_RES) or SS\$CANCEL status. If the latter, it takes appropriate steps to respond to the request cancellation. It must preserve all registers it uses other than R0 through R5 and exit with an RSB instruction.

When it regains control from the driver callback routine, IOC\$CANCEL_CNT_RES returns SS\$NORMAL status to its caller.

IOC\$CRAM_CMD

Generates values for the command, mask, and remote I/O interconnect address fields of the hardware I/O mailbox that are specific to the interconnect that is the target of the mailbox operation, inserting these values into the indicated mailbox, buffer, or both.

Module

[CPU_{xxxx}]IO_SUPPORT__{xxxx}†

Macro

CRAM_CMD

Format

IOC\$CRAM_CMD cmd_index ,byte_offset ,adp_ptr [,cram_ptr] [,buffer_ptr]

Context

IOC\$CRAM_CMD conforms to the OpenVMS AXP calling standard. It acquires no spin locks and leaves IPL unchanged. After inserting the hardware I/O mailbox values into the CRAM or specified buffer, IOC\$CRAM_CMD returns to its caller.

Arguments

cmd_index

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

Command index. IOC\$CRAM_CMD uses this index to generate a mailbox command that is specific to the tightly-coupled interconnect that is to be the target of a request using this CRAM.

You can specify any of the following values (defined by the \$CRAMDEF macro), although which of these I/O operations is supported depends on the I/O interconnect that is to be the object of the mailbox operation.

Command Index	Description
CRAMCMD\$K_RDQUAD32	Quadword read in 32-bit space
CRAMCMD\$K_RDLONG32	Longword read in 32-bit space
CRAMCMD\$K_RDWORD32	Word read in 32-bit space
CRAMCMD\$K_RDBYTE32	Byte read in 32-bit space
CRAMCMD\$K_WTQUAD32	Quadword write in 32-bit space
CRAMCMD\$K_WTLONG32	Longword write in 32-bit space
CRAMCMD\$K_WTWORD32	Word write in 32-bit space

† where *xxxx* represents the internal OpenVMS code number for an AXP CPU

Command Index	Description
CRAMCMD\$K_WTBYTE32	Byte write in 32-bit space
CRAMCMD\$K_RDQUAD64	Quadword read in 64 bit space
CRAMCMD\$K_RDLONG64	Longword read in 64 bit space
CRAMCMD\$K_RDWORD64	Word read in 64 bit space
CRAMCMD\$K_RDBYTE64	Byte read in 64 bit space
CRAMCMD\$K_WTQUAD64	Quadword write in 64 bit space
CRAMCMD\$K_WTLONG64	Longword write in 64 bit space
CRAMCMD\$K_WTWORD64	Word write in 64 bit space
CRAMCMD\$K_WTBYTE64	Byte write in 64 bit space

byte_offset

VMS Usage: longword_unsigned
 type: longword (unsigned)
 access: read only
 mechanism: by value

Byte offset of the field to be written or read from the base of device interface register (CSR) space. Calculation of the RBADR and MASK fields of the hardware mailbox depends on the addressing and masking mechanisms provided by the remote bus. The **byte_offset** argument is used by IOC\$CRAM_CMD to calculate the RBADR.

adp_ptr

VMS Usage: longword_unsigned
 type: longword (unsigned)
 access: read only
 mechanism: by reference

Address of ADP associated with this command. IOC\$CRAM_CMD uses this parameter to determine which tightly-coupled I/O interconnect is the object of the mailbox transaction and to construct the mailbox command accordingly.

cram_ptr

VMS Usage: longword_unsigned
 type: longword (unsigned)
 access: read only
 mechanism: by reference

Address of CRAM. IOC\$CRAM_CMD returns the command, mask, and remote bus address values in the corresponding fields of the hardware I/O mailbox.

Returns

VMS Usage: cond_value
 type: longword_unsigned
 access: longword (unsigned)
 mechanism: write only—by value

Status indicating the success or failure of the operation.

System Routines

IOC\$CRAM_CMD

Return Values

SS\$NORMAL	The calculated command, mask, and remote bus address values have been written to the CRAM and/or the specified buffer.
SS\$BADPARAM	Illegal command supplied as input or illegal argument supplied in call
SS\$INSFARG	Insufficient arguments supplied in call

Description

IOC\$CRAM_CMD calculates the COMMAND, MASK, and RBADR fields for a hardware I/O mailbox according to the requirements of a specific I/O interconnect. It performs the following tasks:

- Obtains the address of the command table specific to the given I/O interconnect from ADP\$PS_COMMAND_TBL.
- Uses the value specified in the **command** argument as an index into the command table to determine the corresponding command supported by the I/O interconnect.
- If the command is valid for the I/O interconnect, IOC\$CRAM_CMD writes it to CRAM\$SL_COMMAND, to the specified buffer, or to both. If the command is invalid for the I/O interconnect, IOC\$CRAM_CMD returns SS\$BADPARAM status to its caller.
- Calculates the RBADR and MASK fields based of the hardware I/O mailbox, basing their values on the command, the address of device register interface space (ADP\$Q_CSR or IDB\$Q_CSR, if the **cram** argument is specified), the **byte_offset** argument, and interconnect-specific requirements. It writes these values to CRAM\$B_BYTE_MASK and CRAM\$Q_RBADR, to the specified buffer, or to both.
- Returns SS\$NORMAL status to its caller.

IOC\$CRAM_IO

Queues the hardware I/O mailbox defined within a controller register access mailbox (CRAM) to the mailbox pointer register (MBPR) and awaits the completion of the mailbox transaction.

Module

[SYSLOA]CRAM-IO

Macro

CRAM_IO

Format

IOC\$CRAM_IO cram

Context

IOC\$CRAM_IO conforms to the OpenVMS AXP calling standard. It acquires no spin locks and leaves IPL unchanged. After queuing the request and waiting for its completion, IOC\$CRAM_IO returns to its caller.

Arguments

cram

VMS Usage: address
 type: longword (unsigned)
 access: write only
 mechanism: by reference

Address of CRAM associated with the hardware I/O mailbox transaction.

Returns

VMS Usage: cond_value
 type: longword_unsigned
 access: longword (unsigned)
 mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	CRAM has been successfully queued to the MBPR.
SS\$BADPARAM	Supplied argument is not a CRAM.
SS\$CTRLERR	Error bit set in mailbox transaction.
SS\$INSFARG	No argument supplied in call.
SS\$INTERLOCK	Failed to queue hardware I/O mailbox to MBPR in queue time.

System Routines

IOC\$CRAM_IO

SS\$_TIMEOUT Mailbox operation did not complete in mailbox transaction timeout interval.

Description

IOC\$CRAM_IO performs an entire hardware I/O mailbox transaction from the queuing of the hardware I/O mailbox to the MBPR to the transaction's completion. A call to IOC\$CRAM_IO is the equivalent of independent calls to IOC\$CRAM_QUEUE and IOC\$CRAM_WAIT. Prior to calling IOC\$CRAM_IO, a driver typically calls IOC\$CRAM_CMD to insert a command, mask, and remote interconnect address into the hardware I/O mailbox portion of the CRAM. For CRAMs involved in writes to device interface registers, the driver must also insert the data to be written into CRAM\$Q_WDATA,

IOC\$CRAM_IO initiates an I/O operation to a device in remote I/O space by writing the physical address of the hardware I/O mailbox portion of a CRAM to the MBPR. If it is not able to post the mailbox to the MBPR in the MBPR queue timeout interval (CRAM\$Q_QUEUE_TIME), it returns SS\$_INTERLOCK status to its caller.

If it does successfully queue the mailbox, it sets the CRAM\$V_IN_USE bit in CRAM\$B_CRAM_FLAGS and repeatedly checks the done bit in the hardware I/O mailbox (CRAM\$V_MBX_DONE in CRAM\$W_MBX_FLAGS):

- If the done bit is not set in the mailbox transaction timeout interval (CRAM\$Q_WAIT_TIME), IOC\$CRAM_IO leaves the CRAM\$V_IN_USE bit in CRAM\$B_CRAM_FLAGS set and returns SS\$_TIMEOUT status to its caller.
- If the done bit is set, but the error bit in the mailbox (CRAM\$V_MBX_ERROR in CRAM\$W_MBX_FLAGS) is also set, IOC\$CRAM_IO clears CRAM\$V_IN_USE and returns SS\$_CTRLERR status to its caller. Note that, if the disable-error bit (CRAM\$V_DER) is set, IOC\$CRAM_IO never returns an error (although it may request an IOMBXERR fatal bugcheck in the event of an error).
- If the done bit is set and the error bit is clear, IOC\$CRAM_IO clears CRAM\$V_IN_USE and returns SS\$_NORMAL status to its caller. If IOC\$CRAM_IO returns SS\$_NORMAL status for read mailbox operations, the requested data has been returned to CRAM\$Q_RDATA. A return of SS\$_NORMAL status for mailbox write operations does not necessarily guarantee that the data placed in CRAM\$Q_WDATA has been successfully written to the device register.

IOC\$CRAM_QUEUE

Queues the hardware I/O mailbox defined within a controller register access mailbox (CRAM) to the mailbox pointer register (MBPR).

Module

[SYSLOA]CRAM-IO

Macro

CRAM_QUEUE

Format

IOC\$CRAM_QUEUE cram

Context

IOC\$CRAM_QUEUE conforms to the OpenVMS AXP calling standard. It acquires no spin locks and leaves IPL unchanged. After queuing the request, IOC\$CRAM_QUEUE returns to its caller. It is expected that the caller will eventually call IOC\$CRAM_WAIT to await completion of the request.

Arguments

cram
VMS Usage: address
type: longword (unsigned)
access: write only
mechanism: by reference

Address of CRAM to be queued.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	CRAM has been successfully queued to the MBPR.
SS\$BADPARAM	Supplied argument is not a CRAM.
SS\$INSFARG	No argument supplied in call
SS\$INTERLOCK	Failed to queue hardware I/O mailbox to MBPR in queue time.

System Routines

IOC\$CRAM_QUEUE

Description

IOC\$CRAM_QUEUE initiates an I/O operation to a device in remote I/O space by writing the physical address of the hardware I/O mailbox portion of a CRAM to the MBPR. Prior to calling IOC\$CRAM_QUEUE, a driver typically calls IOC\$CRAM_CMD to insert a command, mask, and remote interconnect address into the hardware I/O mailbox portion of the CRAM. For CRAMs involved in writes to device interface registers, the driver must also insert the data to be written into CRAM\$Q_WDATA,

If it is not able to post the mailbox to the MBPR in the MBPR queue timeout interval (CRAM\$Q_QUEUE_TIME), IOC\$CRAM_QUEUE returns SSS_INTERLOCK status to its caller. If the disable-error bit (CRAM\$V_DER) is set, IOC\$CRAM_QUEUE does not return an error (although it may request an IOMBXERR fatal bugcheck in the event of an error).

If IOC\$CRAM_QUEUE does successfully queue the mailbox, it sets the CRAM\$V_IN_USE bit in CRAM\$B_CRAM_FLAGS and returns SSS_NORMAL.

IOC\$CRAM_WAIT

Awaits the completion of a hardware I/O mailbox transaction to a tightly-coupled I/O interconnect.

Module

[SYSLOA]CRAM-IO

Macro

CRAM_WAIT

Format

IOC\$CRAM_WAIT cram

Context

IOC\$CRAM_WAIT conforms to the OpenVMS AXP calling standard. It acquires no spin locks and leaves IPL unchanged. After queuing the request, IOC\$CRAM_WAIT returns to its caller.

IOC\$CRAM_WAIT assumes that its caller has previously called IOC\$CRAM_QUEUE to post to the MBPR the hardware I/O mailbox defined within the specified CRAM for an I/O operation.

Arguments

cram

VMS Usage: address
type: longword (unsigned)
access: write only
mechanism: by reference

Address of CRAM associated with a previously-queued hardware I/O mailbox transaction.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	CRAM has been successfully queued to the MBPR.
SS\$BADPARAM	Supplied argument is not a CRAM.
SS\$CTRLERR	Error bit set in mailbox transaction.

System Routines

IOC\$CRAM_WAIT

SS\$INSFARG	No argument supplied in call.
SS\$TIMEOUT	Mailbox operation did not complete in mailbox transaction timeout interval.

Description

IOC\$CRAM_WAIT checks the done bit in the hardware I/O mailbox (CRAM\$V_MBX_DONE in CRAM\$W_MBX_FLAGS):

- If CRAM\$V_MBX_DONE is not set in the mailbox transaction timeout interval (CRAM\$Q_WAIT_TIME), IOC\$CRAM_WAIT leaves the CRAM\$V_IN_USE bit in CRAM\$B_CRAM_FLAGS set and returns SS\$TIMEOUT status to its caller.
- If CRAM\$V_MBX_DONE is set, but the error bit in the mailbox (CRAM\$V_MBX_ERROR in CRAM\$W_MBX_FLAGS) is also set, IOC\$CRAM_WAIT clears CRAM\$V_IN_USE and returns SS\$CTRLERR status to its caller. In this case, CRAM\$W_ERROR_BITS contains a device-specific encoding of additional status information.
- If the done bit is set and the error bit is clear, IOC\$CRAM_WAIT clears CRAM\$V_IN_USE and returns SS\$NORMAL status to its caller. If IOC\$CRAM_WAIT returns SS\$NORMAL status for read mailbox operations, the requested data has been returned to CRAM\$Q_RDATA. A return of SS\$NORMAL status for mailbox write operations does not necessarily guarantee that the data placed in CRAM\$Q_WDATA has been successfully written to the device register.

Note

If the disable-error bit (CRAM\$V_DER) is set, IOC\$CRAM_WAIT does not return an error (although it may request an IOMBXERR fatal bugcheck in the event of an error).

IOC\$DEALLOC_CNT_RES

Deallocates the requested number of items of a counted resource.

Module

DEALLOC_CNT_RES

Format

IOC\$DEALLOC_CNT_RES crab ,crctx

Context

IOC\$DEALLOC_CNT_RES conforms to the OpenVMS AXP calling standard. Its caller must be executing at fork IPL, holding the corresponding fork lock.

Arguments

crab

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of CRAB.

crctx

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of CRCTX structure.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$BADPARAM	CRCTX\$SL_ITEM_CNT and CRCTX\$SL_ITEM_NUM fields are invalid.

System Routines

IOC\$DEALLOC_CNT_RES

Description

IOC\$DEALLOC_CNT_RES deallocates a requested number of items of a counted resource. The resource request is described in the CRCTX structure; the counted resource itself is described in the CRAB. After deallocating the items, IOC\$DEALLOC_CNT_RES attempts to restart any waiters for the resource.

IOC\$DEALLOC_CNT_RES performs the following tasks:

1. It examines CRCTX\$V_ITEM_VALID in CRCTX\$SL_FLAGS. If it is clear, IOC\$DEALLOC_CNT_RES returns SSS_BADPARAM status to its caller.
2. It acquires the spin lock indicated by CRAB\$SL_SPINLOCK, raising IPL to IPL\$_IOLOCKLL in the process.
3. It scans the CRAB allocation array for a descriptor into which the items being deallocated (indicated by CRCTX\$SL_ITEM_CNT) can be merged.
4. It adjusts the CRAB allocation array and CRAB\$SL_VALID_DESC_CNT to reflect the deallocation.
5. If there are waiters for the counted resource, IOC\$DEALLOC_CNT_RES removes the CRCTX of the first waiter from the CRAB wait queue (CRAB\$SL_WQFL) and calls IOC\$ALLOC_CNT_RES to grant the requested number of resources.

If this attempt succeeds, IOC\$DEALLOC_CNT_RES restores the context of the stalled waiter (R3 through R5), releases the spin lock indicated by CRAB\$SL_SPINLOCK (upon the condition that the caller of IOC\$DEALLOC_CNT_RES did not already own this spin lock at the time of the call), and issues a standard call to the callback routine indicated by CRCTX\$SL_CALLBACK, passing it the address of the CRAB; the address of the CRCTX; the values stored in CRCTX\$SQ_FR3, CRCTX\$SQ_FR4, and CRCTX\$SQ_FR5; and SSS_NORMAL status.

IOC\$DEALLOC_CNT_RES continues to attempt to restart waiters in this manner until an allocation request fails. When this occurs, IOC\$DEALLOC_CNT_RES replaces its CRCTX in the CRAB wait queue, conditionally releases the spin lock indicated by CRAB\$SL_SPINLOCK, and returns SSS_NORMAL status to its caller.

6. If there are no waiters for the counted resource, IOC\$DEALLOC_CNT_RES conditionally releases the spin lock indicated by CRAB\$SL_SPINLOCK, and returns SSS_NORMAL status to its caller.

IOC\$DEALLOC_CRAB

Deallocates a counted resource allocation block (CRAB).

Module

ALLOC_CNT_RES

Format

IOC\$DEALLOC_CRAB crab

Context

IOC\$DEALLOC_CRAB conforms to the OpenVMS AXP calling standard. Because IOC\$DEALLOC_CRAB calls EXE\$DEANONPAGED, its caller cannot be executing above IPL\$SYNCH.

Arguments

crab
VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of CRAB to be deallocated.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL The routine completed successfully.

Description

A driver calls IOC\$DEALLOC_CRAB to deallocate a CRAB. IOC\$DEALLOC_CRAB passes the address of the CRAB to EXE\$DEANONPAGED and returns SS\$NORMAL status to its caller.

IOC\$DEALLOC_CRCTX

Deallocates a counted resource context block (CRCTX).

Module

ALLOC_CNT_RES

Format

IOC\$DEALLOC_CRCTX crctx

Context

IOC\$DEALLOC_CRCTX conforms to the OpenVMS AXP calling standard. Because IOC\$DEALLOC_CRCTX calls EXE\$DEANONPAGED, its caller cannot be executing above IPL\$_SYNCH.

Arguments

crctx

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of CRCTX to be deallocated.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$_NORMAL The routine completed successfully.

Description

A driver calls IOC\$DEALLOC_CRCTX to deallocate a CRCTX. IOC\$DEALLOC_CRCTX passes the address of the CRCTX to EXE\$DEANONPAGED and returns SS\$_NORMAL status to its caller.

IOC\$DEALLOCATE_CRAM

Deallocates a controller register access mailbox.

Module

CRAM-ALLOC

Macro

CRAM_DEALLOC

Format

IOC\$DEALLOCATE_CRAM cram

Context

IOC\$DEALLOCATE_CRAM conforms to the OpenVMS AXP calling standard. Its caller must be executing at or below IPL 8 and must not hold spin locks ranked higher than IO_MISC.

IOC\$DEALLOCATE_CRAM acquires and releases the IO_MISC spin lock and returns to its caller at its caller's IPL.

Arguments

cram

VMS Usage: address
type: longword (unsigned)
access: write only
mechanism: by reference

Address of CRAM to be deallocated by IOC\$DEALLOCATE_CRAM

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	CRAM has been successfully deallocated.
SS\$BADPARAM	Supplied argument is not a CRAM.
SS\$INSFARG	Insufficient arguments supplied in call

Description

IOC\$DEALLOCATE_CRAM deallocates a single controller register access mailbox (CRAM).

IOC\$KP_REQCHAN

Stalls a kernel process in such a manner that it can be resumed by the granting of a device controller channel.

Module

KERNEL_PROCESS_MIN, KERNEL_PROCESS_MON

Macro

KP_STALL_REQCHAN

Format

IOC\$KP_REQCHAN kpb ,priority

Context

IOC\$KP_REQCHAN conforms to the OpenVMS AXP calling standard. It can only be called by a kernel process.

A kernel process calls IOC\$KP_REQCHAN at fork IPL holding the appropriate fork lock.

If the requested channel is busy, either the channel-requesting routine IOC\$PRIMITIVE_REQCHANH or IOC\$PRIMITIVE_REQCHANL preserves the contents of its caller's R3 in UCBSQ_FR3 (contents of caller's R3). IOC\$RELCHAN eventually issues a JSB instruction to the fork routine upon granting the channel request. At this time, the kernel process is provided with the contents of UCBSQ_FR3 in R3, the IDB address in R4, and the UCB address in R5.

Arguments

kpb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of the caller's KPB which must be a VEST KPB. KPB\$PS_UCB must contain the address of a UCB and KPB\$PS_IRP must contain the address of an IRP.

priority

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

Priority of the request for the controller channel. You must specify one of the following symbolic constants:

Constant	Meaning
KPBSK_LOW	Insert fork block of UCB requesting controller channel at the tail of the channel-wait queue.
KPBSK_HIGH	Insert fork block of UCB requesting controller channel at the head of the channel-wait queue.

Returns

VMS Usage: cond_value
 type: longword_unsigned
 access: longword (unsigned)
 mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$_NORMAL	The routine completed successfully.
SS\$_BADPARAM	The kp argument does not specify a VEST KP, or an illegal value was supplied in the priority argument.
SS\$_INSFARG	Not all of the required arguments were specified.

Description

IOC\$KP_REQCHAN first checks the CRB to determine if the controller channel is busy. If the CRB is not busy (CRB\$V_BSY in CRB\$B_MASK is clear), IOC\$KP_REQCHAN grants the channel request immediately by placing the UCB address in IDB\$L_OWNER and returning SS\$_NORMAL status to its caller.

If the CRB is busy, IOC\$KP_REQCHAN performs the following tasks to initiate a stall of the kernel process:

1. Copies the **priority** argument to KPBSIS_CHANNEL_DATA.
2. Inserts the procedure descriptor of subroutine STALL_REQCHAN in KPB\$PS_SCH_STALL_RTN, thus making it the kernel process scheduling stall routine.
3. Clears KPB\$PS_SCH_RESTART, thus indicating that there is no kernel process scheduling restart routine.
4. Calls EXE\$KP_STALL_GENERAL, passing to it the address of the KP.

Note that, having stalled the kernel process, the STALL_REQCHAN kernel process scheduling stall routine returns control to EXE\$KP_STALL_GENERAL, which returns to the initiator of the kernel process thread (that is, the caller of EXE\$KP_START or EXE\$KP_RESTART). When the controller channel request is ultimately granted, STALL_REQCHAN calls EXE\$KP_RESTART which, in turn, passes control back to IOC\$KP_REQCHAN. IOC\$KP_REQCHAN then returns to the kernel process that called it.

IOC\$KP_WFIKPCH, IOC\$KP_WFIRLCH

Stall a kernel process in such a manner that it can be resumed by device interrupt processing.

Module

KERNEL_PROCESS_MIN, KERNEL_PROCESS_MON

Macro

KP_STALL_WFIKPCH
KP_STALL_WFIRLCH

Format

IOC\$KP_WFIKPCH kpb ,time ,newipl
IOC\$KP_WFIRLCH kpb ,time ,newipl

Context

IOC\$KP_WFIKPCH and IOC\$KP_WFIRLCH conform to the OpenVMS AXP calling standard. They can only be called by a kernel process.

When called, IOC\$KP_WFIKPCH or IOC\$KP_WFIRLCH assumes that the local processor has obtained the appropriate synchronization with the device database by securing the appropriate device lock, as recorded in the unit control block (UCBSL_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.

Before exiting, the wait-for-interrupt routine (IOCSPRIMITIVE_WFIKPCH or IOCSPRIMITIVE_WFIRLCH) conditionally releases the device lock, so that if the initiator of the kernel process thread previously owned the device lock, it will continue to hold it when it regains control. IOCSPRIMITIVE_WFIKPCH or IOCSPRIMITIVE_WFIRLCH also lowers the local processor's IPL to the IPL specified in the **newipl** argument.

Arguments

kpb

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of the caller's KPB (which must be a VEST KPB). KPB\$PS_UCB must contain the address of a UCB and KPB\$PS_IRP must contain the address of an IRP.

time

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

Timeout value in seconds.

newipl

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

IPL to which to lower before returning to the initiator of the kernel process thread (that is, the caller of EXESKP_START or EXESKP_RESTART). This IPL must be the fork IPL associated with device processing and at which the kernel process was executing prior to invoking the DEVICELock macro.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$BADPARAM	The kp argument does not specify a VEST KPB.
SS\$INSFARG	Not all of the required arguments were specified.
SS\$TIMEOUT	A timeout has occurred.

Description

IOC\$KP_WFIKPCH and IOC\$KP_WFIRLCH perform the following tasks to initiate a stall of the kernel process:

1. Copy the **time** argument to KPB\$IS_TIMEOUT_TIME and the **newipl** argument to KPB\$IS_RESTORE_IPL.
2. Move the symbolic constant KPB\$K_KEEP (for IOC\$KP_WFIKPCH) or KPB\$K_RELEASE (for IOC\$KP_WFIRLCH) to KPB\$IS_CHANNEL_DATA.
3. Insert the procedure descriptor of subroutine STALL_WFIXXCH in KPB\$PS_SCH_STALL_RTN, this making it the kernel process scheduling stall routine.
4. Clear KPB\$PS_SCH_RESTART, thus indicating that there is no kernel process scheduling restart routine.
5. Call EXESKP_STALL_GENERAL, passing to it the address of the KPB.

Note that, having stalled the kernel process, the STALL_WFIXXCH kernel process scheduling stall routine returns control to EXESKP_STALL_GENERAL, which returns to the initiator of the kernel process thread (that is, the caller of EXESKP_START or EXESKP_RESTART). When interrupt servicing transfers control back to STALL_WFIXXCH, or a timeout occurs, STALL_WFIXXCH calls EXESKP_RESTART which, in turn, passes control back to IOC\$KP_WFIKPCH or IOC\$KP_WFIRLCH. The kernel process wait-for-interrupt stall routine then returns to the kernel process that called it.

IOC\$LOAD_MAP

Loads a set of adapter-specific map registers.

Module

[CPU_{xxxx}]MAPREG__{xxxx}†

Format

IOC\$LOAD_MAP adp ,crctx ,svapte ,boff ,dma_address_ref

Context

IOC\$LOAD_MAP conforms to the OpenVMS AXP calling standard.

Arguments

adp

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of ADP for adapter which provides the map registers.

crctx

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of CRCTX that describes a map register allocation (that is, a CRCTX that has been obtained by a call to IOC\$ALLOC_CRCTX and supplied in a call to IOC\$ALLOC_CNT_RES for the CRAB that manages this adapter's map registers).

svapte

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

System virtual address of the PTE for the first page to be used in the transfer.

boff

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

Byte offset into the first page of the transfer buffer.

† where *xxxx* represents the internal OpenVMS code number for an AXP CPU

dma_address_ref

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of a location to receive a port-specific DMA address. For DEC 3000-500 systems, this address is a function of the starting map register and the byte offset. A DEC 3000-500 system port driver must strip off two lower bits when loading the address register of the DMA device.

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$INSMEM	Memory allocation failure.

Description

A driver calls IOC\$LOAD_MAP to load a set of adapter-specific map registers. The driver must have previously allocated the map registers (including an extra two to serve as guard pages) in calls to IOC\$ALLOC_CRCTX and IOC\$ALLOC_CNT_RES.

IOC\$LOAD_MAP computes a port-specific DMA address and returns it to the driver for use in a hardware I/O mailbox operation that loads the address register of a DMA device.

IOC\$MAP_IO

IOC\$MAP_IO maps I/O bus physical address space into an address region accessible by the processor. The caller of this routine can express the mapping request in terms of the bus address space without regard to address swizzle space, dense space, or sparse space.

IOC\$MAP_IO is supported on PCI, EISA, TURBOchannel, or PCI systems. It is not supported on XMI systems.

Description

The routine prototype is as follows:

```
int ioc$map_io (ADP      *adp,  
               int      node,  
               uint64   *physical_offset,  
               int      num_bytes,  
               int      attributes,  
               uint64   *iohandle)
```

Inputs

adp Address of bus ADP. Driver can get this from IDB\$PS_ADP.

node Bus node number of device. Bus specific interpretation. Available to driver in CRB\$L_NODE (driver must be loaded with /NODE qualifier).

physical_offset Address of a quadword cell. For EISA, PCI, and Futurebus, the quadword cell should contain the starting bus physical address to be mapped. For Turbochannel, the quadword cell should contain the physical offset from the Turbochannel slot base address.

num_bytes Number of bytes to be mapped. Expressed in terms of the bus/device without regard to the platform hardware addressing tricks.

attributes Specifies desired attributes of space to be mapped. From [lib]iocdef. One of the following:

IOC\$K_BUS_IO_BYTE_GRAN

Request mapping in a platform address space which corresponds to bus I/O space and provides byte granularity access. In general, if you are mapping device control registers that exist in bus I/O space, you should specify this attribute. For example, drivers for PCI devices with registers in PCI I/O space or EISA devices with EISA I/O port addresses should request mapping with this attribute.

IOC\$K_BUS_MEM_BYTE_GRAN

Request mapping in a platform address space which corresponds to bus memory space and provides byte granularity access. In general, if you are mapping device registers that exist in bus memory space, you should specify this attribute. For example, drivers for PCI devices with registers in PCI memory space should request mapping with this attribute.

IOC\$K_BUS_DENSE_SPACE

Request mapping in a platform address space that corresponds to bus memory space and provides coarse access granularity. IOC\$K_BUS_DENSE_SPACE is suitable for mapping device memory buffers such as graphics frame buffers. In IOC\$K_BUS_DENSE_SPACE, there must be no side effects on reads and it may be possible for the processor to merge writes. Thus you should not map device registers in dense space.

iohandle Pointer to a 64 bit cell. A 64 bit magic number is written to this cell by IOC\$MAP_IO when the mapping request is successful. The caller must save the iohandle, as it is an input to IOC\$CRAM_CMD and to the new platform independent access routines IOC\$READ_IO and IOC\$WRITE_IO.

Outputs

SS\$_NORMAL Success. The address space is mapped. A 64 bit IOHANDLE is written to the caller's buffer.

SS\$_BADPARAM Bad input argument. For example, the requested bus address may not be accessible from the CPU, or the attribute may be unrecognized.

SS\$_UNSUPPORTED Address space with the requested attributes not available on this platform. For example, the Jensen platform does not support EISA memory dense space.

SS\$_INSFSPTS Not enough PTEs to satisfy mapping request.

IOC\$NODE_FUNCTION

Performs node-specific functions on behalf of a driver, such as enabling or disabling interrupts from a bus slot.

Module

[SYSLOA]MISC_SUPPORT

Format

IOC\$NODE_FUNCTION crb_addr ,function_code

Context

IOC\$NODE_FUNCTION conforms to the OpenVMS AXP calling standard. It may be called in kernel mode at any IPL and may acquire the MEGA spin lock (SPL\$C_MEGA), raising IPL to IPL\$_MEGA in the process, depending on the function code.

Arguments

crb_addr

VMS Usage: address
type: longword (unsigned)
access: read only
mechanism: by reference

Address of CRB.

function_code

VMS Usage: longword_unsigned
type: longword (unsigned)
access: read only
mechanism: by value

Function to be effected for the bus node indicated by the **crb_addr** argument. You can specify one of the following values (defined by the \$IOCDEF macro in SYSS\$LIBRARY:LIB.MLB). Note that not all function codes are supported by all adapters.

Code	Action
IOC\$K_ENABLE_INTR	Enable interrupts
IOC\$K_DISABLE_INTR	Disable interrupts
IOC\$K_ENABLE_SG	Enable scatter/gather map
IOC\$K_DISABLE_SG	Disable scatter/gather map
IOC\$K_ENABLE_PAR	Enable parity
IOC\$K_DISABLE_PAR	Disable parity
IOC\$K_ENABLE_BLK	Enable block mode
IOC\$K_DISABLE_BLK	Disable block mode

Returns

VMS Usage: cond_value
type: longword_unsigned
access: longword (unsigned)
mechanism: write only—by value

Status indicating the success or failure of the operation.

Return Values

SS\$NORMAL	The routine completed successfully.
SS\$ILLIOFUNC	Requested function not available on this platform or bus.

Description

IOC\$NODE_FUNCTION locates the ADP associated with the specified CRB (from VEC\$PS_ADP) and calls the adapter-specific node function routine specified in ADP\$PS_NODE_FUNCTION. The node function routine performs the function indicated by the **function_code** argument.

Drivers request the node-specific functions as follows:

- IOC\$K_ENABLE_INTR, IOC\$K_DISABLE_INTR

On both DEC 3000-500 and DEC 3000-300 systems, when the console transfers control to OpenVMS AXP, TURBOchannel interrupts from all slots are disabled. The controller or unit initialization routine of a driver for a TURBOchannel devices must call IOC\$NODE_FUNCTION, specifying the IOC\$K_ENABLE_INTR function code, to enable interrupts for the TURBOchannel slot in which the device resides. The field CRB\$SL_NODE of the specified CRB contains this slot number.

Calling IOC\$NODE_FUNCTION with the IOC\$K_DISABLE_INTR code disables interrupts from the node.

- IOC\$K_ENABLE_SG, IOC\$K_DISABLE_SG

On DEC 3000-500 systems, calling IOC\$NODE_FUNCTION with function code IOC\$K_ENABLE_SG, allows DMA transactions from a device to use the DEC 3000-500 system scatter/gather map. The TURBOchannel slot of the device is indicated by the field CRB\$SL_NODE in the specified CRB.

Calling IOC\$NODE_FUNCTION with the IOC\$K_DISABLE_SG code disables the scatter/gather map.

DEC 3000-300 systems have no scatter/gather map. IOC\$NODE_FUNCTION returns SS\$ILLIOFUNC if it is called on a DEC 3000-300 system with either an IOC\$K_ENABLE_SG or IOC\$K_DISABLE_SG function code.

- IOC\$K_ENABLE_PAR, IOC\$K_DISABLE_PAR

On DEC 3000-500 systems, calling IOC\$NODE_FUNCTION with function code IOC\$K_ENABLE_PAR causes parity to be generated on TURBOchannel transactions directed to a device, and causes parity to be checked on TURBOchannel transactions coming from the device. The TURBOchannel slot of the device is indicated by the field CRB\$SL_NODE in the specified CRB.

System Routines

IOC\$NODE_FUNCTION

If an adapter supports TURBOchannel parity, a driver controller or unit initialization routine enable it by calling IOC\$NODE_FUNCTION with the IOC\$K_ENABLE_PAR function code.

Calling IOC\$NODE_FUNCTION with the IOC\$K_DISABLE_PAR code disables TURBOchannel parity.

DEC 3000-300 systems do not support TURBOchannel parity. IOC\$NODE_FUNCTION returns SSS_ILLIOFUNC if it is called on a DEC 3000-300 system with either an IOC\$K_ENABLE_PAR or IOC\$K_DISABLE_PAR function code.

- IOC\$K_ENABLE_BLK, IOC\$K_DISABLE_BLK

On DEC 3000-500 systems, calling IOC\$NODE_FUNCTION with function code IOC\$K_ENABLE_BLK causes block mode to be used on TURBOchannel transactions to and from the device indicated by the field CRB\$L_NODE in the specified CRB. Most drivers have no need to enable block mode.

DEC 3000-300 systems do not support TURBOchannel block mode. IOC\$NODE_FUNCTION returns SSS_ILLIOFUNC if it is called on a DEC 3000-300 system with either an IOC\$K_ENABLE_BLK or IOC\$K_DISABLE_BLK function code.

IOC\$READ_IO

Reads a value from a previously mapped location in I/O address space. This routine requires that the I/O space to be accessed has been previously mapped by a call to IOC\$MAP_IO.

IOC\$READ_IO is supported on PCI, EISA, TURBOchannel, and PCI systems. It is not supported on XMI systems.

Description

The routine prototype for IOC\$READ_IO is as follows:

```
int ioc$read_io (ADP      *adp,
                uint64   *iohandle,
                int      offset,
                int      length,
                void     *read_data)
```

Inputs

- adp Address of bus ADP. Driver can get this from IDB\$PS_ADP.
- iohandle Pointer to a 64 bit IOHANDLE. The 64 bit IOHANDLE is obtained by calling the platform independent mapping routine IOC\$MAP_IO.
- offset Offset in device space of field to be read or written. This should be specified as an offset from the base of the space that was previously mapped by the call to IOC\$MAP_IO. The offset is specified in terms of the device or bus without regard to any hardware address trickery.
- length Length of field to be read or written. Should be 1 (byte), 2 (word), 3 (tribyte), 4 (longword) or 8 (quadword). Note that not all of these lengths are supported on all buses.
- read_data Pointer to a data cell. For ioc\$read_io, the data read from the device will be returned in this cell. If the requested data length was 1, 2, 3, or 4, a longword is written to the data cell with valid data in the byte lane(s) corresponding to the requested length and offset. If the requested data length was 8, a quadword is written to the data cell.
- write_data Pointer to a data cell. The data cell should contain the data to be written to the device. For lengths of 1, 2, 3 or 4, the ioc\$write_io routine reads a longword from the data cell and writes this longword to the bus with the proper byte enables set according to the length and offset. The actual data to be written must be positioned in the proper byte lane(s) according to the requested length and offset. For a length 8 transfer, the ioc\$write_io routine reads a quadword from the data cell.

Outputs

- SS\$_NORMAL Success. If IOC\$READ_IO, data is returned in the caller's buffer. If IOC\$WRITE_IO, data is written to device.
- SS\$_BADPARAM Bad input argument, such as an illegal length.

System Routines

IOC\$READ_IO

SS\$UNSUPPORTED A transaction length not supported by this bus
or platform.

IOC\$UNMAP_IO

Unmaps a previously mapped I/O address space, returning the IOHANDLE and the PTEs to the system. The caller's quadword cell containing the IOHANDLE is cleared.

Description

The routine prototype is as follows:

```
int ioc$unmap_io (ADP *adp,  
                 uint64 *iohandle)
```

IOC\$WRITE_IO

Writes a value to a previously mapped location in I/O address space. IOC\$WRITE_IO requires that the I/O space to be accessed has been previously mapped by a call to IOC\$MAP_IO.

Description

The routine prototype is as follows:

```
int ioc$write_io (ADP      *adp,  
                 uint64  *iohandle,  
                 int      offset,  
                 int      length,  
                 void     *write_data)
```

Inputs

- adp Address of bus ADP. Driver can get this from IDB\$PS_ADP.
- iohandle Pointer to a 64 bit IOHANDLE. The 64 bit IOHANDLE is obtained by calling the platform independent mapping routine IOC\$MAP_IO.
- offset Offset in device space of field to be read or written. This should be specified as an offset from the base of the space that was previously mapped by the call to IOC\$MAP_IO. The offset is specified in terms of the device or bus without regard to any hardware address trickery.
- length Length of field to be read or written. Should be 1 (byte), 2 (word), 3 (tribyte), 4 (longword) or 8 (quadword). Note that not all of these lengths are supported on all buses.
- read_data Pointer to a data cell. For ioc\$read_io, the data read from the device will be returned in this cell. If the requested data length was 1, 2, 3, or 4, a longword is written to the data cell with valid data in the byte lane(s) corresponding to the requested length and offset. If the requested data length was 8, a quadword is written to the data cell.
- write_data Pointer to a data cell. The data cell should contain the data to be written to the device. For lengths of 1, 2, 3 or 4, the ioc\$write_io routine reads a longword from the data cell and writes this longword to the bus with the proper byte enables set according to the length and offset. The actual data to be written must be positioned in the proper byte lane(s) according to the requested length and offset. For a length 8 transfer, the ioc\$write_io routine reads a quadword from the data cell.

Outputs

- SS\$NORMAL Success. If ioc\$read_io, data is returned in the caller's buffer. If ioc\$write_io, data is written to device.
- SS\$BADPARAM Bad input argument, such as an illegal length.
- SS\$UNSUPPORTED A transaction length not supported by this bus or platform.

IOC_STD\$ALTREQCOM

Completes an I/O request for a device using the disk or tape class drivers.

Module

IOSUBNPAG

Format

IOC_STD\$ALTREQCOM (iost1, iost2, cdrp, irp_p, ucb_p)

Arguments

Argument	Type	Access	Mechanism	Status
iost1	integer	input	value	required
iost2	integer	input	value	required
cdrp	CDRP	input	reference	required
irp_p	pointer	output	reference	required
ucb_p	pointer	output	reference	required

iost1

First longword of I/O status.

iost2

Second longword of I/O status.

cdrp

Class driver request packet.

irp_p

Address at which IOC_STD\$ALTREQCOM writes the address of the I/O request packet.

ucb_p

Address at which IOC_STD\$ALTREQCOM writes the address of the unit control block.

Context

IOC_STD\$ALTREQCOM is typically called at fork IPL with the corresponding fork lock held in an OpenVMS multiprocessing system.

Description

For Digital internal use only.

System Routines

IOC_STD\$ALTREQCOM

Macro

CALL_ALTREQCOM

In a Step 2 driver, the CALL_ALTREQCOM macro simulates a JSB to IOC\$ALTREQCOM in a Step 1 driver. CALL_ALTREQCOM calls IOC_STD\$ALTREQCOM, using the current contents of R0, R1, and R5 as the **iost1**, **iost2**, and **cdrp** arguments, respectively. When IOC_STD\$ALTREQCOM returns, the macro returns the address of the IRP in R3 and the address of the UCB in R4.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$ALTREQCOM replaces IOC\$ALTREQCOM (used by Step 1 and OpenVMS VAX drivers). Unlike IOC\$ALTREQCOM, IOC_STD\$ALTREQCOM does not return the addresses of the IRP and UCB in R3 and R5, respectively.

IOC_STD\$BROADCAST

Broadcasts the specified message to a given terminal.

Module

IOSUBNPAG

Format

status = IOC_STD\$BROADCAST (msglen, msg_p, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
msglen	integer	input	value	required
msg_p	address	input	reference	required
ucb	UCB	input	reference	required

msglen
Message length.

msg_p
Message.

ucb
Address of target terminal's UCB.

Return Values

SS\$_ILLIOFUNC	The specified term_ucb is not associated with a terminal.
SS\$_INSFMEM	Insufficient dynamic nonpaged pool to satisfy the request.
SS\$_NORMAL	The broadcast completed successfully.

Context

IOC_STD\$BROADCAST is typically called at fork IPL with the corresponding fork lock held in an OpenVMS multiprocessing system.

Description

For Digital internal use only.

System Routines

IOC_STD\$BROADCAST

Macro

CALL_BROADCAST [save_r1]

where:

save_r1 indicates that the macro should preserve register R1 across the call to IOC_STD\$BROADCAST. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

In a Step 2 driver, the CALL_BROADCAST macro simulates a JSB to IOC\$BROADCAST in a Step 1 driver. CALL_BROADCAST calls IOC_STD\$BROADCAST, using the current contents of R1, R2, and R5 as the **msglen**, **msg_p**, and **ucb** arguments, respectively. It returns status in R0. Unless you specify **save_r1=NO**, the macro preserves the quadword register R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$BROADCAST replaces IOC\$BROADCAST (used by Step 1 and OpenVMS VAX drivers). Unlike IOC\$BROADCAST, IOC_STD\$BROADCAST does not preserve R1 across the call.

IOC_STD\$CANCELIO

Conditionally marks a UCB so that its current I/O request will be canceled.

Module

IOSUBNPAG

Format

IOC_STD\$CANCELIO (chan, irp, pcb, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
chan	integer	input	value	required
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required

chan

Channel index number.

irp

I/O request packet. IOC_STD\$CANCELIO reads the following IRP fields:

Field	Contents
IRP\$L_PID	Process identification of the process that queued the I/O request
IRP\$L_CHAN	I/O request channel index number

pcb

Current process control block.

ucb

Unit control block. IOC_STD\$CANCELIO reads UCB\$L_STS to determine if the device is busy (UCB\$V_BSY set) or idle (UCB\$V_BSY clear). IOC_STD\$CANCELIO sets UCB\$V_CANCEL if the I/O request should be canceled.

Context

IOC_STD\$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 multiprocessing environment.

System Routines

IOC_STD\$CANCELIO

Description

IOC_STD\$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 (UCBSV_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\$L_CHAN).
4. It sets the cancel-I/O bit in the UCB status longword (UCBSV_CANCEL in UCB\$L_STS).

Macro

CALL_CANCELIO [save_r0r1]

where:

save_r0r1 indicates that the macro should preserve registers R0 and R1 across the call to IOC_STD\$CANCELIO. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

In a Step 2 driver, the CALL_CANCELIO macro simulates a JSB to IOC\$CANCELIO in a Step 1 driver. CALL_CANCELIO calls IOC_STD\$CANCELIO, using the current contents of R2, R3, R4, and R5 as the **chan**, **irp**, **pcb**, and **ucb** arguments, respectively. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$CANCELIO replaces IOC\$CANCELIO (used by Step 1 and OpenVMS VAX drivers). Unlike IOC\$CANCELIO, IOC_STD\$CANCELIO does not preserve R0 and R1 across the call.

IOC_STD\$CLONE_UCB

Copies a template UCB and links it to the appropriate DDB list.

Module

UCBCREDEL

Format

status = IOC_STD\$CLONE_UCB (tmpl_ucb, new_ucb_p)

Arguments

Argument	Type	Access	Mechanism	Status
tmpl_ucb	UCB	input	reference	required
new_ucb_p	pointer	output	reference	required

tmpl_ucb

Template unit control block.

new_ucb_p

Location into which IOC_STD\$CLONE_UCB writes the address of the newly-created unit control block.

Return Values

SS\$NORMAL

UCB cloning was successful.

SS\$INSFMEM

Insufficient nonpaged pool to copy UCB.

Context

A driver calls IOC_STD\$CLONE_UCB at or below IPL\$MAILBOX with the I/O database locked for write access.

Description

For Digital internal use only.

Macro

CALL_CLONE_UCB [interface_warning=YES]

where:

interface_warning=YES, the default, specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. **interface_warning=NO** suppresses the warning.

System Routines

IOC_STD\$CLONE_UCB

In a Step 2 driver, `CALL_CLONE_UCB` simulates a JSB to `IOC$CLONE_UCB`. It calls `IOC_STD$CLONE_UCB` using the current contents of R5 as the **tmpl_uctb** argument. `CALL_CLONE_UCB` returns status in R0 and the address of the newly-created UCB in R2, but does not return the address of the UCBs that precede and follow it on the DDB chain in R3 and R1, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- `IOC_STD$CLONE_UCB` replaces `IOC$CLONE_UCB` (used by Step 1 drivers and OpenVMS VAX drivers). `IOC_STD$CLONE_UCB` does not return the addresses of the UCBs that precede and follow the newly-created UCB on the DDB chain.

IOC_STD\$COPY_UCB

Copies and initializes a template UCB and ORB.

Module

UCBCREDEL

Format

status = IOC_STD\$COPY_UCB (src_ucb, new_ucb)

Arguments

Argument	Type	Access	Mechanism	Status
src_ucb	UCB	input	reference	required
new_ucb	pointer	output	reference	required

src_ucb

Template unit control block.

new_ucb

Location into which IOC_STD\$COPY_UCB writes the address of the newly-created duplicate unit control block.

Return Values

SS\$NORMAL

UCB copy was successful.

SS\$INSFMEM

Insufficient nonpaged pool to copy UCB.

Context

A driver calls IOC_STD\$COPY_UCB at or below IPL\$MAILBOX with the I/O database locked for write access.

Description

For Digital internal use only.

Macro

CALL_COPY_UCB

In a Step 2 driver, CALL_COPY_UCB simulates a JSB to IOC\$COPY_UCB. It calls IOC_STD\$COPY_UCB using the current contents of R5 as the **src_ucb** argument. CALL_CLONEUCB returns the address of the newly-created UCB in R2.

System Routines

IOC_STD\$COPY_UCB

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note that IOC_STD\$COPY_UCB replaces IOC\$COPY_UCB (used by Step 1 drivers and OpenVMS VAX drivers). IOC_STD\$COPY_UCB does not preserve the contents of R3 and R4 across the call.

IOC_STD\$CREDIT_UCB

Credits the UCB charges associated with a given UCB against the process identified by the contents of UCB\$LCPID.

Module

UCBCREDEL

Format

IOC_STD\$CREDIT_UCB (ucb)

Arguments

Argument	Type	Access	Mechanism	Status
ucb	UCB	input	reference	required

ucb
Unit control block.

Context

A driver calls IOC_STD\$CREDIT_UCB at IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

CALL_CREDIT_UCB

In a Step 2 driver, CALL_CREDIT_UCB simulates a JSB to IOC\$CREDIT_UCB. It calls IOC_STD\$CREDIT_UCB using the current contents of R5 as the **ucb** argument.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$CREDIT_UCB replaces IOC\$CREDIT_UCB (used by Step 1 drivers and OpenVMS VAX drivers).

System Routines

IOC_STD\$CVT_DEVNAM

IOC_STD\$CVT_DEVNAM

Converts a device name and unit number to a physical device name string.

Module

IOSUBNPAG

Format

status = IOC_STD\$CVT_DEVNAM (buflen, buf, form, ucb, outlen_p)

Arguments

Argument	Type	Access	Mechanism	Status
buflen	integer	input	value	required
buf	address	input	reference	required
form	integer	input	value	required
ucb	UCB	reference	input	required
outlen_p	pointer	output	reference	required

buflen

Size of output buffer in bytes.

buf

Output buffer.

form

Name string formation mode, as follows:

Mode	Description
-2 (DVIS_DISPLAY_DEVNAM)	Name suitable for displays but not suitable for \$ASSIGN: "\$allocclass\$ddcn: (host1[, host2])", "node\$ddcn", or "ddcn"
-1 (DVIS_DEVNAM)	Name suitable for displays: "node\$ddcn" for non-local devices or "node\$ddcn" or "ddcn" for local devices
0 (DVIS_FULLDEVNAM)	Name with appropriate node information: either "\$allocclass\$ddcn" or "node\$ddcn"
1 (DVIS_ALLDEVNAM)	Name with allocation class information: either "\$allocclass\$ddcn" or "node\$ddcn"
2 (no GETDVI item code)	Old-fashioned name: "ddcn"
3 (no GETDVI item code)	Secondary path name for displays (same as -1 except secondary path name is returned)

Mode	Description
4 (no GETDVI item code)	Path controller name for displays (same as -1 except no unit number is appended)

ucb

Unit control block for device.

outlen_p

Address of location in which IOC_STD\$CVT_DEVNAM returns the length of the conversion string.

Return Values

SS\$_BUFFEROVF	Successful completion, but specified buffer cannot hold the entire device name string.
SS\$_NORMAL	Normal, successful completion.

Context

IOC_STD\$CVT_DEVNAM is typically called at fork IPL with the corresponding fork lock held in an OpenVMS multiprocessing system.

Description

For Digital internal use only.

Macro

CALL_CVT_DEVNAM

In a Step 2 driver, the CALL_CVT_DEVNAM macro simulates a JSB to IOC\$CVT_DEVNAM in a Step 1 driver. CALL_CVT_DEVNAM calls IOC_STD\$CVT_DEVNAM, using the current contents of R0, R1, R4, and R5 as the **bufLen**, **buf**, **form**, and **ucb** arguments, respectively. When IOC_STD\$CVT_DEVNAM returns, the macro returns status in R0 and the length of the conversion string in R1.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$CVT_DEVNAM replaces IOC\$CVT_DEVNAM (used by Step 1 and OpenVMS VAX drivers). Unlike IOC\$CVT_DEVNAM, IOC_STD\$CVT_DEVNAM does not return the length of the conversion string in R1.

IOC_STD\$CVTLOGPHY

Conditionally converts a logical block number to a physical disk address and stores the result in the I/O request packet.

Module

IOSUBRAMS

Format

IOC_STD\$CVTLOGPHY (lbn, irp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
lbn	integer	input	value	required
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

lbn

Logical block number to be converted.

irp

I/O request packet.

ucb

Unit control block.

Context

A driver calls IOC_STD\$CVTLOGPHY at fork IPL with the corresponding fork lock held in a multiprocessing system.

Description

For Digital internal use only.

Macro

CALL_CVTLOGPHY

In a Step 2 driver, the CALL_CVTLOGPHY macro simulates a JSB to IOC\$CVTLOGPHY in a Step 1 driver. CALL_CVTLOGPHY calls IOC_STD\$CVTLOGPHY, using the current contents of R0, R3, and R5 as the **lbn**, **irp** and **ucb** arguments, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- `IOC_STD$CVTLOGPHY` replaces `IOC$CVTLOGPHY` (used by Step 1 and OpenVMS VAX drivers). Unlike `IOC$CVTLOGPHY`, `IOC_STD$CVTLOGPHY` does not preserve R3 across the call.

System Routines

IOC_STD\$DELETE_UCB

IOC_STD\$DELETE_UCB

Deletes the specified UCB if its reference count is zero and UCBSV_DELETEUCB is set in UCBSL_STS.

Module

UCBCREDEL

Format

IOC_STD\$DELETE_UCB (ucb)

Arguments

Argument	Type	Access	Mechanism	Status
ucb	UCB	input	reference	required

ucb
Unit control block.

Context

A driver calls IOC_STD\$DELETE_UCB with the I/O database locked for write access.

Description

For Digital internal use only.

Macro

CALL_DELETE_UCB

In a Step 2 driver, CALL_DELETE_UCB simulates a JSB to IOC\$DELETE_UCB. It calls IOC_STD\$DELETE_UCB using the current contents of R5 as the **ucb** argument.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$DELETE_UCB replaces IOC\$DELETE_UCB (used by Step 1 drivers and OpenVMS VAX drivers).

IOC_STD\$DIAGBUFILL

Fills a diagnostic buffer if the original \$QIO request specified such a buffer.

Module

IOSUBNPAG

Format

IOC_STD\$DIAGBUFILL (driver_param, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
driver_param	unspecified	input	reference	required
ucb	UCB	input	reference	required

driver_param

Parameter to be passed to the driver's register dumping routine. Typically, a driver supplies the address of a CRAM in this register.

ucb

Unit control block. IOC_STD\$DIAGBUFILL reads the final error retry count from UCBSL_ERTCNT. It obtains the address of the current IRP from UCBSL_IRP and reads the following IRP fields:

Field	Contents
IRPSL_STS	IRPSV_DIAGBUF set if a diagnostic buffer exists
IRPSL_DIAGBUF	Address of diagnostic buffer, if one is present

IOC_STD\$DIAGBUFILL obtains the address of the DDB from UCBSL_DDB and the address of the DDT from DDBSL_DDT. The procedure value of driver's register dumping routine is obtained from DDTSL_REGDUMP.

Context

The caller of IOC_STD\$DIAGBUFILL may be executing at or above fork IPL and must hold the corresponding fork lock. IOC_STD\$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.

System Routines

IOC_STD\$DIAGBUFILL

Description

A device driver fork process calls IOC_STD\$DIAGBUFILL at the end of I/O processing but before releasing the I/O channel. IOC_STD\$DIAGBUFILL stores the I/O completion time and the final error retry count in the diagnostic buffer. (IOC_STD\$INITIATE has already placed the I/O initiation time [from EXESGQ_SYSTIME] in the first quadword of the buffer.) IOC_STD\$DIAGBUFILL then calls the driver's register dumping routine, passing to it in the **buffer** argument an address within the diagnostic buffer in which the routine can place the register values it retrieves from device interface register space by means of hardware mailbox read transactions. It also passes the contents of the **driver_param** and **ucb** arguments. The register dumping routine fills the remainder of the buffer, and returns to IOC_STD\$DIAGBUFILL, which returns to its caller.

Macro

CALL_DIAGBUFILL

In a Step 2 driver, the CALL_DIAGBUFILL macro simulates a JSB to IOC\$DIAGBUFILL in a Step 1 driver. CALL_DIAGBUFILL calls IOC_STD\$DIAGBUFILL, using the current contents of R4 and R5 as the **driver_parm** and **ucb** arguments, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$DIAGBUFILL replaces IOC\$DIAGBUFILL (used by Step 1 and OpenVMS VAX drivers).
- Prior to calling IOC_STD\$DIAGBUFILL, the driver places a parameter, which the routine passes to the driver's register dumping routine, in R4. On OpenVMS AXP systems, this parameter is often the address of a CRAM (obtained, for instance, from UCB\$PS_CRAM or CRB\$PS_CRAM). On OpenVMS VAX systems, the parameter similarly would contain the address of the device's CSR.
- The contents of R2 and R3 are destroyed when the caller of IOC_STD\$DIAGBUFILL regains control; on OpenVMS VAX systems, these registers contain the DDT address and IRP address respectively.

System Routines

IOC_STD\$FILSPT

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$FILSPT replaces IOC\$FILSPT (used by Step 1 drivers and OpenVMS VAX drivers).

System Routines

IOC_STD\$GETBYTE

Macro

CALL_GETBYTE

In a Step 2 driver, CALL_GETBYTE simulates a JSB to IOC\$GETBYTE. CALL_GETBYTE calls IOC_STD\$GETBYTE, passing the current contents of R0 and R5 as the **sva** and **ucb** arguments, respectively. It returns in R0 the byte of data (not zero-extended) returned from the user buffer. It returns in R1 the updated system virtual address.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note that IOC_STD\$GETBYTE replaces IOC\$GETBYTE (used by Step 1 drivers and OpenVMS VAX drivers). Unlike IOC\$GETBYTE, IOC_STD\$GETBYTE returns the byte of data (and not the updated system virtual address) in R0.

System Routines

IOC_STD\$INITBUFWIND

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$INITBUFWIND replaces IOCS\$INITBUFWIND (used by Step 1 drivers and OpenVMS VAX drivers).

IOC_STD\$INITIATE

Initiates the processing of the next I/O request for a device unit.

Module

IOSUBNPAG

Format

IOC_STD\$INITIATE (irp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

irp

I/O request packet. IOC_STD\$INITIATE reads the following IRP fields:

Field	Contents
IRP\$S_SVAPTE	Address of system buffer (buffered I/O) or system virtual address of the PTE that maps process buffer (direct I/O).
IRP\$S_BOFF	Byte offset of start of buffer.
IRP\$S_BCNT	Size in bytes of transfer.
IRP\$W_STS	IRP\$V_DIAGBUF set if a diagnostic buffer exists.
IRP\$S_DIAGBUF	Address of diagnostic buffer, if one is present. IOC_STD\$INITIATE writes the current system time from EXE\$GQ_SYSTIME into the first quadword of this buffer.

ucb

Unit control block. IOC_STD\$INITIATE reads the following UCB fields:

Field	Contents
UCB\$S_DDB	Address of DDB.
UCB\$S_DDT	Address of DDT. DDT\$PS_START contains the procedure value of the driver's start-I/O routine.
UCB\$S_AFFINITY	Device's affinity mask.

IOC_STD\$INITIATE writes the following UCB fields:

System Routines

IOC_STD\$INITIATE

Field	Contents
UCB\$L_IRP	Address of IRP
UCB\$L_SVAPTE	IRP\$L_SVAPTE
UCB\$L_BOFF	IRP\$L_BOFF
UCB\$L_BCNT	IRP\$L_BCNT
UCB\$L_STS	UCB\$V_CANCEL and UCB\$V_TIMEOUT cleared

Context

IOC_STD\$INITIATE is called at fork IPL with the corresponding fork lock held in a multiprocessing system. Within this context, it transfers control to the driver's start-I/O routine.

Description

IOC_STD\$INITIATE creates the context in which a driver fork process services an I/O request. IOC_STD\$INITIATE creates this context and activates the driver's start-I/O routine in the following steps:

1. 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_STD\$INITIATE takes steps to initiate the I/O function on another processor in a multiprocessing system. It then returns to its caller.
2. Stores the address of the current IRP in UCB\$L_IRP.
3. Copies the transfer parameters contained in the IRP into the UCB:
 - a. 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
 - b. Copies the byte offset within the page from IRP\$L_BOFF to UCB\$L_BOFF
 - c. Copies the byte count from IRP\$L_BCNT to UCB\$L_BCNT
4. Clears the cancel-I/O and timeout bits in the UCB status longword (UCB\$V_CANCEL and UCB\$V_TIMEOUT in UCB\$L_STS).
5. If the I/O request specifies a diagnostic buffer, as indicated by IRP\$V_DIAGBUF in IRP\$L_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).
6. Transfers control to the driver's start-I/O routine.

Macro

CALL_INITIATE

In a Step 2 driver, the CALL_INITIATE macro simulates a JSB to IOC\$INITIATE in a Step 1 driver. CALL_INITIATE calls IOC_STD\$INITIATE, using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$INITIATE replaces IOC\$INITIATE (used by Step 1 drivers and OpenVMS VAX drivers).

System Routines

IOC_STD\$LINK_UCB

IOC_STD\$LINK_UCB

Searches the UCB list attached to the device data block identified by the specified UCB and links the specified UCB into the list in ascending unit number order.

Module

UCBCREDEL

Format

status = IOC_STD\$LINK_UCB (ucb)

Arguments

Argument	Type	Access	Mechanism	Status
ucb	UCB	input	reference	required

ucb
Unit control block.

Return Values

SS\$NORMAL	Link operation was successful.
SS\$OPINCOMPL	Link operation failed due to the presence of a UCB with the same unit number as the specified UCB.

Context

A driver calls IOC_STD\$LINK_UCB with the I/O database locked for write access.

Description

For Digital internal use only.

Macro

CALL_LINK_UCB [interface_warning=YES]

where:

interface_warning=YES, the default, specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. **interface_warning=NO** suppresses the warning.

In a Step 2 driver, `CALL_LINK_UCB` simulates a JSB to `IOC$LINK_UCB`. It calls `IOC_STD$LINK_UCB` using the current contents of R5 as the **ucb** argument. `CALL_LINK_UCB` returns status in R0 and the address of the newly created UCB in R2, but does not return the address of the UCBs that precede and follow it on the DDB chain in R3 and R1, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- `IOC_STD$LINK_UCB` replaces `IOC$LINK_UCB` (used by Step 1 drivers and OpenVMS VAX drivers). `IOC_STD$LINK_UCB` does not return the addresses of the UCBs that precede and follow the newly-created UCB on the DDB chain.

System Routines

IOC_STD\$MAPVBLK

IOC_STD\$MAPVBLK

Maps a virtual block to a logical block using a mapping window.

Module

IOSUBRAMS

Format

status = IOC_STD\$MAPVBLK (vbn, numbytes, wcb, irp, ucb, lbn_p, notmapped_p, new_ucb_p)

Arguments

Argument	Type	Access	Mechanism	Status
vbn	integer	input	value	required
numbytes	integer	input	value	required
wcb	WCB	input	reference	required
irp	IRP	input	reference	required
ucb	UCB	input	reference	required
lbn_p	pointer	output	value	required
notmapped_p	pointer	output	value	required
new_ucb_p	pointer	output	value	required

vbn

Virtual block number.

numbytes

Number of bytes to map.

wcb

Window control block.

irp

I/O request packet.

ucb

Unit control block.

lbn_p

Address at which IOC_STD\$MAPVBLK writes the logical block number of the first block it maps.

notmapped_p

Address at which IOC_STD\$MAPVBLK writes the number of unmapped bytes.

new_ucb_p

Address at which IOC_STD\$MAPVBLK writes the address of the updated UCB.

Return Values

status	Low bit set indicates partial map with all output parameters valid, low bit clear indicates total mapping failure with only the notmapped_p parameter valid.
--------	---

Context

IOC_STD\$MAPVBLK raises IPL to IPL\$_FILSYS and obtains the corresponding spin lock to perform the mapping. As a result, it cannot be called by a driver executing above IPL 8, or by a driver is executing at IPL 8 but holds the IOLOCK8 fork lock.

Description

For Digital internal use only.

Macro

CALL_MAPVBLK

In a Step 2 driver, the CALL_MAPVBLK macro simulates a JSB to IOC\$MAPVBLK in a Step 1 driver. CALL_MAPVBLK calls IOC_STD\$MAPVBLK, using the current contents of R0, R1, R2, R3, and R5 as the **vbn**, **numbytes**, **wcb**, **irp** and **ucb** arguments, respectively. It returns status in R0, the address of the logical block number of the first block mapped in R1, the number of unmapped bytes in R2, and the address of the updated UCB in R3. If the low bit of the status value in R0 is clear, signifying failure status, only the value in R2 is valid.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$MAPVBLK replaces IOC\$MAPVBLK (used by Step 1 and OpenVMS VAX drivers). Unlike IOC\$MAPVBLK, IOC_STD\$MAPVBLK does not preserve R3 across the call.

IOC_STD\$MNTVER

Assists a driver with mount verification.

Module

IOSUBNPAG

Format

IOC_STD\$MNTVER (irp, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
ucb	UCB	input	reference	required

irp

I/O request packet, or 0. If **irp** contains the address of an IRP, EXE_STD\$MNTVER inserts the IRP at the head of the pending-I/O queue in the UCB. If it contains zero, EXE_STD\$MNTVER removes the IRP from the head of the pending-I/O queue and attempts to initiate I/O processing.

ucb

Unit control block.

Context

IOC_STD\$MNTVER is called at fork IPL with the corresponding fork lock held in a multiprocessing system.

Description

For Digital internal use only.

Macro

CALL_MNTVER

In a Step 2 driver, the CALL_MNTVER macro simulates a JSB to IOCSMNTVER in a Step 1 driver. CALL_MNTVER calls IOC_STD\$MNTVER, using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$MNTVER replaces IOCSMNTVER (used by Step 1 and OpenVMS VAX drivers).

IOC_STD\$MOVFRUSER, IOC_STD\$MOVFRUSER2

Move data from a user buffer to an internal buffer.

Module

BUFFERCTL

Format

pointer = IOC_STD\$MOVFRUSER (sysbuf, numbytes, ucb, sysbuf_p)

pointer = IOC_STD\$MOVFRUSER2 (sysbuf, numbytes, ucb, sva, sysbuf_p)

Arguments

Argument	Type	Access	Mechanism	Status
sysbuf	address	input	reference	required
numbytes	integer	input	value	required
ucb	UCB	input	reference	required
sva	address	input	reference	required
sysbuf_p	pointer	output	value	required

sysbuf

Address of internal buffer.

numbytes

Number of bytes to move.

ucb

Unit control block. IOC_STD\$MOVFRUSER and IOC_STD\$MOVFRUSER2 read the following UCB fields:

Field	Contents
UCB\$\$_SVAPTE	System virtual address of PTE that maps the first page of the user buffer
UCB\$\$_SVPN	System virtual page number of SPTE allocated to driver
UCB\$\$_BOFF	Byte offset within the first page to start of user buffer (IOC_STD\$MOVFRUSER only)

sva

System virtual address of the byte in the user buffer after the last byte moved (IOC_STD\$MOVFRUSER2 only).

bufptr

System virtual address of the byte in the user buffer after the last byte moved. IOC_STD\$MOVFRUSER and IOC_STD\$MOVFRUSER2 write this field.

System Routines

IOC_STD\$MOVFRUSER, IOC_STD\$MOVFRUSER2

Return Values

pointer	System virtual address of the byte in the internal buffer after the last byte moved.
---------	--

Context

The caller of IOC_STD\$MOVFRUSER or IOC_STD\$MOVFRUSER2 may be executing at fork IPL or above and must hold the corresponding fork lock in a 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_STD\$MOVFRUSER and IOC_STD\$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).

To use either routine, the driver must have set bit DPT\$V_SVP in the driver prologue table, typically by using the **flags** argument of the DPTAB macro. This causes OpenVMS to allocate a system page-table entry (SPTE) for driver use. (See the description of the DPTAB macro in Chapter 4 for additional information.)

In order to accomplish the move, IOC_STD\$MOVFRUSER and IOC_STD\$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.

IOC_STD\$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_STD\$MOVFRUSER. For each subsequent piece, the driver calls IOC_STD\$MOVFRUSER2.

Macro

CALL_MOVFRUSER
CALL_MOVFRUSER2

In a Step 2 driver, CALL_MOVFRUSER and CALL_MOVFRUSER2 simulate a JSB to IOC\$MOVFRUSER and IOC\$MOVFRUSER2 respectively. CALL_MOVFRUSER calls IOC_STD\$MOVFRUSER, and CALL_MOVFRUSER2 calls IOC_STD\$MOVFRUSER2, passing the current contents of R1, R2, and R5 as the **sysbuf**, **numbytes**, and **ucb** arguments. \$MOVFRUSER2 also passes the current contents of R0 as the **sva** argument. Both macros return in R0 and R1, respectively, the system virtual addresses of the bytes in the internal buffer and user buffer after the last byte moved.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- **IOC_STD\$MOVFRUSER** and **IOC_STD\$MOVFRUSER2** replace **IOC\$MOVFRUSER** and **IOC\$MOVFRUSER2** (used by Step 1 drivers and OpenVMS VAX drivers). Unlike the corresponding OpenVMS VAX routines, both OpenVMS AXP routines destroy R1 across the call. The order in which formal parameters are passed to **IOC_STD\$MOVFRUSER2** differs from the order in which they are provided in registers to the Step 1 routine **IOC\$MOVFRUSER2**.

System Routines

IOC_STD\$MOVTOUSER, IOC_STD\$MOVTOUSER2

IOC_STD\$MOVTOUSER, IOC_STD\$MOVTOUSER2

Move data from an internal buffer to a user buffer.

Module

BUFFERCTL

Format

pointer = IOC_STD\$MOVTOUSER (sysbuf, numbytes, ucb, sysbuf_p)

pointer = IOC_STD\$MOVTOUSER2 (sysbuf, numbytes, ucb, sva, sysbuf_p)

Arguments

Argument	Type	Access	Mechanism	Status
sysbuf	address	input	reference	required
numbytes	integer	input	value	required
ucb	UCB	input	reference	required
sva	address	input	reference	required
sysbuf_p	pointer	output	value	required

sysbuf

Address of internal buffer.

numbytes

Number of bytes to move.

ucb

Unit control block. IOC_STD\$MOVTOUSER and IOC_STD\$MOVTOUSER2 read the following UCB fields:

Field	Contents
UCB\$\$_SVAPTE	System virtual address of PTE that maps the first page of the user buffer
UCB\$\$_SVPN	System virtual page number of SPTE allocated to driver
UCB\$\$_BOFF	Byte offset within the first page to start of user buffer (IOC_STD\$MOVTOUSER only)

sva

System virtual address of the byte in the user buffer after the last byte moved (IOC_STD\$MOVTOUSER2 only).

bufptr

System virtual address of the byte in the user buffer after the last byte moved. IOC_STD\$MOVTOUSER and IOC_STD\$MOVTOUSER2 write this field.

Return Values

pointer	System virtual address of the byte in the internal buffer after the last byte moved.
---------	--

Context

The caller of IOC_STD\$MOVTOUSER or IOC_STD\$MOVTOUSER2 may be executing at fork IPL or above and must hold the corresponding fork lock in a 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_STD\$MOVTOUSER and IOC_STD\$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.

To use either routine, the driver must have set bit DPT\$V_SVP in the driver prologue table, typically by using the **flags** argument of the DPTAB macro. This causes OpenVMS to allocate a system page-table entry (SPTE) for driver use. (See the description of the DPTAB macro in Chapter 4 for additional information.)

In order to accomplish the move, IOC_STD\$MOVTOUSER and IOC_STD\$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.

IOC_STD\$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_STD\$MOVTOUSER. For each subsequent buffer to move, the driver calls IOC_STD\$MOVTOUSER2.

Macro

CALL_MOVTOUSER
CALL_MOVTOUSER2

In a Step 2 driver, CALL_MOVTOUSER and CALL_MOVTOUSER2 simulate a JSB to IOC\$MOVTOUSER and IOC\$MOVTOUSER2 respectively. CALL_MOVTOUSER calls IOC_STD\$MOVTOUSER, and CALL_MOVTOUSER2 calls IOC_STD\$MOVTOUSER2, passing the current contents of R1, R2, and R5 as the **sysbuf**, **numbytes**, and **ucb** arguments. CALL_MOVTOUSER2 also passes the current contents of R0 as the **sva** argument. Both macros return in R0 and R1, respectively, the system virtual addresses of the bytes in the internal buffer and user buffer after the last byte moved.

System Routines

IOC_STD\$MOVTOUSER, IOC_STD\$MOVTOUSER2

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- **IOC_STD\$MOVTOUSER** and **IOC_STD\$MOVTOUSER2** replace **IOC\$MOVTOUSER** and **IOC\$MOVTOUSER2** (used by Step 1 drivers and OpenVMS VAX drivers). Unlike the corresponding OpenVMS VAX routines, both OpenVMS AXP routines destroy R1 across the call. The order in which formal parameters are passed to **IOC_STD\$MOVTOUSER2** differs from the order in which they are provided in registers to the Step 1 routine **IOC\$MOVTOUSER2**.

IOC_STD\$PARSDEVNAM

Parses a device name string, checking its syntax and extracting the node name, allocation class number, and unit number.

Module

IOSUBNPAG

Format

status = IOC_STD\$PARSDEVNAM (devnamlen, devnam, flags, unit_p, scslen_p, devnamlen_p, devnam_p, flags_p)

Arguments

Argument	Type	Access	Mechanism	Status
devnamlen	integer	input	value	required
devnam	address	input	reference	required
flags	integer	input	value	required
unit_p	pointer	output	reference	required
scslen_p	pointer	output	reference	required
devnamlen_p	pointer	output	reference	required
devnam_p	pointer	output	reference	required
flags_p	pointer	output	reference	required

devnamlen

Size of the name string.

devnam

Name string.

flags

Flags.

unit_p

Address at which IOC_STD\$PARSDEVNAM writes an integer representing the unit number.

scslen_p

Address at which IOC_STD\$PARSDEVNAM writes an integer representing either the length of the SCS node name, the allocation class number, or the device type code.

devnamlen_p

Address at which IOC_STD\$PARSDEVNAM writes an integer representing the size of the name string.

System Routines

IOC_STD\$PARSDEVNAM

devnam_p

Address at which IOC_STD\$PARSDEVNAM writes the address of the name string.

flags_p

Address at which IOC_STD\$PARSDEVNAM writes an integer that contains the flags.

Return Values

SS\$_IVDEVNAM	Invalid device name string.
SS\$_NORMAL	Valid device name string.

Context

IOC_STD\$PARSDEVNAM is typically called at fork IPL with the corresponding fork lock held in an OpenVMS multiprocessing system.

Description

For Digital internal use only.

Macro

CALL_PARSDEVNAM

In a Step 2 driver, the CALL_PARSDEVNAM macro simulates a JSB to IOC\$PARSDEVNAM in a Step 1 driver. CALL_PARSDEVNAM calls IOC_STD\$PARSDEVNAM, using the current contents of R8, R9, and R10 as the **devnamlen**, **devnam**, and **flags** arguments, respectively. When IOC_STD\$PARSDEVNAM returns, the macro returns status in R0; the unit number in R2; the length of the SCS node name at the beginning of the name string, allocation class number, or device type code in R3; the size of the name string in R8, the address of the name string in R9, and the flags in R10.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$PARSDEVNAM replaces IOC\$PARSDEVNAM (used by Step 1 and OpenVMS VAX drivers). Unlike IOC\$PARSDEVNAM, IOC_STD\$PARSDEVNAM does not preserve the contents of R8, R9, and R10 across the call.

IOC_STD\$POST_IRP

Inserts an I/O request packet in a CPU-specific I/O postprocessing queue.

Module

IOSUBNPAG

Format

IOC_STD\$POST_IRP (irp)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required

irp
I/O request block.

Context

Mount verification processing calls IOC_STD\$POST_IRP at or above IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

CALL_POST_IRP

In a Step 2 driver, CALL_POST_IRP simulates a JSB to IOC\$POST_IRP. It calls IOC_STD\$POST_IRP using the current contents of R3 as the **irp** argument.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$POST_IRP replaces IOC\$POST_IRP (used by Step 1 drivers and OpenVMS VAX drivers).

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- **IOC_STD\$PTETOPFN** replaces **IOC\$PTETOPFN** (used by Step 1 drivers and OpenVMS VAX drivers). Note that, the page-table entry input argument is passed by value (in R3) to **IOC\$PTETOPFN**, but passed by reference to **IOC_STD\$PTETOPFN**.

System Routines

IOC_STD\$QNXTSEG1

IOC_STD\$QNXTSEG1

Queues the next segment of a virtual I/O request that did not map to a single contiguous I/O request.

Module

IOCIOPST

Format

IOC_STD\$QNXTSEG1 (vbn, bcnt, wcb, irp, pcb, ucb, ucb_p)

Arguments

Argument	Type	Access	Mechanism	Status
vbn	integer	output	value	required
bcnt	integer	output	value	required
wcb	WCB	output	reference	required
irp	IRP	output	reference	required
pcb	PCB	output	reference	required
ucb	UCB	output	reference	required
ucb_p	pointer	input	reference	required

vbn

Virtual block number of the start of the next segment.

bcnt

Required byte count of next segment.

wcb

Window control block.

irp

I/O request packet.

pcb

Process control block.

ucb

Unit control block.

ucb_p

Address at which IOC_STD\$QNXTSEG1 writes the address of the unit control block.

Context

The caller of IOC_STD\$QNXTSEG1 typically executes at or above fork IPL. IOC_STD\$QNXTSEG1 executes at its caller's IPL and returns control at that IPL. The caller retains any spin locks it held at the time of the call.

Description

For Digital internal use only.

Macro

CALL_QNXTSEG1

In a Step 2 driver, CALL_QNXTSEG1 simulates a JSB to IOC\$QNXTSEG1. It calls IOC_STD\$QNXTSEG1 using the current contents of R0, R1, R2, R3, R4, and R5 as the **vbn**, **bcnt**, **wcb**, **irp**, **pcb**, and **ucb** arguments. It returns the address of the updated UCB in R5.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$QNXTSEG1 replaces IOC\$QNXTSEG1 (used by Step 1 drivers and OpenVMS VAX drivers). Unlike IOC\$QNXTSEG1, IOC_STD\$QNXTSEG1 does not return the address of the updated UCB in R5.

System Routines

IOC_STD\$PRIMITIVE_REQCHANH, IOC_STD\$PRIMITIVE_REQCHANL

IOC_STD\$PRIMITIVE_REQCHANH, IOC_STD\$PRIMITIVE_REQCHANL

Request a controller's data channel and, if unavailable, place process in channel wait queue.

Module

IOSUBNPAG

Format

status = IOC_STD\$PRIMITIVE_REQCHANH (irp, ucb, idb_p)

status = IOC_STD\$PRIMITIVE_REQCHANL (irp, ucb, idb_p)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
ucb	UCB	input	reference	required
idb_p	pointer	output	reference	required

irp

I/O request packet.

ucb

Unit control block. IOC_STD\$PRIMITIVE_REQPCHANH and IOC_STD\$PRIMITIVE_REQPCHANL read the following UCB fields:

Field	Contents
UCB\$L_FPC	Procedure value of fork routine to be executed when the channel is granted if the channel cannot be granted immediately

IOC_STD\$PRIMITIVE_REQCHANH, IOC_STD\$PRIMITIVE_REQCHANL

Field	Contents
UCB\$\$_CRB	Address of controller request block (CRB). IOC_STD\$REQCHANH and IOC_STD\$REQCHANL access the following CRB fields:
Field	Contents
CRB\$_MASK	CRB\$_BSY set if the channel is busy
CRB\$_INTD+VEC\$_IDB	Address of IDB
CRB\$_WQFL	Head of queue of UCBs waiting for the controller channel
CRB\$_WQBL	Tail of queue of UCBs waiting for the controller channel

IOC_STD\$REQCHANH and IOC_STD\$REQCHANL write the contents of the **irp** parameter in UCB\$_FR3, and the address of the UCB in IDB\$_OWNER.

If the channel is busy, IOC_STD\$REQCHANH and IOC_STD\$REQCHANL update CRB\$_WQFL and CRB\$_WQBL.

idb_p

Address of location in which IOC_STD\$REQCHANH and IOC_STD\$REQCHANL write the address of the interrupt dispatch block (IDB).

Return Values

SS\$_NORMAL	Channel has been granted immediately.
0	Channel is busy and UCB fork block has been queued on channel-wait queue.

Context

A driver calls IOC_STD\$PRIMITIVE_REQCHANH or IOC_STD\$PRIMITIVE_REQCHANL at fork IPL holding the appropriate fork lock. Either IOC_STD\$PRIMITIVE_REQCHANH or IOC_STD\$PRIMITIVE_REQCHANL, unlike the corresponding OpenVMS VAX system routine, returns to its caller and not to its caller's caller. Each assumes that, prior to the call, its caller has placed the procedure value of the fork routine into UCB\$_FPC.

If the requested channel is busy, either IOC_STD\$PRIMITIVE_REQCHANH or IOC_STD\$PRIMITIVE_REQCHANL preserves the contents of the **irp** parameter in UCB\$_FR3 . IOC_STD\$RELCHAN eventually calls the fork routine upon granting the channel request, passing the **irp**, **idb**, and **ucb** parameters.

System Routines

IOC_STD\$PRIMITIVE_REQCHANH, IOC_STD\$PRIMITIVE_REQCHANL

Description

A driver fork process calls IOC_STD\$PRIMITIVE_REQCHANH or IOC_STD\$PRIMITIVE_REQCHANL to acquire ownership of the controller's 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\$PS_OWNER and returning successfully with the IDB address in the location specified by the **idb_p** parameter.

If the data channel is busy, the routine saves process context by placing the IRP address, as specified in the **irp** parameter, into the UCB fork block. IOC_STD\$REQCHANH then inserts the UCB at the head of the channel wait queue (CRB\$SL_WQFL); IOC_STD\$REQCHANL inserts the UCB at the tail of the queue (CRB\$SL_WQBL). Finally, the routine returns control to its caller.

When the controller channel is available to a waiting fork process, IOC_STD\$RELCHAN resumes the suspended fork process at its channel grant routine, passing to it the **irp**, **idb**, and **ucb** parameters.

Macro

REQCHAN

REQPCHAN

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$PRIMITIVE_REQPCHANH and IOC_STD\$PRIMITIVE_REQPCHANL replace the Step 1 routines IOC\$PRIMITIVE_REQPCHANH and IOC\$PRIMITIVE_REQPCHANL. The OpenVMS VAX routines IOC\$REQPCHAN and IOC\$REQPCHANL are not provided on OpenVMS AXP systems.

IOC_STD\$PRIMITIVE_WFIKPCH, IOC_STD\$PRIMITIVE_WFIRLCH

Suspend a driver fork thread and fold its context into a fork block in anticipation of a device interrupt or timeout.

Module

IOSUBNPAG

Format

IOC_STD\$PRIMITIVE_WFIKPCH (irp, fr4, ucb, tmo, restore_ipl)

IOC_STD\$PRIMITIVE_WFIRLCH (irp, fr4, ucb, tmo, restore_ipl)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
fr4	int64	input	value	required
ucb	UCB	input	reference	required
tmo	integer	input	value	required
restore_ipl	int	input	value	required

irp
I/O request packet.

fr4
Parameter to be passed to the interrupt service routine or timeout handling routine.

ucb
Unit control block. IOC_STD\$PRIMITIVE_WFIKPCH and IOC_STD\$PRIMITIVE_WFIRLCH read the following UCB fields:

Field	Contents
UCB\$L_FPC	Procedure value of fork routine which may be the destination of a JSB instruction issued by either the driver's interrupt service routine or EXE\$TIMEOUT
UCB\$B_FLCK	Fork lock index

IOC_STD\$PRIMITIVE_WFIKPCH and IOC_STD\$PRIMITIVE_WFIRLCH write the following UCB fields:

Field	Contents
UCB\$L_DUETIM	Sum of timeout value and EXE\$GL_ABSTIM

System Routines

IOC_STD\$PRIMITIVE_WFIKPCH, IOC_STD\$PRIMITIVE_WFIRLCH

Field	Contents
UCB\$L_STS	UCB\$V_INT is set to indicate that interrupts are expected on the device; UCB\$V_TIM is set to indicate device I/O is being timed; and UCB\$V_TIMEOUT is cleared to indicate that unit has not yet timed out.
UCB\$Q_FR3	R3 of caller
UCB\$Q_FR4	R4 of caller

tmo

Timeout value in seconds.

restore_ipl

IPL to which to lower before returning to caller. This IPL must be the fork IPL associated with device processing and at which the driver was executing prior to invoking the DEVICELOCK macro.

Context

When it is called, IOC_STD\$PRIMITIVE_WFIKPCH or IOC_STD\$PRIMITIVE_WFIRLCH assumes that the local processor has obtained the appropriate synchronization with the device database by securing 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.

Before exiting, IOC_STD\$PRIMITIVE_WFIKPCH or IOC_STD\$PRIMITIVE_WFIRLCH 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 it regains control. IOC_STD\$PRIMITIVE_WFIKPCH or IOC_STD\$PRIMITIVE_WFIRLCH also lowers the local processor's IPL to the IPL specified in the **restore_ipl** parameter.

Description

A driver fork process calls IOC_STD\$PRIMITIVE_WFIKPCH to wait for an interrupt while keeping ownership of the controller's data channel; IOC_STD\$PRIMITIVE_WFIRLCH, by contrast, releases the channel.

Either routine performs the following operations:

1. Moves contents of the **irp** and **fr4** parameters into the UCB fork block.
2. Sets UCB\$V_INT to indicate an expected interrupt from the device unit.
3. Sets UCB\$V_TIM to indicate that OpenVMS should check for timeouts from the device unit.
4. Determines the timeout due time by adding the timeout value specified in R1 to EXE\$GL_ABSTIM and storing the result in UCB\$L_DUETIM.
5. Clears UCB\$V_TIMEOUT to indicate that the unit has not yet timed out.
6. Invokes the DEVICEUNLOCK macro to conditionally release the device lock associated with the device unit and to lower IPL to the IPL specified in the **restore_ipl** parameter. These actions presume that the DEVICELOCK macro has been issued prior to the wait-for-interrupt invocation.

System Routines

IOC_STD\$PRIMITIVE_WFIKPCH, IOC_STD\$PRIMITIVE_WFIRLCH

7. Returns to its caller.

Note that IOC_STD\$PRIMITIVE_WFIRLCH exits by transferring control to IOC_STD\$RELCHAN. IOC_STD\$RELCHAN releases the controller data channel and eventually issues an RSB instruction to IOC_STD\$PRIMITIVE_WFIRLCH which returns to its caller. Because the release of the channel occurs at fork IPL, an interrupt service routine cannot reliably distinguish between operations initiated by IOC_STD\$PRIMITIVE_WFIKPCH and IOC_STD\$PRIMITIVE_WFIRLCH by examining the ownership of the CRB.

Macro

WFIKPCH

WFIRLCH

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note that IOC_STD\$PRIMITIVE_WFIKPCH and IOC_STD\$PRIMITIVE_WFIRLCH replace the Step 1 routines IOC\$PRIMITIVE_WFIKPCH and IOC\$PRIMITIVE_WFIRLCH.

The OpenVMS VAX routines IOC\$WFIKPCH and IOC\$WFIRLCH are not provided on OpenVMS AXP systems.

IOC_STD\$RELCHAN

Releases device ownership of all controller data channels.

Module

IOSUBNPAG

Format

IOC_STD\$RELCHAN (ucb)

Arguments

Argument	Type	Access	Mechanism	Status
ucb	UCB	input	reference	required

ucb

Unit control block. IOC_STD\$RELCHAN reads UCB\$\$_CRB to obtain the address of the controller request block (CRB) in order to access the following CRB fields:

Field	Contents
CRB\$_MASK	CRB\$_BSY set if the channel is busy. IOC_STD\$RELCHAN clears this bit if no driver is waiting for the controller channel.
CRB\$_INTD+VEC\$_IDB	Address of IDB. IOC_STD\$RELCHAN obtains the address the UCB that owns the controller channel from IDB\$_OWNER. IOC_STD\$RELCHAN clears IDB\$_OWNER if no driver is waiting for the controller channel.
CRB\$_WQFL	Head of queue of UCBs waiting for the controller.

Context

A driver fork process calls IOC_STD\$RELCHAN at fork IPL, holding the corresponding fork lock in a multiprocessing environment. IOC_STD\$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_STD\$RELCHAN to release all controller data channels assigned to a device.

If the channel wait queue contains waiting fork processes, IOC_STD\$RELCHAN dequeues a process, assigns the channel to that process and calls the suspended fork process at its channel grant routine, passing to it the **irp**, **idb**, and **ucb** parameters.

Macro

CALL_RELCHAN

In a Step 2 driver, CALL_RELCHAN simulates a JSB to IOC\$RELCHAN. It calls IOC_STD\$RELCHAN using the current contents of R5 as the **ucb** argument.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$RELCHAN replaces IOC\$RELCHAN (used by Step 1 drivers and OpenVMS VAX drivers).
- IOC\$RELCHAN resumes the fork routine with the address of a device's controller and status register (CSR) in R4. Because OpenVMS AXP device drivers access device CSRs by means of a controller register access mailbox (CRAM), IOC_STD\$RELCHAN provides the IDB address as input to the reactivated fork routine. The fork routine uses the IDB address as input to the driver macros and routines that manipulate CSRs by means of the CRAM.

IOC_STD\$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

Format

IOC_STD\$REQCOM (iost1, iost2, ucb)

Arguments

Argument	Type	Access	Mechanism	Status
iost1	integer	input	value	required
iost2	integer	input	value	required
ucb	UCB	input	reference	required

iost1

First longword of I/O status.

iost2

Second longword of I/O status.

ucb

Unit control block. IOC_STD\$REQCOM accesses the following UCB fields:

Field	Contents
UCB\$_ERTCNT	Final error count.
UCB\$_ERTMAX	Maximum error retry count.
UCB\$_EMB	Address of error message buffer.
UCB\$_IRP	Address of IRP. IOC_STD\$REQCOM writes iost1 and iost2 into IRP\$_IOST1 and IRP\$_IOST2, respectively.
UCB\$_DEVCLASS	DC\$_DISK and DC\$_TAPE devices are subject to mount verification checks.
UCB\$_IOQFL	Device unit's pending-I/O queue. IOC_STD\$REQCOM updates this field.

Field	Contents										
UCB\$L_STS	<p>If error logging is in progress (that is, UCB\$V_ERLOGIP is set), IOC_STD\$REQCOM writes the following fields in the error message buffer:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Field</th> <th style="text-align: left;">Contents</th> </tr> </thead> <tbody> <tr> <td>EMB\$L_DV_STS</td> <td>UCB\$L_STS.</td> </tr> <tr> <td>EMB\$L_DV_ERTCNT</td> <td>UCB\$L_ERTCNT.</td> </tr> <tr> <td>EMB\$L_DV_ERTCNT+1</td> <td>UCB\$L_ERTMAX.</td> </tr> <tr> <td>EMB\$Q_DV_IOSB</td> <td>Quadword of I/O status.</td> </tr> </tbody> </table> <p>IOC_STD\$REQCOM then clears UCB\$V_BSY and UCB\$V_ERLOGIP.</p>	Field	Contents	EMB\$L_DV_STS	UCB\$L_STS.	EMB\$L_DV_ERTCNT	UCB\$L_ERTCNT.	EMB\$L_DV_ERTCNT+1	UCB\$L_ERTMAX.	EMB\$Q_DV_IOSB	Quadword of I/O status.
Field	Contents										
EMB\$L_DV_STS	UCB\$L_STS.										
EMB\$L_DV_ERTCNT	UCB\$L_ERTCNT.										
EMB\$L_DV_ERTCNT+1	UCB\$L_ERTMAX.										
EMB\$Q_DV_IOSB	Quadword of I/O status.										
UCB\$L_OPCNT	Unit operations count. IOC_STD\$REQCOM increases this field.										

Context

A driver fork process calls IOC_STD\$REQCOM at fork IPL, holding the corresponding fork lock in a multiprocessing environment. IOC_STD\$REQCOM transfers control to IOC_STD\$RELCHAN, which may call the OpenVMS fork dispatcher to resume another driver fork process. When it regains control, IOC_STD\$REQCOM returns to the driver fork process.

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.

IOCS\$REQCOM performs the following tasks:

1. 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_STD\$RELEASEMB.
2. Increases the device unit's operations count (UCB\$L_OPCNT).
3. 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 (EOF) errors (SS\$_ENDOFFILE) are exempt from these checks. For a tape device with success status, checks to determine if CRC must be generated.
4. Writes final I/O status (R0 and R1) into IRP\$L_IOST1 and IRP\$L_IOST2.
5. Inserts the IRP in systemwide I/O postprocessing queue.
6. Requests a software interrupt from the local processor at IPL\$_IOPOST.
7. Attempts to remove an IRP from the device's pending-I/O queue (at UCB\$L_IOQFL). If successful, it transfers control to IOC_STD\$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.

System Routines

IOC_STD\$REQCOM

8. Exits by transferring control to IOC_STD\$RELCHAN.

Macro

CALL_REQCOM

In a Step 2 driver, the CALL_REQCOM macro simulates a JSB to IOCSREQCOM in a Step 1 driver. CALL_REQCOM calls IOC_STD\$REQCOM, using the current contents of R0, R1, and R5 as the **iost1**, **iost2**, and **ucb** arguments, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$REQCOM replaces IOCSREQCOM (used by Step 1 and OpenVMS VAX drivers). Unlike IOCSREQCOM, IOC_STD\$REQCOM does not return the addresses of the IRP and UCB in R3 and R5, respectively.
- The Step 1 REQCOM macro issues a JMP instruction to IOCSREQCOM, the effect of which is to effect a return to the caller of the driver fork process when IOCSREQCOM returns. The Step 2 REQCOM macro returns control to the driver fork process, which itself must issue the return to its caller.

IOC_STD\$SEARCHDEV

Searches the I/O database for a specific physical device.

Module

IOSUBPAGD

Format

status = IOC_STD\$SEARCHDEV (descr_p, ucb_p, ddb_p, sb_p)

Arguments

Argument	Type	Access	Mechanism	Status
descr_p	pointer	input	reference	required
ucb_p	pointer	output	reference	required
ddb_p	pointer	output	reference	required
sb_p	pointer	output	reference	required

descr_p

Descriptor of device logical name.

ucb_p

Address at which IOC_STD\$SEARCHDEV writes the unit control block (UCB) address.

ddb_p

Address at which IOC_STD\$SEARCHDEV writes the device data block (DDB) address.

sb_p

Address at which IOC_STD\$SEARCHDEV writes the system block (SB) address.

Return Values

SS\$_ACCVIO	Name string is not readable.
SS\$_DEVALLOC	Device is allocated to another user.
SS\$_DEVMOUNT	Device already mounted.
SS\$_DEVOFFLINE	Device marked offline.
SS\$_IVDEVNAM	Invalid device name string.
SS\$_IVLOGNAM	Invalid logical name.
SS\$_NODEVAVL	Device exists but is not available.
SS\$_NONLOCAL	Nonlocal device.
SS\$_NOPRIV	Insufficient privilege to access device.

System Routines

IOC_STD\$SEARCHDEV

SS\$ _NORMAL	Device found.
SS\$ _NOSUCHDEV	Device not found.
SS\$ _TEMPLATEDEV	Cannot allocate template device.
SS\$ _TOOMANYLNAM	Maximum logical name recursion limit exceeded.

Context

A driver calls IOC_STD\$SEARCHDEV at IPL\$_ASTDEL holding the I/O database mutex.

Description

For Digital internal use only.

Macro

CALL_SEARCHDEV

In a Step 2 driver, the CALL_SEARCHDEV macro simulates a JSB to IOC\$SEARCHDEV in a Step 1 driver. CALL_SEARCHDEV calls IOC_STD\$SEARCHDEV, using the current contents of R1 as the **descr_p** argument. When IOC_STD\$SEARCHDEV returns, the macro returns status in R0, the UCB address in R1, the DDB address in R2, and the SB address in R3.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$SEARCHDEV replaces IOC\$SEARCHDEV (used by Step 1 and OpenVMS VAX drivers). Unlike IOC\$SEARCHDEV, IOC_STD\$SEARCHDEV does not provide the addresses of the UCB, DDB, and SB in R1, R2, and R3, respectively.

IOC_STD\$SEARCHINT

Searches the I/O database for the specified device, using specified search rules.

Module

IOSUBNPAG

Format

status = IOC_STD\$SEARCHINT (unit, scslen, devnamlen, devnam, flags, ucb_p, ddb_p, sb_p, lock_val_p)

Arguments

Argument	Type	Access	Mechanism	Status
unit	integer	input	value	required
scslen	integer	input	value	required
devnamlen	integer	input	value	required
devnam	address	input	reference	required
flags	integer	input	value	required
ucb_p	pointer	output	reference	required
ddb_p	pointer	output	reference	required
sb_p	pointer	output	reference	required
lock_val_p	pointer	output	reference	required

unit

Unit number.

scslen

Integer representing either the length of the SCS node name, the allocation class number, or the device type code.

devnamlen

Size of the name string.

devnam

Name string.

flags

Flags.

ucb_p

Address at which IOC_STD\$SEARCHINT writes the UCB address.

ddb_p

Address at which IOC_STD\$SEARCHINT writes the DDB address.

System Routines

IOC_STD\$SEARCHINT

sb_p

Address at which IOC_STD\$SEARCHINT writes the system block (SB) address.

lock_val_p

Address at which IOC_STD\$SEARCHINT writes the address of the lock value block.

Return Values

SS\$_DEVMOUNT	Device already mounted.
SS\$_DEVOFFLINE	Device marked offline.
SS\$_NODEVAVL	Device exists but is not available.
SS\$_NOPRIV	Insufficient privilege to access device.
SS\$_NORMAL	Device found.
SS\$_NOSUCHDEV	Device not found.
SS\$_TEMPLATEDEV	Cannot allocate template device.

Context

A driver calls IOC_STD\$SEARCHINT at IPL\$ASTDEL holding the I/O database mutex. It may be called at elevated IPL only for searches specifying IOC\$V_ANY.

Description

For Digital internal use only.

Macro

CALL_SEARCHINT

In a Step 2 driver, the CALL_SEARCHINT macro simulates a JSB to IOC\$SEARCHINT in a Step 1 driver. CALL_SEARCHINT calls IOC_STD\$SEARCHINT, using the current contents of R2, R3, R8, R9 and R10 as the **unit**, **scslen**, **devnamlen**, **devnam**, and **flags** arguments, respectively. When IOC_STD\$SEARCHINT returns, the macro returns status in R0, the UCB address in R5, the DDB address in R6, and the SB address in R7.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$SEARCHINT replaces IOC\$SEARCHINT (used by Step 1 and OpenVMS VAX drivers). Unlike IOC\$SEARCHINT, IOC_STD\$SEARCHINT does not provide the addresses of the UCB, DDB, and SB in R5, R6, and R7, respectively.

IOC_STD\$SENSEDISK

Copies the disk's size in logical blocks from the device's UCB into the second longword of the I/O status block (IOSB) specified in a \$QIO system service call, and completes the I/O operation successfully.

Module

IOSUBRAMS

Format

status = IOC_STD\$SENSEDISK (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp
I/O request packet for the current I/O request.

pcb
Process control block of the current process.

ucb
Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb
Channel control block that describes the process-I/O channel

Return Values

SS\$FDT_COMPL	Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.
---------------	---

Status in FDT_CONTEXT

SS\$NORMAL	The routine completed successfully.
------------	-------------------------------------

System Routines

IOC_STD\$SENSEDISK

Context

FDT dispatching code in the \$QIO system service calls IOC_STD\$SENSEDISK as an upper-level FDT action routine at IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine IOC\$SENSEDISK (used by OpenVMS VAX and Step 1 device drivers) expects as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to IOC_STD\$SENSEDISK.
- IOC_STD\$SENSEDISK returns control to the system service dispatcher, passing it the final \$QIO system service status (SS\$NORMAL) in R0. IOC_STD\$SENSEDISK returns to its caller, passing it SS\$FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

IOC_STD\$SEVER_UCB

Removes the specified UCB from the UCB list of the device data block identified within the specified UCB.

Module

UCBCREDEL

Format

IOC_STD\$SEVER_UCB (ucb)

Arguments

Argument	Type	Access	Mechanism	Status
ucb	UCB	input	reference	required

ucb
Unit control block.

Context

A driver calls IOC_STD\$SEVER_UCB with the I/O database locked for write access.

Description

For Digital internal use only.

Macro

CALL_SEVER_UCB

In a Step 2 driver, CALL_SEVER_UCB simulates a JSB to IOC\$SEVER_UCB. It calls IOC_STD\$SEVER_UCB using the current contents of R5 as the **ucb** argument.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$SEVER_UCB replaces IOC\$SEVER_UCB (used by Step 1 drivers and OpenVMS VAX drivers).

IOC_STD\$SIMREQCOM

Completes an I/O operation by setting an event flag, modifying an I/O status block (IOSB), setting an event flag, or queuing an AST to the process requesting the I/O. The caller of this routine is responsible for checking quotas and updating the I/O count.

Module

IOCIOPST

Format

status = IOC_STD\$SIMREQCOM (iosb, pri, efn, iost, acb, acmode)

Arguments

Argument	Type	Access	Mechanism	Status
iosb	IOSB	input	reference	optional
pri	integer	input	value	optional
efn	integer	input	value	optional
iost	unspecified	input	unspecified	required
acb	ACB	input	reference	optional
acmode	integer	input	value	optional

iosb

I/O status block. If this parameter contains the address of an IOSB, IOC_STD\$SIMREQCOM checks for write access to the IOSB. If it contains a zero, IOC_STD\$SIMREQCOM makes no IOSB modifications.

pri

Priority boost class to be passed directly to SCH\$POSTEF and SCH\$QAST. If an IOSB address is supplied to the **iosb** parameter, this parameter has no effect. If this parameter contains a zero, there is no priority boost.

efn

Common or local event flag to be set. If this parameter contains -1, no event flag is set.

iost

Internal process identification (IPID) of the target process (if the **iosb** parameter is zero); address of a quadword containing the new contents of the user's IOSB (if the **iosb** is non-zero).

acb

AST control block. If this parameter is zero, no AST is delivered. When the **acb** parameter is non-zero and ACB\$L_AST is zero, IOC_STD\$SIMREQCOM checks ACB\$V_NODELETE. If ACB\$V_NODELETE is clear, IOC_STD\$SIMREQCOM uses ACB\$W_SIZE to return the ACB and any structure in which it is embedded to nonpaged pool.

acmode

Access mode of the process originally requesting the I/O operation. IOC_STD\$SIMREQCOM uses this value to probe the IOSB (if specified) for write access. If the **iosb** parameter is zero, this parameter is ignored.

Return Values

SS\$ILLEFC	Illegal cluster number.
SS\$NONEXPR	Nonexistent process.
SS\$NORMAL	Normal, successful completion.
SS\$UNASEFC	Unassigned cluster number.
SS\$WASCLR	Specified event flag was clear initially.
SS\$WASSET	Specified event flag was set initially.

Context

If supplying a non-zero value for the **iosb** parameter, the caller of IOC_STD\$SIMREQCOM must be executing in the context of the target process.

Description

For Digital internal use only.

Macro

CALL_SIMREQCOM

In a Step 2 driver, the CALL_SIMREQCOM macro simulates a JSB to IOC\$SIMREQCOM in a Step 1 driver. CALL_SIMREQCOM calls IOC_STD\$SIMREQCOM, using the current contents of R1, R2, R3, R4, R5, and R6 as the **iosb**, **pri**, **efn**, **iost**, **acb**, and **acmode** arguments, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- IOC_STD\$SIMREQCOM replaces IOC\$SIMREQCOM (used by Step 1 and OpenVMS VAX drivers).

System Routines

IOC_STD\$THREADCRB

IOC_STD\$THREADCRB

Threads a controller request block (CRB) onto the CRB timeout queue chain headed by IOCSGL_CRBTMOUT.

Module

IOSUBNPAG

Format

IOC_STD\$THREADCRB (crb)

Arguments

Argument	Type	Access	Mechanism	Status
crb	CRB	input	reference	required

crb
Controller request block.

Context

Mount verification processing calls IOC_STD\$THREADCRB at or above IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

CALL_THREADCRB [save_r0]

where:

save_r0 indicates that the macro should preserve register R0 across the call to IOC_STD\$THREADCRB. If **save_r0** is blank or **save_r0=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r0=NO**, R0 is not saved.)

In a Step 2 driver, CALL_THREADCRB simulates a JSB to IOC\$THREADCRB. It calls IOC_STD\$THREADCRB using the current contents of R3 as the **crb** argument. Unless you specify **save_r1=NO**, the macro preserves the quadword register R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- **IOC_STD\$THREADCRB** replaces **IOC\$THREADCRB** (used by Step 1 drivers and OpenVMS VAX drivers). Unlike **IOC\$THREADCRB**, **IOC_STD\$THREADCRB** routines does not preserve R0 across the call.

System Routines

MMG_STD\$IOLOCK

MMG_STD\$IOLOCK

Locks process pages in memory.

Module

IOLOCK

Format

status = MMG_STD\$IOLOCK (buf, bufsize, is_read, pcb, svapte_p)

Arguments

Argument	Type	Access	Mechanism	Status
buf	address	input	reference	required
bufsize	integer	input	value	required
is_read	integer	input	value	required
pcb	PCB	input	reference	required
svapte_p	pointer	output	reference	required

buf
Buffer.

bufsize
Size of output buffer in bytes.

is_read
Transfer direction indicator, as follows:

Value	Description
0	Write from memory to I/O device
1	Read into memory from I/O device
5	Write from and read into memory from I/O device

pcb
Process control block.

svapte_p
Address of location in which MMG_STD\$IOLOCK returns either the system virtual address of the first page-table entry (if the returned status is SSS_NORMAL) or the address of a page to be faulted into memory (if the returned status is 0).

Return Values

SS\$_ACCVIO	Specified buffer is not a process buffer, but does not fully reside in system space; or process buffer overruns balance set slots.
SS\$_INSFWSL	Insufficient working set list.
SS\$_NORMAL	Normal, successful completion.
0	Virtual address must be faulted into memory.

Context

MMG_STD\$IOLOCK must be called at IPL\$_ASTDEL.

Description

For Digital internal use only.

Macro

CALL_IOLOCK

In a Step 2 driver, CALL_IOLOCK simulates a JSB to MMG\$IOLOCK. It calls MMG_STD\$IOLOCK using the current contents of R0, R1, R2, and R4 as the **buf**, **bufsize**, **is_read**, and **pcb** arguments, respectively.

CALL_IOLOCK returns status in R0. If R0 contains SS\$_NORMAL, R1 contains the system virtual address of the first page-table entry. If R0 contains zero, R1 contains the address of a page to be faulted into memory. R0 can also contain a system-level status.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- MMG_STD\$IOLOCK replaces MMG\$IOLOCK (used by Step 1 drivers and OpenVMS VAX drivers).

MMG_STD\$UNLOCK

Unlocks process pages previously locked for a direct-I/O operation.

Module

IOLOCK

Format

MMG_STD\$UNLOCK (npages, svapte)

Arguments

Argument	Type	Access	Mechanism	Status
npages	integer	input	value	required
svapte	integer	input	value	required

npages

Number of buffer pages to unlock.

svapte

System virtual address of PTE for the first buffer page.

Context

Because MMG_STD\$UNLOCK raises IPL to IPL\$_SYNCH, and obtains the MMG spin lock in a multiprocessing environment, its caller cannot be executing above IPL\$_SYNCH or hold any higher ranked spin locks. MMG_STD\$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_STD\$UNLOCK. At the completion of a direct-I/O transfer, IOC_STD\$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_STD\$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_STD\$READLOCK, EXE_STD\$WRITELOCK, and EXE_STD\$MODIFYLOCK) allow a driver to specify an error-handling callback routine that can call MMG_STD\$UNLOCK to unlock all previously locked regions and deallocate the IRPE using EXE_STD\$DEANONPAGED.

Macro

CALL_UNLOCK

In a Step 2 driver, CALL_UNLOCK simulates a JSB to MMG\$UNLOCK. It calls MMG_STD\$UNLOCK using the current contents of R1 and R3 as the **npages** and **svapte** arguments, respectively.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- MMG_STD\$UNLOCK replaces MMG\$UNLOCK (used by Step 1 drivers and OpenVMS VAX drivers).

System Routines

MT_STD\$CHECK_ACCESS

MT_STD\$CHECK_ACCESS

Checks access rights for magtape control write functions.

Module

MTFDT

Format

status = MT_STD\$CHECK_ACCESS (irp, pcb, ucb, ccb)

Arguments

Argument	Type	Access	Mechanism	Status
irp	IRP	input	reference	required
pcb	PCB	input	reference	required
ucb	UCB	input	reference	required
ccb	CCB	input	reference	required

irp
I/O request packet.

pcb
Process control block of the current process.

ucb
Unit control block of the device assigned to the process-I/O channel specified as an argument to the \$QIO request.

ccb
Channel control block that describes the process-I/O channel.

Return Values

SS\$FDT_COMPL Warning-level status indicating that FDT processing is complete. The routine that receives this status can no longer safely access the IRP.

Status in FDT_CONTEXT

SS\$ACCVIO Process does not have write access to volume.
SS\$NORMAL I/O request has been successfully queued to the driver's start-I/O routine.
SS\$NOPRIV Process has insufficient privileges to perform a control write function.
SS\$WRITLCK Device software is write locked.

Context

FDT dispatching code in the \$QIO system service calls MT_STD\$CHECK_ACCESS as an upper-level FDT action routine at IPL\$ASTDEL.

Description

For Digital internal use only.

Macro

None.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The FDT routine MT\$CHECK_ACCESS (used by OpenVMS VAX and Step 1 device drivers) expects as input in R7 the number of the bit that specifies the code for the requested I/O function from R7, and, in R8, the address of the entry in the function decision table from which it received control.
R0, R7, and R8 are not provided as input to MT_STD\$CHECK_ACCESS.

- Upon a successful return from MT\$CHECK_ACCESS, its OpenVMS VAX and Step 1 callers needed to call EXE\$ZEROPARM to queue the request to the driver's start-I/O routine.

If the volume is not write-locked and the requesting process has write access to the volume. MT_STD\$CHECK_ACCESS automatically invokes the CALL_QIODRVPKT macro.

- MT\$CHECK_ACCESS returns control to the system service dispatcher, passing it the final \$QIO system service status in R0. MT_STD\$CHECK_ACCESS returns to its caller, passing it SS\$_FDT_COMPL status in R0 and storing the final \$QIO system service status in the FDT_CONTEXT structure. The \$QIO system service retrieves the status from this structure.

SCH_STD\$IOLOCKR

Locks the I/O database mutex on behalf of its caller for read access.

Module

MUTEX

Format

pointer = SCH_STD\$IOLOCKR (pcb)

Arguments

Argument	Type	Access	Mechanism	Status
pcb	PCB	input	reference	required

pcb
Process control block.

Return Values

pointer Address of I/O database mutex.

Context

SCH_STD\$IOLOCKR must be called at or below IPL\$_SYNCH. It returns to its caller at IPL\$_ASTDEL.

Description

For Digital internal use only.

Macro

CALL_IOLOCKR [save_r1]

where:

save_r1 indicates that the macro should preserve register R1 across the call to SCH_STD\$IOLOCKR. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

In a Step 2 driver, CALL_IOLOCKR simulates a JSB to SCH\$IOLOCKR. It calls SCH_STD\$IOLOCKR using the current contents of R4 as the **pcb** argument.

CALL_IOLOCKR returns the address of the I/O database mutex in R0. Unless you specify **save_r1=NO**, the macro preserves R1 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- SCH_STD\$IOLOCKR replaces SCH\$IOLOCKR (used by Step 1 drivers and OpenVMS VAX drivers). Unlike SCH\$IOLOCKR, SCH_STD\$IOLOCKR destroys the contents of R1 through R3 across the call.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- SCH_STD\$IOLOCKW replaces SCH\$IOLOCKW (used by Step 1 drivers and OpenVMS VAX drivers). Unlike SCH\$IOLOCKW, SCH_STD\$IOLOCKW destroys the contents of R1 through R3 across the call.

SCH_STD\$IOUNLOCK

Releases ownership of the I/O database mutex and, if the mutex has thus become available to waiting processes, reactivates the next eligible process.

Module

MUTEX

Format

SCH_STD\$IOUNLOCK (pcb)

Arguments

Argument	Type	Access	Mechanism	Status
pcb	PCB	input	reference	required

pcb
Process control block.

Context

SCH_STD\$IOUNLOCK must be called below IPL\$_SCHED. It returns to its caller at its caller's IPL.

Description

For Digital internal use only.

Macro

CALL_IOUNLOCK

In a Step 2 driver, CALL_IOUNLOCK simulates a JSB to SCH\$IOUNLOCK. It calls SCH_STD\$IOUNLOCK using the current contents of R4 as the **pcb** argument.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- SCH_STD\$IOUNLOCK replaces SCH\$IOUNLOCK (used by Step 1 drivers and OpenVMS VAX drivers). Unlike SCH\$IOUNLOCK, SCH_STD\$IOUNLOCK destroys the contents of R1 through R3 across the call.

Data Structures

Because a driver and the operating system cooperate to process an I/O request, they must have a common and current source of information about the request and the components of the I/O subsystem involved in servicing the request. This information source consists of a set of data structures collectively known as the **I/O database**.

Components of the I/O database include the following:

- Structures that describe individual hardware components, such as devices, controllers, adapters, and widgets. In this category are the following:

Structure	Description	Associated Structures
Unit control block (UCB)	Records the current status of an I/O device unit attached to the OpenVMS system	Object rights block (ORB), Controller register access mailbox (CRAM), Fork block (FKB)
Device data block (DDB)	Describes the common characteristics of devices of the same type connected to a particular controller	—
Channel request block (CRB)	Describes the current state of an I/O controller	Interrupt transfer vector block (VEC), Fork block (FKB)
Interrupt dispatch block (IDB)	Provides information that supplements that contained in the CRB, enabling the system to correctly dispatch and service interrupts from a device unit attached to a controller	Vector list extension (VLE), Controller register access mailbox (CRAM)
Adapter control block (ADP)	Describes the processor-memory interconnect (PMI), a tightly coupled I/O interconnect, or a multichannel I/O widget	Adapter bus array (BUSARRAY)

- Driver tables that allow the system to load drivers, validate device functions, and call driver routines at their entry points. In this category are the following:

Data Structures

Structure	Description	Associated Structures
Driver prologue table (DPT)	Contains information that allows the driver-loading procedure to load the driver into memory and initialize the I/O database according to the number and type of devices supported by the driver	—
Driver dispatch table (DDT)	Contains procedure values representing all external driver entry points (with the exception of the interrupt service routine) and the address of the driver's function decision table (FDT)	—
Function decision table (FDT)	Identifies those I/O functions supported by a device and associates valid function codes with the addresses of I/O preprocessing routines (also known as FDT routines)	—

- Structures that describe the context of a request for I/O activity. In this category are the following:

Structure	Description	Associated Structures
Channel control block (CCB)	Describes the software I/O channel that links a process to the target device of an I/O operation	—
I/O request packet (IRP)	Describes a pending or in-progress I/O request	I/O request packet extension (IRPE)

- Miscellaneous structures, such as the following:

Structure	Description	Associated Structures
Kernel process block (KPB)	Describes the scheduling and suspension mechanisms associated with a kernel process and records its suspended context	Fork block (FKB)
Counted resource allocation block (CRAB)	Records the number and type of a counted shared resource, such as a set of map registers, available to drivers	Counted resource context block (CRCTX)
Controller register access mailbox (CRAM)	Describes a read or write transaction to device interface register space	—

Figure 3–1 shows the relationships among the principal data structures in the I/O database.

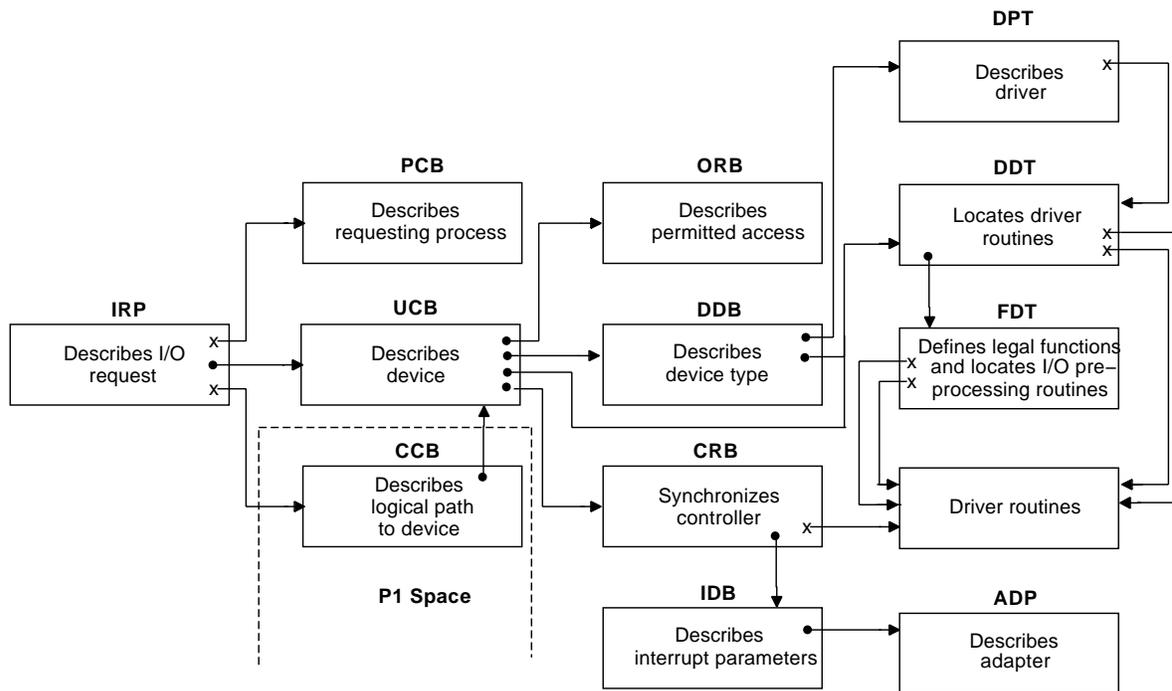
This chapter describes those structures referenced by driver code. It lists their fields in the order in which they appear in the structures. All data structures discussed in this chapter, with the exception of the channel control block (CCB), exist in nonpaged system memory.

Notes

Fields marked “Reserved” or “Unused” are reserved 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 OpenVMS operating system. The figures in this chapter list OpenVMS AXP Version 6.1 numeric offsets to aid in driver debugging.

Figure 3–1 I/O Database



3.1 ADP (Adapter Control Block)

An adapter control block (ADP) represents a hardware block that connects one interconnect to another. OpenVMS AXP I/O configuration code creates an ADP for the processor-memory interconnect (PMI), each tightly coupled I/O interconnect, and each multichannel I/O widget.

Data Structures

3.1 ADP (Adapter Control Block)

The system ADP represents the PMI. Any other ADP represents either a tightly coupled I/O interconnect or a multichannel I/O widget.

- An ADP for a tightly coupled I/O interconnect contains information related to hardware mailbox support, system topology, adapter interrupts, and related items. It also contains information about the I/O adapter that connects the interconnect to the PMI or to a parent tightly coupled I/O interconnect. The adjective **parent** in this context describes the tightly coupled I/O interconnect that is closer to the PMI.
- Although information relating to an I/O widget is normally maintained only in a widget-specific data structure defined and used by the widget's driver, information that is common to all loosely coupled I/O interconnects that connect to a multichannel I/O widget is maintained in an ADP.

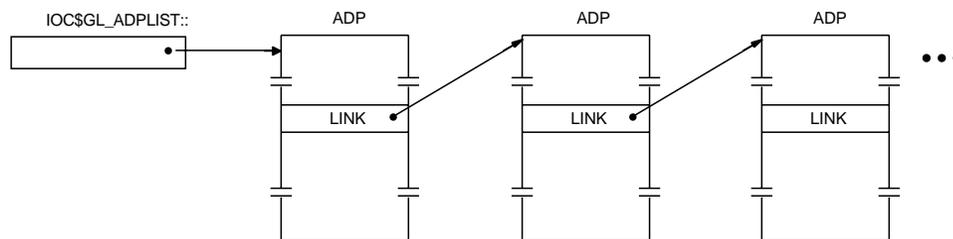
Table 3-1 defines the fields that appear in an ADP. Bus-specific extensions start at offset ADP\$*L*_XBIA_CSR in the ADP.

An ADP can have up to four auxiliary data structures:

- An adapter bus array (BUSARRAY), pointed to by ADP\$PS_BUS_ARRAY
- An adapter command table (CMDTABLE), pointed to by ADP\$PS_COMMAND_TBL
- A counted resource allocation block (CRAB), pointed to by ADP\$L_CRAB
- An indirect interrupt vector dispatch table, pointed to by ADP\$L_VECTOR

IOC\$GL_ADPLIST is the listhead for the list of all ADPs in the system. The first ADP in the ADP list (as depicted in Figure 3-2) is the system ADP. Offset ADP\$L_LINK in each ADP points to the next ADP in this list. The last ADP in the list contains a zero in this field. The SYSMAN command IO SHOW ADAPTER traverses this list and displays its contents.

Figure 3-2 ADP List



The hierarchy of tightly coupled I/O interconnects in a system is represented by the interconnection between the ADPs in the ADP list. In conjunction with the auxiliary BUSARRAY structure of each ADP, this information represents a system's configuration. Figure 3-3 shows how these structures represent the configuration of a DEC 4000 AXP system.

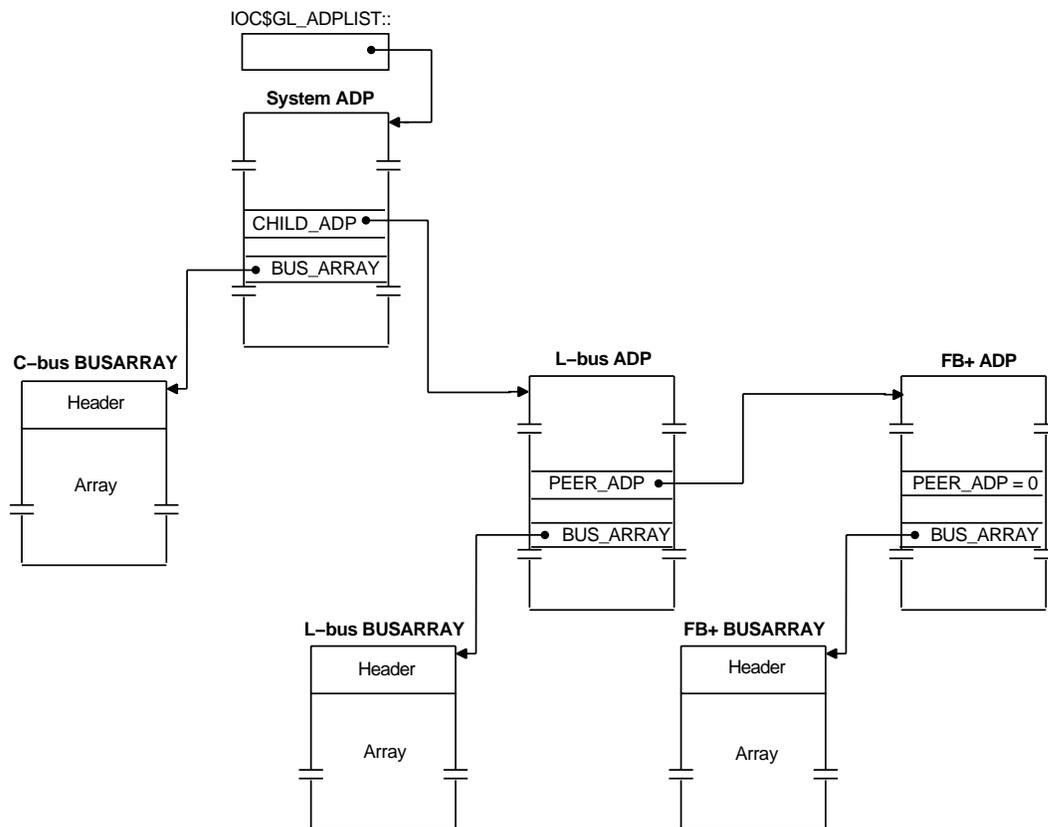
At the root of the hierarchical ADP list is the system ADP. Offset ADP\$PS_CHILD_ADP in the system ADP points to an ADP for a tightly coupled I/O interconnect at the next level in the hierarchy — one that connects to the PMI directly: that is, without other intervening interconnects. In Figure 3-3, the system ADP's ADP\$PS_CHILD_ADP field points to the L-bus ADP.

Data Structures

3.1 ADP (Adapter Control Block)

Offset `ADP$PS_PEER_ADP` in the system ADP is always zero because the system ADP has no peers. The DEC 4000 AXP system local bus (L-bus) and Futurebus+ are both tightly coupled I/O interconnects that are directly connected to the C-bus through the DEC 4000 AXP system I/O module. Offset `ADP$PS_PEER_ADP` in the L-bus ADP points to the Futurebus+ ADP, because the Futurebus+ is the L-bus's peer, connecting to the system at the same level as the L-bus. `ADP$PS_CHILD_ADP` in each of the L-bus and Futurebus+ ADPs contains a zero.

Figure 3-3 ADP Hierarchy and System Configuration



Data Structures

3.1 ADP (Adapter Control Block)

Table 3–1 Contents of Adapter Control Block

Field	Use
ADP\$Q_CSR	<p>Address of adapter control and status register (CSR), which marks the base of adapter register space on the remote tightly coupled I/O interconnect. This may be either a virtual or physical address, depending upon the adapter.</p> <p>The OpenVMS adapter initialization routine writes this field. The IOC\$CRAM_CMD uses the CSR address in calculations that set up driver transactions to and from remote adapter I/O space by means of hardware I/O mailboxes.</p> <p>For single-channel adapters, the contents of ADP\$Q_CSR and IDB\$Q_CSR are often the same. For multichannel adapters, ADP\$Q_CSR contains the base address of the common adapter register space, and individual IDBs point to the specific adapter registers associated with individual channels.</p>
ADP\$W_SIZE	<p>Size of ADP in bytes. 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. The OpenVMS adapter initialization routine writes this field when the routine creates the ADP.</p>
ADP\$B_TYPE	<p>Type of data structure. The OpenVMS adapter initialization routine writes the symbolic constant DYN\$C_ADP into this field when the routine creates the ADP.</p>
ADP\$B_NUMBER	<p>Number of this type of adapter. This field is currently unused in OpenVMS AXP systems.</p>
ADP\$L_LINK	<p>Pointer to the next ADP in the ADP list (headed by IOC\$GL_ADPLIST). The last ADP in the list contains a zero in this field.</p>
ADP\$L_TR	<p>Nexus number of adapter. The OpenVMS adapter initialization routine assigns a nexus number to each node it encounters as it probes an I/O interconnect.</p> <p>When processing an SYSMAN IO CONNECT command which specifies the /ADAPTER qualifier the driver-loading procedure compares the specified nexus number with this field of each ADP in the system to locate the adapter to which the device serviced by the driver is attached.</p>
ADP\$L_ADPTYPE	<p>Type of ADP. The OpenVMS adapter initialization routine writes a symbolic constant (defined by the SDCDEF macro in SYSS\$LIBRARY:STARLET.MLB) into this field when the routine creates an ADP.</p>

(continued on next page)

Data Structures

3.1 ADP (Adapter Control Block)

Table 3–1 (Cont.) Contents of Adapter Control Block

Field	Use
ADP\$SL_VECTOR	<p>Address of indirect interrupt vector dispatch table. For adapters that service indirect interrupts, the OpenVMS adapter initialization routine sets ADP\$V_INDIRECT_VECTOR in ADP\$SL_ADAPTER_FLAGS, and allocates sufficient nonpaged dynamic memory for this table. Each entry in this table consists of a longword pointer to the VEC substructure of a CRB of a device for which the system dispatches interrupts through this ADP.</p> <p>ADPs that service indirectly-vectorized device interrupts include a VEC substructure at ADP\$SL_INTD (as described in Section 3.5) that contains the code address (VEC\$PS_ISR_CODE), procedure descriptor address (VEC\$PS_ISR_PD), and parameter field (VEC\$SL_IDB, which contains the address of the ADP) of the adapter's indirect interrupt service routine. The SCB entries assigned to devices that interrupt indirectly contain the code address of the common interrupt dispatcher and, as the parameter, the address of ADP\$SL_INTD. The common interrupt dispatcher issues a standard call to the ADP's indirect interrupt service routine, which determines the interrupt vector of the interrupting device, using it as an index into the indirect interrupt vector dispatch table. The ADP's indirect interrupt service routine thereby locates the appropriate device driver's interrupt service routine and calls it, passing it the address of the IDB as the only parameter.</p>
ADP\$SL_CRB	<p>Address of controller request block (CRB) associated with the ADP. In the case of an ADP that describes a multichannel I/O widget, this field represents the head of a singly-linked list of CRBs linked together by the field CRB\$PS_CRB_LINK.</p>
ADP\$PS_MBPR	<p>Virtual address of mailbox pointer register (MBPR). The OpenVMS adapter initialization routine initializes this field.</p>
ADP\$SQ_QUEUE_TIME	<p>Timeout value for mailbox queuing operation. The OpenVMS adapter initialization routine initializes this field with the number of nanoseconds it takes to write the physical address of a hardware I/O mailbox to the MBPR without a timeout occurring.</p>
ADP\$SQ_WAIT_TIME	<p>Timeout value for the completion of a hardware I/O mailbox transaction. The OpenVMS adapter initialization routine initializes this field with the number of nanoseconds a thread should wait, before timing out, for the hardware I/O mailbox DON bit to be set.</p>
ADP\$PS_PARENT_ADP	<p>Address of the ADP in the preceding level of the system's ADP hierarchy that is related to this ADP and its peers. In the system ADP, this field contains a zero.</p> <p>See Figure 3–3, and the discussion at the beginning of Section 3.1, for an example of parent, child, and peer ADP relationships.</p>

(continued on next page)

Data Structures

3.1 ADP (Adapter Control Block)

Table 3–1 (Cont.) Contents of Adapter Control Block

Field	Use
ADP\$PS_PEER_ADP	<p>Address of the next ADP in the list of ADPs that are children of a common parent ADP in the preceding level of the system's ADP hierarchy, and headed by field ADP\$PS_CHILD_ADP in that parent ADP. This field contains a zero if the ADP has no peers.</p> <p>See Figure 3–3, and the discussion at the beginning of Section 3.1, for an example of parent, child, and peer ADP relationships.</p>
ADP\$PS_CHILD_ADP	<p>Listhead of the ADPs that are related to this ADP in the succeeding level of the ADP hierarchy, or zero if the ADP has no children. At this lower level, the child ADPs of a common parent ADP are linked together by the contents of their ADP\$PS_PEER_ADP fields.</p> <p>See Figure 3–3, and the discussion at the beginning of Section 3.1, for an example of parent, child, and peer ADP relationships.</p>
ADP\$SL_PROBE_CMD	<p>Index into the adapter command table that EXESTEST_CSR uses to determine which command to use when probing the interconnect described by this ADP.</p>
ADP\$PS_BUS_ARRAY	<p>Address of BUSARRAY describing the nodes on the tightly coupled interconnect or the ports of a multichannel I/O widget or controller associated with this ADP.</p>
ADP\$PS_COMMAND_TBL	<p>Address of the adapter command table specific to the I/O interconnect described by this ADP. The OpenVMS adapter initialization routine constructs this table.</p> <p>IOC\$CRAM_CMD refers to this field to locate the table when it calculates the COMMAND, MASK, and RBADR fields of a hardware I/O mailbox involved in a transaction to a device interface register.</p>
ADP\$PS_SPINLOCK	<p>Address of device lock synchronizing access to the CSRs of the devices associated with this ADP. The OpenVMS adapter initialization routine allocates this device lock and places its address in this field, IDB\$PS_SPL, and CRB\$PS_DLCK.</p>
ADP\$W_PRIM_NODE_NUM	<p>Node number of the I/O adapter (or widget) on the local interconnect (for instance, the node number of the DEC 7000 AXP Model 600 system bus [PMI] to XMI bus adapter on the PMI).</p>
ADP\$W_SEC_NODE_NUM	<p>Node number of the I/O adapter on the remote interconnect (for instance, the node number of the DEC 7000 AXP Model 600 system bus [PMI] to XMI bus adapter on the XMI).</p>
ADP\$B_HOSE_NUM	<p>Hose number associated with the I/O adapter. OpenVMS adapter initialization routine writes this field.</p>
ADP\$SL_CRAB	<p>Address of CRAB used to manage map registers, if the AXP system provides map registers for this adapter.</p>

(continued on next page)

Data Structures

3.1 ADP (Adapter Control Block)

Table 3–1 (Cont.) Contents of Adapter Control Block

Field	Use
ADP\$SL_ADAPTER_FLAGS	The following bit is defined within ADP\$SL_ADAPTER_FLAGS: <div style="margin-left: 20px;"> ADP\$V_INDIRECT_VECTOR Adapter services indirectly vectored interrupts for its associated devices. ADP\$V_ONLINE Adapter is online. ADP\$V_BOOT_ADP Adapter is boot adapter. </div>
ADP\$SL_VPORTSTS	CI-VAX port status bits. The following bits are defined within ADP\$SL_VPORTSTS: <div style="margin-left: 20px;"> ADP\$V_SHUTDOWN CI-adapter microcode is stopped. ADP\$V_PORTONLY CI-port restart only—no adapter restart. ADP\$V_STRUCT_ALLOCATED CI/SCSI-adapter-wide structures allocated. </div>
ADP\$PS_NODE_FUNCTION	Procedure value of the node-specific function routine that services driver calls to IOC\$NODE_FUNCTION.
ADP\$SL_AVECTOR	Address of first SCB vector for adapter.
ADP\$Q_SCRATCH_BUF_PA	Physical address of adapter scratch buffer.
ADP\$PS_SCRATCH_BUF_VA	Virtual address of a physically contiguous scratch buffer used in an adapter-specific manner.
ADP\$SL_SCRATCH_BUF_LEN	Size of adapter scratch buffer.
ADP\$SL_LSDUMP	Address of physical contiguous memory for the adapter memory dump.
ADP\$PS_PROBE_CSR	Procedure value of adapter-specific routine that checks for the existence of devices on an I/O interconnect. EXE\$PROBE_CSR issues a standard call to this routine.
ADP\$PS_PROBE_CSR_CLEANUP	Procedure value of adapter probe CSR cleanup routine. The adapter-specific probe CSR routine calls the cleanup routine when an error occurs during its attempts to probe an I/O interconnect.
ADP\$PS_LOAD_MAP_REG	Procedure value of adapter load map register routine.
ADP\$PS_SHUTDOWN	Procedure value of adapter shutdown routine.
ADP\$PS_CONFIG_TABLE	Pointer to autoconfiguration table.
ADP\$PS_MAP_REG_BASE	Base virtual address of adapter map registers.
ADP\$PS_ADP_SPECIFIC	Address of adapter auxiliary data structure.
ADP\$PS_DISABLE_INTERRUPTS	Address of adapter-specific interrupt disabling routine.
ADP\$PS_STARTUP	Address of adapter-specific startup routine.
ADP\$PS_INIT	Address of adapter-specific initialization routine.
ADP\$Q_HARDWARE_TYPE	Saved hardware device type information. The interpretation of this information is adapter-specific.

(continued on next page)

Data Structures

3.1 ADP (Adapter Control Block)

Table 3–1 (Cont.) Contents of Adapter Control Block

Field	Use
ADPSQ_HARDWARE_REV	Saved hardware device revision information. The interpretation of this information is adapter-specific.
ADPSL_INTD	Interrupt transfer vector. For adapters that service indirect interrupts (ADPSV_INDIRECT_VECTOR in ADPSL_ADAPTER_FLAGS is set), this 4-longword field (described in Section 3.5) provides information used by OpenVMS AXP to service a device interrupt, such as the location of the ADP and its indirect interrupt service routine. See the description of the ADPSL_VECTOR field for additional information on how the adapter services indirect interrupts.

3.1.1 BUSARRAY (Bus Array)

The bus array (BUSARRAY) contains information about the nodes on a tightly coupled I/O interconnect or the ports of a multichannel I/O widget. The BUSARRAY consists of a fixed portion and an array of entries. The fixed portion records the interconnect type, the number of nodes on the interconnect, and a pointer to the ADP with which the BUSARRAY is associated. Each array entry records the node number, the node's hardware ID, and a pointer to either an ADP or a CRB.

Table 3–2 describes the fields of the BUSARRAY structure; Table 3–3 describes the contents of each entry in the bus array.

Table 3–2 Contents of Bus Array

Field	Use
BUSARRAY\$PS_PARENT_ADP	Address of ADP for the tightly coupled I/O interconnect or multichannel I/O widget the BUSARRAY describes.
BUSARRAY\$W_SIZE	Size of BUSARRAY in bytes. The adapter initialization routine writes this field when it creates the BUSARRAY.
BUSARRAY\$B_TYPE	Type of data structure. The adapter initialization routine writes the symbolic constant DYN\$C_MISC in this field when it creates the BUSARRAY.
BUSARRAY\$B_SUBTYPE	Structure subtype. The adapter initialization routine writes DYN\$C_BUSARRAY in this field when it creates the BUSARRAY.
BUSARRAY\$L_BUS_TYPE	Type of tightly coupled I/O interconnect or multichannel I/O widget the BUSARRAY describes. The adapter initialization routine writes this field when it creates the BUSARRAY. The following constants (defined by the \$BUSDEF macro in SYSSLIBRARY:LIB.MLB) represent the interconnects supported on OpenVMS AXP systems: BUS\$_FBUS Futurebus BUS\$_XMI XMI

(continued on next page)

Data Structures

3.1 ADP (Adapter Control Block)

Table 3–2 (Cont.) Contents of Bus Array

Field	Use
	BUS\$_LBUS DEC 4000 AXP LBUS
	BUS\$_TURBO TURBOchannel
	BUS\$_CBUS DEC 4000 AXP system bus
	BUS\$_LSB DEC 7000 AXP Model 600 system bus
	BUS\$_SCSI SCSI
	BUS\$_NI Ethernet
	BUS\$_CI CI
	BUS\$_KA0402_ DEC 3000 AXP Model 500 core CORE_IO I/O bus
	BUS\$_KDM70 KDM70
	BUS\$_GENXMI Generic XMI
	BUS\$_BUSLESS_ No bus SYSTEM
BUSARRAY\$\$_BUS_NODE_ CNT	Number of entries in the bus array located at BUSARRAY\$\$_ENTRY_LIST. The OpenVMS adapter initialization routine writes this field when it creates the BUSARRAY.
BUSARRAY\$\$_ENTRY_LIST	Bus array consisting of BUSARRAY\$\$_BUS_NODE_ CNT entries.

Table 3–3 Contents of Bus Array

Field	Use
BUSARRAY\$\$_HW_ID	Hardware ID. The macro \$NDTDEF (in SYSSLIBRARY:LIB.MLB) provides symbolic definitions for the hardware IDs of all possible OpenVMS AXP nodes.
BUSARRAY\$\$_CSR	Base address of the node's CSR. The adapter initialization routine writes this field.
BUSARRAY\$\$_NODE_ NUMBER	Node number. The adapter initialization routine writes this field.
BUSARRAY\$\$_FLAGS	Bus array flags. The only bit that is currently defined, BUSARRAY\$\$_NO_RECONNECT, when set, indicates that a node has been configured properly. A bus- specific routine in an IOGEN configuration building module (ICBM) sets this bit.
BUSARRAY\$\$_PS_CRB	Pointer to node's CRB. This field must be zero if BUSARRAY\$\$_PS_ADP is filled in.
BUSARRAY\$\$_PS_ADP	Pointer to the child ADP of the parent ADP (identified by BUSARRAY\$\$_PS_PARENT_ADP) with which this node is associated. If there is no such child ADP, this field must be zero.
BUSARRAY\$\$_AUTOCONFIG	Reserved for the Autoconfiguration facility.

(continued on next page)

Data Structures

3.1 ADP (Adapter Control Block)

Table 3–3 (Cont.) Contents of Bus Array

Field	Use
BUSARRAY\$\$_CTRLTR	A bus-specific routine in an IOGEN configuration building modules writes this field by calling IOGEN\$ASSIGN_CONTROLLER.

3.2 CCB (Channel Control Block)

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 channel control blocks (CCBs) preallocated to the process. 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 described in Table 3–4.

Table 3–4 Contents of Channel Control Block

Field	Use												
CCB\$\$_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\$\$_WIND	Address of window control block (WCB) for file-structured device assignment. This field is written by an ancillary control process (ACP) or the extended QIO processor (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\$\$_STS	Channel status. The following bits are defined within CCB\$\$_STS: <table border="0"><tr><td>CCB\$\$_AMB</td><td>Mailbox associated with channel.</td></tr><tr><td>CCB\$\$_IMGTMP</td><td>Temporary image.</td></tr><tr><td>CCB\$\$_RDCHKDON</td><td>Read protection check completed.</td></tr><tr><td>CCB\$\$_WRTCHKDON</td><td>Write protection check completed.</td></tr><tr><td>CCB\$\$_LOGCHKDON</td><td>Logical I/O access check done.</td></tr><tr><td>CCB\$\$_PHYCHKDON</td><td>Physical I/O access check done.</td></tr></table>	CCB\$\$_AMB	Mailbox associated with channel.	CCB\$\$_IMGTMP	Temporary image.	CCB\$\$_RDCHKDON	Read protection check completed.	CCB\$\$_WRTCHKDON	Write protection check completed.	CCB\$\$_LOGCHKDON	Logical I/O access check done.	CCB\$\$_PHYCHKDON	Physical I/O access check done.
CCB\$\$_AMB	Mailbox associated with channel.												
CCB\$\$_IMGTMP	Temporary image.												
CCB\$\$_RDCHKDON	Read protection check completed.												
CCB\$\$_WRTCHKDON	Write protection check completed.												
CCB\$\$_LOGCHKDON	Logical I/O access check done.												
CCB\$\$_PHYCHKDON	Physical I/O access check done.												
CCB\$\$_AMOD	Access mode plus 1 of the channel. EXE\$ASSIGN writes the access mode value into this field.												

(continued on next page)

Table 3–4 (Cont.) Contents of Channel Control Block

Field	Use
CCBSL_IOC	Number of outstanding I/O requests on channel. EXESQIO 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 EXESDASSGN read this field.
CCBSL_DIRP	Address of I/O request packet (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, EXESQIO holds the deaccess request until all other outstanding I/O requests are processed.
CCBSL_CHAN	Associated channel number.

3.3 CRAM (Controller Register Access Mailbox)

The controller register access mailbox (CRAM) contains information that describes a specific hardware I/O mailbox transaction. To facilitate mailbox operations within the operating system, the CRAM contains information required by the operating system as well as the hardware I/O mailbox itself. For example, mailbox operations require the physical address of the hardware mailbox itself as well as the virtual address of the corresponding mailbox pointer register (MBPR). Additionally, the timeout values for both the queuing and waiting portions of a mailbox operation are kept in the CRAM.

CRAMs are allocated from pages obtained from the memory management free list. Once the pages have been allocated from the free list, they are managed privately by the CRAM allocation and deallocation code. Each page of CRAMs begins with a structure known as a controller register access mailbox header (CRAMH); the set of pages is maintained as a linked list starting at IOC\$GQ_CRAMH_HDR.

The controller register access mailbox is described in Table 3–5.

Data Structures

3.3 CRAM (Controller Register Access Mailbox)

Table 3–5 Contents of Controller Register Access Mailbox

Field	Use
CRAM\$L_FLINK	Forward link to next CRAM in list (headed by IDB\$PS_CRAM or UCB\$PS_CRAM). The driver-loading procedure initializes this field when the driver preallocates CRAMs by specifying the idb_cramps or ucb_cramps argument to the DPTAB macro. The contents of this field are unpredictable and must be managed by the driver when it spontaneously allocates CRAMs.
CRAM\$L_BLINK	Backward link to next CRAM in list (headed by IDB\$PS_CRAM or UCB\$PS_CRAM). The driver-loading procedure initializes this field when the driver preallocates CRAMs by specifying the idb_cramps or ucb_cramps argument to the DPTAB macro. The contents of this field are unpredictable and must be managed by the driver when it spontaneously allocates CRAMs.
CRAM\$W_SIZE	Size of CRAM in bytes. IOC\$ALLOCATE_CRAM writes the symbolic constant CRAM\$K_LENGTH in this field when it initializes the CRAM.
CRAM\$B_TYPE	Structure type. IOC\$ALLOCATE_CRAM initializes this field to DYN\$C_MISC.
CRAM\$B_SUBTYPE	Structure subtype. IOC\$ALLOCATE_CRAM initializes this field to DYN\$C_CRAM.
CRAM\$L_MBPR	Virtual address of mailbox pointer register (MBPR). When IOC\$ALLOCATE_CRAM is called by the driver-loading procedure, or when it is called independently with the idb parameter, it initializes this field from the contents of ADP\$PS_MBPR. Otherwise, it places a zero in this field.
CRAM\$Q_HW_MBX	Physical address of hardware mailbox. IOC\$ALLOCATE_CRAM initializes this field.
CRAM\$Q_QUEUE_TIME	MBPR queue timeout interval in nanoseconds. If IOC\$CRAM_QUEUE or IOC\$CRAM_CMD cannot queue the hardware I/O mailbox defined in this CRAM to the MBPR in this amount of time, it returns SSS_INTERLOCK status to its caller. When IOC\$ALLOCATE_CRAM is called by the driver-loading procedure, or when it is called independently with the idb parameter, it initializes this field from the contents of ADP\$Q_QUEUE_TIME. Otherwise, it places a zero in this field.
CRAM\$Q_WAIT_TIME	Mailbox transaction wait timeout interval in nanoseconds. If IOC\$CRAM_IO or IOC\$CRAM_WAIT does not see the done or error bit set in the hardware mailbox in this interval, it returns SSS_TIMEOUT status to its caller. When IOC\$ALLOCATE_CRAM is called by the driver-loading procedure, or when it is called independently with the idb parameter, it initializes this field from the contents of ADP\$Q_WAIT_TIME. Otherwise, it places a zero in this field.
CRAM\$L_DRIVER	Spare longword for driver use.

(continued on next page)

3.3 CRAM (Controller Register Access Mailbox)

Table 3–5 (Cont.) Contents of Controller Register Access Mailbox

Field	Use				
CRAM\$SL_IDB	Pointer to IDB. IOC\$ALLOCATE_CRAM initializes this field when called from the driver-loading procedure, and when called with a nonzero idb parameter. Otherwise, it places a zero in this field.				
CRAM\$SL_UCB	Pointer to UCB. IOC\$ALLOCATE_CRAM initializes this field when called from the driver-loading procedure (if the ucb_cram argument is supplied to the DPTAB macro), and when called with a nonzero ucb parameter. Otherwise, it places a zero in this field.				
CRAM\$SL_CRAM_FLAGS	The following bits are defined within CRAM\$SL_CRAM_FLAGS: <table border="0"> <tr> <td style="padding-right: 20px;">CRAM\$V_CRAM_IN_USE</td> <td>CRAM is valid. IOC\$CRAM_QUEUE and IOC\$CRAM_IO set this bit when they have successfully posted the hardware I/O mailbox portion of the CRAM to the MBPR. IOC\$CRAM_IO and IOC\$CRAM_WAIT clear this bit when the mailbox transaction is completed (either successfully or unsuccessfully) within the mailbox transaction timeout interval (CRAM\$Q_WAIT_TIME).</td> </tr> <tr> <td>CRAM\$V_DER</td> <td>Disable error reporting.</td> </tr> </table>	CRAM\$V_CRAM_IN_USE	CRAM is valid. IOC\$CRAM_QUEUE and IOC\$CRAM_IO set this bit when they have successfully posted the hardware I/O mailbox portion of the CRAM to the MBPR. IOC\$CRAM_IO and IOC\$CRAM_WAIT clear this bit when the mailbox transaction is completed (either successfully or unsuccessfully) within the mailbox transaction timeout interval (CRAM\$Q_WAIT_TIME).	CRAM\$V_DER	Disable error reporting.
CRAM\$V_CRAM_IN_USE	CRAM is valid. IOC\$CRAM_QUEUE and IOC\$CRAM_IO set this bit when they have successfully posted the hardware I/O mailbox portion of the CRAM to the MBPR. IOC\$CRAM_IO and IOC\$CRAM_WAIT clear this bit when the mailbox transaction is completed (either successfully or unsuccessfully) within the mailbox transaction timeout interval (CRAM\$Q_WAIT_TIME).				
CRAM\$V_DER	Disable error reporting.				
CRAM\$SL_COMMAND	Command to the remote I/O interconnect command specifying a read or write transaction. The local I/O adapter delivers this command to the remote interconnect to which the target widget is connected. The command may also include fields such as address only, address width, and data width. This field, aligned on a 64-byte boundary, indicates the beginning of the hardware I/O mailbox structure in this CRAM. The characters "MBZ" (must be zero) indicate that the field must contain a zero when it is supplied in a CRAM operation. Given a command index, IOC\$CRAM_CMD initializes this field in a manner specific to the I/O interconnect that is to be the target of an operation using this CRAM.				
CRAM\$B_BYTE_MASK	Byte mask that indicates which bytes within the remote bus address (CRAM\$Q_RBADR) are to be written for mailbox write operations. IOC\$CRAM_CMD, on behalf of a device driver, writes the size of the target location (byte, word, longword, or quadword) in this field. Given a byte offset to an address in remote I/O space, IOC\$CRAM_CMD initializes this field in a manner specific to the masking mode of the I/O interconnect that is to be the target of an operation using this CRAM.				

(continued on next page)

Data Structures

3.3 CRAM (Controller Register Access Mailbox)

Table 3–5 (Cont.) Contents of Controller Register Access Mailbox

Field	Use
CRAM\$B_HOSE	<p>I/O bus number, or hose. This field specifies the remote I/O interconnect to be accessed by the mailbox transaction described by this CRAM.</p> <p>When IOCSALLOCATE_CRAM is called by the driver-loading procedure, or when it is called independently with the idb parameter, it initializes this field from the contents of ADPSB_HOSE_NUM. Otherwise, it places a zero in this field.</p>
CRAM\$Q_RBADR	<p>Remote bus address. A device driver calls IOC\$CRAM_CMD to write a value in this field that represents the physical address of the device interface register to be accessed. IOC\$CRAM_CMD calculates this value from IDB\$Q_CSR (or ADP\$Q_CSR if IDB\$Q_CSR is not available) and the byte_offset input argument.</p>
CRAM\$Q_WDATA	<p>Data to be written. If CRAM\$SL_COMMAND indicates a write transaction to the remote interconnect, the driver initializes this field with the data to be written to the target device interface register. If CRAM\$SL_COMMAND indicates a read transaction, this field is not used.</p>
CRAM\$Q_RDATA	<p>Returned read data. If CRAM\$SL_COMMAND indicates a read transaction to the remote interconnect, the remote adapter returns the requested data in this field. If CRAM\$SL_COMMAND indicates a write transaction, the contents of this field are unpredictable.</p>
CRAM\$W_MBX_FLAGS	<p>The following bits are defined within CRAM\$W_MBX_FLAGS:</p> <p>CRAM\$V_MBX_DONE Mailbox operation completed. IOC\$CRAM_WAIT and IOC\$CRAM_IO check this bit to determine the completion of a hardware I/O mailbox transaction. For both read and write commands, this bit, when set, indicates that the CRAM\$V_MBX_ERROR, CRAM\$W_ERROR_BITS, and CRAM\$Q_RDATA fields are valid. The mailbox structure may then be safely modified by software (reused). Note that the setting of the DON bit does not guarantee that a remote I/O space write has actually completed at the bridge.</p>

(continued on next page)

3.3 CRAM (Controller Register Access Mailbox)

Table 3–5 (Cont.) Contents of Controller Register Access Mailbox

Field	Use
CRAMSV_MBX_ERROR	Error in operation. IOC\$CRAM_WAIT and IOC\$CRAM_IO check this bit to determine whether an error occurred during a hardware I/O mailbox transaction. If set on a read command, indicates that an error was encountered and that the CRAM\$W_ERROR_BITS field contains additional information. This bit is valid only when CRAMSV_MBX_DONE is set.
CRAM\$W_ERROR_BITS	Device-specific error bits that indicate the completion status of a mailbox transaction described by this CRAM.

3.4 CRB (Channel Request Block)

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 driver fork processes waiting to gain access to a device through the controller. It also contains one interrupt transfer vector (VEC) for each of the controller's interrupt vectors.

The channel request block is described in Table 3–6.

Table 3–6 Contents of Channel Request Block

Field	Use
CRB\$SL_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 CRAMs or nonpaged 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\$SL_FQBL	Fork queue backward link. The link points to the previous entry in the fork queue.
CRB\$W_SIZE	Size of CRB in bytes. 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.

(continued on next page)

Data Structures

3.4 CRB (Channel Request Block)

Table 3–6 (Cont.) Contents of Channel Request Block

Field	Use				
CRBSL_FPC	Procedure value of routine at which execution resumes when the fork dispatcher dequeues the fork block. EXESPRIMITIVE_FORK writes this field when called to suspend driver execution.				
CRBSQ_FR3	Value of R3 at the time that the executing code requests the operating system to create a fork block. EXESPRIMITIVE_FORK writes this field when called to suspend driver execution.				
CRBSQ_FR4	Value of R4 at the time that the executing code requests OpenVMS to create a fork block. EXESPRIMITIVE_FORK writes this field when called to suspend driver execution.				
CRBSB_TT_TYPE	Controller type.				
CRBSL_REFC	Unit control block (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.				
CRBSB_MASK	Mask that describes controller status. The following fields are defined in CRBSB_MASK:				
	<table border="0"> <tr> <td style="padding-right: 20px;">CRBSV_BSY</td> <td>Busy bit. IOCSPRIMITIVE_REQCHANy 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. IOCSRELCHAN clears the busy bit if no driver is waiting to acquire the channel.</td> </tr> <tr> <td>CRBSV_UNINIT</td> <td>Indication, when set, that the OpenVMS driver loading procedure has yet to call the driver's controller initialization routine. The driver loading procedure reads this bit to determine whether to call the controller initialization routine and clears it when the initialization routine completes.</td> </tr> </table>	CRBSV_BSY	Busy bit. IOCSPRIMITIVE_REQCHANy 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. IOCSRELCHAN clears the busy bit if no driver is waiting to acquire the channel.	CRBSV_UNINIT	Indication, when set, that the OpenVMS driver loading procedure has yet to call the driver's controller initialization routine. The driver loading procedure reads this bit to determine whether to call the controller initialization routine and clears it when the initialization routine completes.
CRBSV_BSY	Busy bit. IOCSPRIMITIVE_REQCHANy 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. IOCSRELCHAN clears the busy bit if no driver is waiting to acquire the channel.				
CRBSV_UNINIT	Indication, when set, that the OpenVMS driver loading procedure has yet to call the driver's controller initialization routine. The driver loading procedure reads this bit to determine whether to call the controller initialization routine and clears it when the initialization routine completes.				
CRBSPS_BUSARRAY	Address of BUSARRAY that describes the devices residing on loosely coupled I/O interconnects (for instance, a SCSI port).				
CRBSQ_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.				
CRBSQ_LAN_STRUC	Address of auxiliary data structure used by local area network drivers.				
CRBSQ_SSB_STRUC	Address of auxiliary data structure used by system communications services drivers.				

(continued on next page)

Table 3–6 (Cont.) Contents of Channel Request Block

Field	Use
CRB\$SL_TIMELINK	Forward link in queue of CRBs waiting for periodic wakeups. This field points to the CRB\$SL_TIMELINK field of the next CRB in the list. The CRB\$SL_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\$SL_NODE	Bus-slot number of the controller node. The OpenVMS AXP driver-loading procedure initializes this field, which is used by IOC\$NODE_FUNCTION to enable or disable functionality for the node.
CRB\$SL_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 required 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\$SL_TOUTROUT	Procedure value 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\$SL_DUETIME if another periodic wakeup request is desired. Use of this field is reserved to Digital.
CRB\$PS_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\$PS_CRB_LINK	Pointer to next CRB on ADP.
CRB\$PS_CTRLR_SHUTDOWN	Procedure value of driver controller shutdown routine.
CRB\$SL_INTD	Interrupt transfer vector. This 4-longword field (described in Section 3.5) contains information used by the operating system to service a device interrupt, such as the location of the device's interrupt service routine and its associated interrupt dispatch block (IDB).
CRB\$SL_INTD2	Second interrupt transfer vector for devices with multiple interrupt vectors.

3.5 VEC (Interrupt Transfer Vector Block)

An interrupt transfer vector block (VEC) exists in OpenVMS only as a substructure of a CRB or an ADP. A VEC stores information that allows OpenVMS to correctly dispatch and service the interrupts of devices that share a common controller or adapter. The VEC substructures of ADPs are of interest only to OpenVMS-supplied device drivers.

By default, the driver-loading procedure creates a single VEC within a given CRB. (Adapter initialization code generates the VECs associated with an ADP.) You can control the number of VECs created by specifying a value in the /NUMVEC qualifier of an SYSMAN IO CONNECT command.

Data Structures

3.5 VEC (Interrupt Transfer Vector Block)

The OpenVMS driver-loading procedure initializes the contents of each VEC's IDB and ADP pointers and connects the VEC to the appropriate vector offsets within the system control block (SCB). A device driver must initialize the VEC\$PS_ISR_CODE and VEC\$PS_ISR_PD fields in each VEC by invoking the DPT_STORE_ISR macro, as described in Chapter 4.

Although the OpenVMS interrupt dispatching mechanism passes the address of the device's IDB to a driver's interrupt service routine as its sole parameter, other driver routines must determine the location of the IDB by directly accessing VEC\$SL_IDB in a VEC substructure. The data structure definition macro \$CRBDEF supplies symbolic offsets so that a driver can easily locate the first two VECs. For additional VECs, the driver must employ the following formula, where *n* represents the vector number:

$$\text{CRB\$L_INTD} + ((n-1) * \text{VEC\$K_LENGTH})$$

The following table lists the symbolic location of the first three VECs for a given controller:

Vector Number	Symbolic Offset to VEC
1	CRB\$SL_INTD
2	CRB\$SL_INTD2
3	CRB\$SL_INTD + <2*VEC\$K_LENGTH>

Table 3–7 describes the contents of the VEC substructure.

Table 3–7 Contents of Interrupt Transfer Vector Block (VEC)

Field	Use
VEC\$PS_ISR_CODE	Address of the code entry point of a driver interrupt service routine (ISR). The driver specifies an ISR by using the DPT_STORE_ISR macro, which initializes this field.
VEC\$PS_ISR_PD	Address of the procedure descriptor of a driver ISR. The driver specifies an ISR by using the DPT_STORE_ISR macro, which initializes this field.
VEC\$SL_IDB	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 addresses of IDB CRAMs. When a driver's interrupt service routine gains control, it receives this value as its only parameter.
VEC\$PS_ADP	Address of ADP. The SYSMAN command IO CONNECT must specify the nexus number of the adapter used by a controller. The driver-loading procedure writes the address of the ADP for the specified adapter into the VEC\$PS_ADP field.

3.6 DDB (Device Data Block)

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 SYSMAN command CONNECT. The procedure initializes all fields in the DDB. All the

Data Structures

3.6 DDB (Device Data Block)

DDBs associated with a given system block (SB) are linked in a singly linked list off that SB. The field DDB\$\$_SB points to the parent SB of any given DDB.

The device data block is described in Table 3–8.

Table 3–8 Contents of Device Data Block

Field	Use								
DDB\$_LINK	Address of next DDB. A zero indicates that this is the last DDB in the DDB chain.								
DDB\$_UCB	Address of UCB for first unit attached to controller.								
DDB\$_SIZE	Size of DDB in bytes. The driver-loading procedure writes the symbolic constant DDB\$_LENGTH in this field when it creates the DDB.								
DDB\$_TYPE	Type of data structure. The driver-loading procedure writes the constant DYN\$_DDB into this field when the procedure creates the DDB.								
DDB\$_DDT	Address of driver dispatch table (DDT). OpenVMS can transfer control to a device driver only through procedure values and entry points listed in the DDT, CRB, and UCB fork block. The driver-loading procedure initializes this field.								
DDB\$_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\$_ACPCCLASS, 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\$_ACPCCLASS by invoking a DPT_STORE macro. Values for DDB\$_ACPCCLASS are as follows: <table style="margin-left: 2em; border: none;"> <tr> <td>DDB\$_PACK</td> <td>Standard disk pack</td> </tr> <tr> <td>DDB\$_CART</td> <td>Cartridge disk pack</td> </tr> <tr> <td>DDB\$_SLOW</td> <td>Floppy disk</td> </tr> <tr> <td>DDB\$_TAPE</td> <td>Magnetic tape that simulates file-structured device</td> </tr> </table>	DDB\$_PACK	Standard disk pack	DDB\$_CART	Cartridge disk pack	DDB\$_SLOW	Floppy disk	DDB\$_TAPE	Magnetic tape that simulates file-structured device
DDB\$_PACK	Standard disk pack								
DDB\$_CART	Cartridge disk pack								
DDB\$_SLOW	Floppy disk								
DDB\$_TAPE	Magnetic tape that simulates file-structured device								
DDB\$_NAME	Name of device. The first byte of this field contains the number of characters in the device name. The remainder of the field contains a string of up to 15 characters representing the device name in the format <i>ddc</i> , where <table style="margin-left: 2em; border: none;"> <tr> <td>dd =</td> <td>device code (up to 9 alphabetic characters)</td> </tr> <tr> <td>c =</td> <td>controller designation (alphabetic)</td> </tr> </table>	dd =	device code (up to 9 alphabetic characters)	c =	controller designation (alphabetic)				
dd =	device code (up to 9 alphabetic characters)								
c =	controller designation (alphabetic)								
DDB\$_DPT	Address of DPT of driver that supports this device.								
DDB\$_DRVLINK	Address of next DDB in singly linked list, headed by DPT\$_DDB_LIST, of DDBs serviced by a particular driver.								
DDB\$_SB	Address of system block.								
DDB\$_CONLINK	Address of next DDB in the connection subchain.								
DDB\$_ALLOCLS	Allocation class of device.								
DDB\$_2P_UCB	Address of the first UCB on the secondary path.								

Data Structures

3.7 DDT (Driver Dispatch Table)

3.7 DDT (Driver Dispatch Table)

Each device driver contains a driver dispatch table (DDT). The DDT lists procedure values for driver entry points that system routines call.

A device driver creates a DDT by invoking the VAX MACRO DDTAB macro. Table 3-9 describes the fields in the driver dispatch table.

Table 3-9 Contents of Driver Dispatch Table

Field	Use
DDT\$PS_START_2	<p>Procedure value of the driver's start-I/O routine. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the start argument to the macro. All drivers must specify a start-I/O routine.</p> <p>When a device unit is idle and an I/O request is pending for that unit, IOC\$INITIATE transfers control to the routine entry point represented by the procedure value in this field.</p> <p>A driver that employs kernel process services typically specifies its start-I/O routine in the kp_startio argument to the DDTAB macro, and the system routine EXE\$KP_STARTIO in the start argument. This allows OpenVMS to set up the kernel process environment prior to transferring control to the driver's start-I/O routine.</p>
DDT\$PS_START_JSB	<p>Procedure value of the driver Start I/O routine when DDTAB JSB_START is used. The DDT\$PS_START field contains a pointer to the IOC\$START_C2J routine.</p>
DDT\$IW_SIZE	<p>Size of DDT in bytes. The DDTAB macro writes the symbolic constant DDT\$SK_LENGTH in this field when creating the DDT.</p>
DDT\$W_DIAGBUF	<p>Size of diagnostic buffer, as specified 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, as specified 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\$W_FDTSIZE	<p>Unused on OpenVMS AXP systems.</p>

(continued on next page)

Data Structures

3.7 DDT (Driver Dispatch Table)

Table 3–9 (Cont.) Contents of Driver Dispatch Table

Field	Use
DDT\$PS_CTRLINIT_2	Procedure value of controller initialization routine. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the ctrlinit argument to the macro.
DDT\$PS_UNITINIT_2	Procedure value of the device's unit initialization routine. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the unitinit argument to the macro.
DDT\$PS_CLONEDUCB_2	Procedure value of cloned UCB routine. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the cloneducb argument to the macro.
DDT\$PS_FDT_2	Address of the driver's FDT. Every driver must specify this address in the functb argument to the DDTAB macro. EXESQIO refers to the FDT to validate I/O function codes, decide which functions are buffered, and call FDT routines associated with function codes.
DDT\$PS_CANCEL_2	Procedure value of the driver's cancel-I/O routine. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the cancel argument to the macro. Some devices require special cleanup processing when a process or a system 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.
DDT\$PS_REGDUMP_2	Procedure value of the driver's register dumping routine. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the regdmp argument to the macro. IOCS\$DIAGBUFILL, ERL\$DEVICERR, and ERL\$DEVICTMO call this routine to write device register contents into a diagnostic buffer or error message buffer.
DDT\$PS_ALTSTART_2	Procedure value of the driver's alternate start-I/O routine. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the altstart argument to the macro. EXESALTQUEPKT transfers control to the alternate start-I/O routine specified in this field.
DDT\$PS_ALTSTART_JSB	Procedure value of the driver Alternate Start I/O routine when DDTAB JSB_ALTSTART is used. The DDT\$PS_ALTSTART field contains a pointer to the IOCSALTSTART_C2J routine.
DDT\$PS_MNTVER_2	Procedure value of the system routine (IOCSMNTVER) called at the beginning and end of mount verification operation. The default value of the mntver argument to the DPTAB macro is the procedure value of this routine. Use of the mntver argument to specify any routine other than IOCSMNTVER is reserved to Digital.

(continued on next page)

Data Structures

3.7 DDT (Driver Dispatch Table)

Table 3–9 (Cont.) Contents of Driver Dispatch Table

Field	Use
DDT\$SL_MNTV_SSSC	<p>Procedure value of the routine that is called when mount verification is performed for a shadow-set state change. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the mntv_sssc argument to the macro.</p> <p>Use of this field is reserved to Digital.</p>
DDT\$SL_MNTV_FOR	<p>Procedure value of the routine that is called when mount verification is performed for a foreign device. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the mntv_for argument to the macro.</p> <p>Use of this field is reserved to Digital.</p>
DDT\$SL_MNTV_SQD	<p>Procedure value of the routine that is called when mount verification is performed for a sequential device. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the mntv_sqd argument to the macro.</p> <p>Use of this field is reserved to Digital.</p>
DDT\$SL_AUX_STORAGE	<p>Address of auxiliary storage area, as specified in the aux_storage argument to the DDTAB macro.</p> <p>Use of this field is reserved to Digital.</p>
DDT\$SL_AUX_ROUTINE	<p>Procedure value of auxiliary routine in the mailbox driver that is called by SYSSASSIGN. The OpenVMS VAX mailbox driver uses this routine to complete the processing of reader-wait and writer-wait set mode requests. (Auxiliary routines have yet to be implemented in OpenVMS AXP systems.) The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the aux_routine argument to the macro.</p> <p>Use of this field is reserved to Digital.</p>
DDT\$PS_CHANNEL_ASSIGN_2	<p>Procedure value of routine, called by SYSSASSIGN, to complete channel assignment in a device-specific manner. (Channel-assignment routines have yet to be implemented in OpenVMS AXP systems.) The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the channel_assign argument to the macro.</p> <p>Use of this field is reserved to Digital.</p>
DDT\$PS_CANCEL_SELECTIVE_2	<p>Procedure value of the routine that cancels a list of I/O requests from the specified channel, including both waiting and active requests. The OpenVMS VAX terminal driver and mailbox driver provide this capability which is not yet implemented in OpenVMS AXP systems. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the cancel_selective argument to the macro.</p> <p>Use of this field is reserved to Digital.</p>

(continued on next page)

Table 3–9 (Cont.) Contents of Driver Dispatch Table

Field	Use
DDT\$IS_STACK_BCNT	Size in bytes of the kernel process stack, as indicated by the kp_stack_size argument to the DDTAB macro. EXESKP_STARTIO uses this value, or KPBSK_MIN_IO_STACK (currently 8KB), whichever is larger, to determine the size of the stack created for the driver's start I/O kernel process thread.
DDT\$IS_REG_MASK	Kernel process register save mask, as indicated by the kp_reg_mask argument to the DDTAB macro. Each time a kernel process is stalled and restarted, any registers that the thread uses other than registers that the calling standard defines as scratch must be saved. EXESKP_STARTIO establishes this set of registers by merging the mask specified in this field with a register save mask (represented by the symbolic constant KPREGSK_MIN_IO_REG_MASK) that includes R2 through R5, R12 through R15, R26, R27, and R29. It then specifies the resulting mask in its call to EXESKP_START. It is this latter mask that EXESKP_START stores in KPBSIS_REG_MASK for the lifetime of the kernel process. Note that R0, R1, R16 through R25, R28, R30, and R31 are never preserved and are illegal in a register save mask. OpenVMS represents the set of these registers by the symbolic constant KPREGSK_ERR_REG_MASK. If any of these registers are indicated by the contents of DDT\$IS_REG_MASK, EXESKP_START removes them from the mask it stores in the KPB.
DDT\$PS_KP_STARTIO	Procedure value of the start-I/O routine of a driver that employs the kernel process services. The DDTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the kp_startio argument to the macro. Such a driver typically specifies the system routine EXESKP_STARTIO in the start argument to the DDTAB macro. EXESKP_STARTIO calls the start-I/O routine specified in this field after setting up the kernel process environment.

3.8 DPT (Driver Prologue Table)

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 DPTAB macro. Table 3–10 describes the driver prologue table.

Data Structures

3.8 DPT (Driver Prologue Table)

Table 3–10 Contents of Driver Prologue Table

Field	Use
DPT\$SL_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\$SL_BLINK	Backward link to previous DPT. The driver-loading procedure writes this field.
DPT\$W_SIZE	Size of DPT in bytes. The DPTAB macro writes the value <i>DPT\$K_BASE_LEN + NAM\$C_MAXRSS</i> in this field when it creates the DPT.
DPT\$B_TYPE	Type of data structure. The DPTAB macro always writes the symbolic constant <i>DYN\$C_DPT</i> into this field.
DPT\$IW_STEP	OpenVMS AXP driver step number. You must indicate that a given driver conforms to the coding practices for a Step 2 driver by supplying step=2 in the DPTAB macro invocation. Consequently, the DPTAB macro writes the symbol constant <i>DPT\$K_STEP_2</i> in this field.
DPT\$IW_STEPVER	Integer signifying the version of Step 2 interface used by this driver. An increment of this value represents a change in the interface between Step 2 drivers and the driver loading procedure that does not require changes in driver source code (for example, a change in the DPT produced by a change in the DPTAB macro). The DPTAB macro writes the symbolic constant <i>DPT\$K_STEP2_V2</i> in this field.
DPT\$W_DEFUNITS	Number of UCBs that the OpenVMS 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 <i>DPT\$PS_DELIVER</i> , this field is also the number of times that the autoconfiguration facility calls the unit delivery routine. The DPTAB macro writes the value 1 in this field by default.
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_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. OpenVMS supplies the symbolic constants described in Table 3–17 to represent UCB size. Drivers that employ their own extended UCBs use one of these constants as a base for calculating the size of their extended UCBs. The driver-loading procedure allocates blocks of nonpaged system memory of the specified size when creating UCBs for devices associated with the driver.
DPT\$IW_IDB_CRAMS	Number of CRAMS to be allocated and associated with the IDB. The driver-loading procedure allocates the number of CRAMS specified in idb_crams argument to the DPTAB macro and inserts them in the linked list headed by <i>IDB\$PS_CRAM</i> .

(continued on next page)

Data Structures

3.8 DPT (Driver Prologue Table)

Table 3–10 (Cont.) Contents of Driver Prologue Table

Field	Use
DPT\$IW_UCB_CRAMS	Number of CRAMS to be allocated and associated with the IDB. The driver-loading procedure allocates the number of CRAMS specified in ucb_crams argument to the DPTAB macro and inserts them in the linked list headed by UCB\$PS_CRAM.
DPT\$SL_FLAGS	<p>Driver-loading 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.</p> <p>The following bits are defined within DPT\$SL_FLAGS:</p> <p>DPT\$V_SUBCNTRL Device is a subcontroller.</p> <p>DPT\$V_SVP Device requires permanent system page to be allocated during driver loading.</p> <p>DPT\$V_NOUNLOAD Driver cannot be reloaded.</p> <p>DPT\$V_SCS SCS code must be loaded with this driver.</p> <p>DPT\$V_DUSHADOW Driver is the shadowing disk class driver.</p> <p>DPT\$V_SCSCI Common SCS/CI subroutines must be loaded with this driver. This bit is ignored on OpenVMS AXP systems.</p> <p>DPT\$V_BVPSUBS Common BVP subroutines must be loaded with this driver. This bit is ignored on OpenVMS AXP systems.</p> <p>DPT\$V_UCODE Driver has an associated microcode image. This bit is ignored on OpenVMS AXP systems.</p> <p>DPT\$V_SMPMOD Driver has been designed to run in an OpenVMS environment.</p> <p>DPT\$V_DECW_DECODE Driver is a DECwindows (class input) driver.</p> <p>DPT\$V_TPALLOC Select the tape allocation class parameter.</p> <p>DPT\$V_SNAPSHOT Driver is certified for system snapshot.</p> <p>DPT\$V_NO_IDB_DISPATCH Tells the driver-loading procedure not to create a list of UCB addresses at the end of the IDB (at IDB\$SL_UCBLST), regardless of the value of the maxunits argument to the DPTAB macro or the maximum units specified in the SYSMAN command IO CONNECT.</p> <p>DPT\$V_SCSI_PORT Driver is a SCSI port driver.</p>

(continued on next page)

Data Structures

3.8 DPT (Driver Prologue Table)

Table 3–10 (Cont.) Contents of Driver Prologue Table

Field	Use
DPT\$IL_ADPTYPE	Type of adapter used by the devices using this driver. The DPTAB macro uses the contents of the adapter to construct a symbolic constant of the form AT\$ adapter , the value of which it inserts in this field.
DPT\$IL_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\$PS_INIT_PD	Procedure value of the driver initialization routine. Every driver must specify a list of values to be written into data structure fields at the time that the driver-loading procedure creates the structures and loads the driver. The driver invokes the DPT_STORE macro once for each value to be written; the macro automatically generates an initialization routine containing code that performs the requested writes, and places its procedure value in this field. The driver-loading procedure calls this initialization routine prior to calling the driver's controller and unit initialization routines.
DPT\$PS_REINIT_PD	Procedure value of the driver reinitialization routine. 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 DPT_STORE macro once for each value to be written; the macro automatically generates a reinitialization routine containing code that performs the requested writes, and places its procedure value in this field. The driver-loading procedure calls the reinitialization routine at driver reloading prior to calling the driver's controller and unit initialization routines. Note that driver reloading is not yet supported on OpenVMS AXP systems.
DPT\$PS_DELIVER_2	Procedure value of the unit delivery routine that the OpenVMS autoconfiguration facility calls once for each of the number of UCBs specified in DPT\$W_DEFUNITS. The DPTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the deliver argument to the macro.
DPT\$PS_UNLOAD	Procedure value of the driver routine to be called when driver is reloaded. The DPTAB macro inserts a procedure value in this field when the driver specifies the routine's address in the unload argument to the macro. The driver-loading procedure calls the driver unloading routine before reinitializing all device units associated with the driver. Note that driver reloading is not yet supported on OpenVMS AXP systems.
DPT\$PS_DDT	Address of DDT.
DPT\$PS_DDB_LIST	Header of singly-linked list of DDBs serviced by this driver. This field contains the address of the first DDB in the list. The field DDB\$PS_DRVLINK in each DDB points to the next DDB in the list.
DPT\$IS_BTORDER	Ordering number for calls to the runtime drivers for boot devices.

(continued on next page)

Data Structures

3.8 DPT (Driver Prologue Table)

Table 3–10 (Cont.) Contents of Driver Prologue Table

Field	Use
DPTSL_VECTOR	Address of a driver-specific vector table. A terminal class or port driver stores the address of its class or port entry vector table in this field. For example, a terminal port driver uses this cell as a pointer to a table of addresses within the driver containing the procedure values of routines in the port driver that are called by the terminal class driver.
DPTST_NAME	<p>Name of the device driver.</p> <p>For each driver, the OpenVMS AXP driver-loading procedure constructs a 16-byte counted ASCII character string that identifies a driver and stores it in this field. The first byte records the length of the name string; the name string can be up to 15 characters.</p> <p>If you specify the /DRIVER_NAME qualifier in the SYSMAN command IO LOAD or IO CONNECT, the driver-loading procedure generates the name by extracting the filename from the full driver image specification. Otherwise, it creates the driver name from the device name (<i>ddcu</i>), appending the string "DRIVER" to the 1 to 9-character device code (<i>dd</i>).</p> <p>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.</p>
DPTSL_ECOLEVEL	ECO level of driver, taken from its image header. If for any reason this information is unavailable, the value of this field is left as zero.
DPTSQ_LINKTIME	Time and date at which driver was linked, taken from its image header.
DPTSIQ_IMAGE_NAME	Character string descriptor representing the full file specification of the driver image that has been loaded. To assist the driver loading procedure, this field is initialized as a string descriptor for the entire space available to hold the driver image file specification. The driver loading procedure writes the appropriate descriptor into this field and the driver image file specification in DPTST_IMAGE_NAME.
DPTSIL_LOADER_HANDLE	Loader handle for driver image. This field is 16-bytes long and reserved for storing a loadable image handle returned by the loadable executive image loading procedures. When the unloading of loadable executive images is implemented, the handle will be an required input to the unloading mechanism.
DPTSL_UCODE	Address of associated microcode image, if DPTSV_UCODE is set in DPTSL_FLAGS. Use of this field is reserved to Digital.
DPTSL_DECW_SNAME	Offset to a counted ASCII string that allows the SET TERMINAL/SWITCH DCL command to locate an alternate or extension DECwindows class input (decoder) driver.
DPTSQ_LMF_1	First of eight quadwords reserved to Digital for the use of the OpenVMS license management facility. (The others are DPTSQ_LMF_2, DPTSQ_LMF_3, DPTSQ_LMF_4, DPTSQ_LMF_5, DPTSQ_LMF_6, DPTSQ_LMF_7, and DPTSQ_LMF_8.)

(continued on next page)

Data Structures

3.8 DPT (Driver Prologue Table)

Table 3–10 (Cont.) Contents of Driver Prologue Table

Field	Use
DPT\$T_IMAGE_NAME	Full file specification of the driver image. This field is NAM\$C_MAXRSS long. The driver loading procedure inserts the file specification in DPT\$T_IMAGE_NAME, and the character string representing it in DPT\$IQ_IMAGE_NAME, when it loads the driver image.

3.9 IDB (Interrupt Dispatch Block)

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 supplies the physical address of the device control and status register (CSR) to the system routines that calculate the values that initialize I/O mailboxes, thus allowing device drivers to access device interface registers.

Table 3–11 describes the interrupt dispatch block.

Table 3–11 Contents of Interrupt Dispatch Block

Field	Use
IDB\$Q_CSR	Physical address of the device control and status register (CSR). IOC\$CRAM_CMD uses the CSR address in calculations that set up driver transactions to and from I/O space by means of hardware I/O mailboxes. When provided with the address of a device's CSR (for instance, in the SYSMAN command IO CONNECT), the driver-loading procedure writes the specified value into this field. The driver-loading procedure does not test the value before writing this field. For remote DSA devices and local pseudo-devices that require SCS (DPT\$IL_ADPTYPE equals AT\$_NULL and DPT\$V_SCS set in DPT\$L_FLAGS), the driver-loading procedure writes a specified SYSID into this field.
IDB\$W_SIZE	Size of IDB in bytes. The driver-loading procedure determines the size of the IDB by calculating the size of the ISB\$L_UCBLST field and adding it to the symbolic constant IDB\$K_BASE_LENGTH. It writes this sum to IDB\$W_SIZE when it creates the IDB.
IDB\$B_TYPE	Type of data structure. The driver-loading procedure writes the symbolic constant DYN\$C_IDB into this field when it creates the IDB.
IDB\$W_UNITS	Maximum number of units connected to the controller. The maximum number of units is specified in the defunits argument to the DPTAB macro and stored in DPT\$W_MAXUNITS. (The default is 8.) This value can be overridden at driver-loading time by the /MAX_UNITS qualifier to the SYSMAN command IO CONNECT. The driver-loading procedure uses this value to determine the size of the IDB\$L_UCBLST field.
IDB\$B_TT_ENABLE	Reserved for use by terminal port drivers.

(continued on next page)

Data Structures

3.9 IDB (Interrupt Dispatch Block)

Table 3–11 (Cont.) Contents of Interrupt Dispatch Block

Field	Use				
IDB\$PS_OWNER	<p>Address of UCB of device that owns controller data channel. IOCS\$PRIMITIVE_REQCHANH and IOCS\$PRIMITIVE_REQCHANL write a UCB address into this field when the routine allocates a controller data channel to a driver. IOCS\$RELCHAN confirms that the proper driver fork process is releasing a channel by comparing the driver's UCB with the UCB stored in the IDB\$PS_OWNER field. If the UCB addresses are the same, IOCS\$RELCHAN 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, IOCS\$RELCHAN 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\$PS_CRAM	<p>Header of singly linked list of CRAMs allocated to the device controller. This field contains the address of the first CRAM in the list. The field CRAM\$SL_FLINK in each CRAM points to the next CRAM in the list.</p>				
IDB\$PS_SPL	<p>Address of device lock. The driver-loading procedure copies the value of CRB\$PS_DLCK to this field.</p>				
IDB\$SL_ADP	<p>Address of the ADP associated with the device controller. The SYSMAN command IO CONNECT 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\$SL_ADP field.</p>				
IDB\$SL_FLAGS	<p>The following bits are defined within IDB\$SL_FLAGS:</p> <table style="width: 100%; border: none;"> <tr> <td style="width: 30%;">IDBSV_CRAM_ALLOC</td> <td>The driver-loading procedure has allocated the number of CRAMs specified by DPT\$IW_IDB_CRAMS and has placed them in the linked list headed by IDB\$PS_CRAM.</td> </tr> <tr> <td>IDBSV_VLE</td> <td>IDB\$SL_VECTOR points to a vector list extension (VLE)</td> </tr> </table>	IDBSV_CRAM_ALLOC	The driver-loading procedure has allocated the number of CRAMs specified by DPT\$IW_IDB_CRAMS and has placed them in the linked list headed by IDB\$PS_CRAM.	IDBSV_VLE	IDB\$SL_VECTOR points to a vector list extension (VLE)
IDBSV_CRAM_ALLOC	The driver-loading procedure has allocated the number of CRAMs specified by DPT\$IW_IDB_CRAMS and has placed them in the linked list headed by IDB\$PS_CRAM.				
IDBSV_VLE	IDB\$SL_VECTOR points to a vector list extension (VLE)				
IDB\$SL_DEVICE_SPECIFIC	<p>Longword field available to drivers for device-specific purposes.</p>				

(continued on next page)

Data Structures

3.9 IDB (Interrupt Dispatch Block)

Table 3–11 (Cont.) Contents of Interrupt Dispatch Block

Field	Use
IDBSL_VECTOR	<p>Offset of interrupt vector for this device controller, or, if IDBSV_VLE in IDBSL_VECTOR is set, the address of a vector list extension (VLE).</p> <p>For device controllers utilizing a single interrupt vector, the driver-loading procedure writes a value into this field using either the autoconfiguration database or the value specified in the /VECTOR qualifier to the SYSMAN command IO CONNECT. This value is a byte offset to device controller's vector location either in the SCB or the ADP vector table.</p> <p>For device controllers utilizing multiple interrupt vectors, the driver-loading procedure writes the address of a vector list extension (VLE) in this field. The field VLE\$SL_VECTOR_LIST in the VLE contains an array of unsigned longwords, each of which contains a byte offset to a vector location either in the SCB or the ADP vector table.</p> <p>Drivers for devices that utilize programmable interrupt vectors (that is, devices that define their interrupt vector addresses through device registers) must use this field (and, possibly, the contents of VLE\$SL_VECTOR_LIST) to load those registers during unit initialization and reinitialization after a power failure.</p>
IDBSL_UCBLST	<p>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 by the /MAX_UNITS qualifier to the SYSMAN command IO CONNECT.</p> <p>The driver-loading procedure writes a UCB address at the end of the list located at this symbolic offset in the IDB every time it creates a new UCB associated with the controller.</p>

3.10 IRP (I/O Request Packet)

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 described in Table 3–12. Note that the the standard IRP is followed by fields required by system multiprocessing code and the OpenVMS class drivers. Under no circumstances should a driver not supplied by Digital use these fields.

Data Structures

3.10 IRP (I/O Request Packet)

Table 3–12 Contents of I/O Request Packet (IRP)

Field	Use
IRP\$ <u>L</u> _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\$ <u>L</u> _IOQBL	I/O queue backward link. EXE\$INSERTIRP and IOC\$REQCOM read and write these fields.
IRP\$ <u>W</u> _SIZE	Size of IRP. EXE\$QIO writes the symbolic constant IRP\$K_LENGTH into this field when the routine allocates and fills an IRP.
IRP\$ <u>B</u> _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.
IRP\$ <u>B</u> _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 PS and writes the value into this field.
IRP\$ <u>L</u> _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\$ <u>L</u> _AST	<p>Procedure value 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.</p> <p>During I/O postprocessing, the special kernel-mode AST routine queues a mode-of-caller AST to the requesting process if this field contains the address of an AST routine.</p>
IRP\$ <u>L</u> _ASTPRM	<p>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.</p> <p>During I/O postprocessing, the special kernel-mode AST routine queues a mode-of-caller AST if the IRP\$<u>L</u>_AST field contains an address, and passes the value in IRP\$<u>L</u>_ASTPRM to the AST routine as an argument.</p>
IRP\$ <u>L</u> _OBOFF	<p>Original byte offset into the first page of a direct-I/O transfer. For segmented I/O transfers, I/O postprocessing must recalculate the value of IRP\$<u>L</u>_BOFF before transferring each segment to account for the difference between the large OpenVMS AXP memory page size and the 512-byte OpenVMS disk block size.</p> <p>FDT routines store the original byte offset in IRP\$<u>L</u>_OBOFF (as well as in IRP\$<u>L</u>_BOFF) so that that I/O postprocessing can use IRP\$<u>L</u>_OBOFF in conjunction with IRP\$<u>L</u>_OBCNT and IRP\$<u>L</u>_SVAPTE to unlock the buffer pages locked for the entire transfer.</p>

(continued on next page)

Data Structures

3.10 IRP (I/O Request Packet)

Table 3–12 (Cont.) Contents of I/O Request Packet (IRP)

Field	Use
IRPSL_WIND	<p>Address of window control block (WCB) that describes the file being accessed in the I/O request. EXESQIO writes this field if the I/O request refers to a file-structured device. An ACP or XQP reads this field.</p> <p>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. EXESQIO stores the address of this WCB in the IRPSL_WIND field.</p>
IRPSL_UCB	Address of UCB for the device assigned to the I/O channel assigned to the process. EXESQIO copies this value from the CCB.
IRPSB_EFN	Event flag number and group specified in I/O request. If the I/O request call does not specify an event flag number, EXESQIO uses event flag 0 by default. EXESQIO writes this field. The I/O postprocessing routine calls SCH\$POSTEF to set this event flag when the I/O operation is complete.
IRPSB_PRI	Base priority of the process that issued the I/O request. EXESQIO obtains a value for this field from the process control block (PCB). EXE\$INSERTIRP reads this field to insert an IRP into a priority-ordered pending-I/O queue.
IRPSB_CLN_INDEX	Shadow clone membership index. Use of this field is reserved to Digital.
IRPSB_SHD_FLAGS	Shadow clone flags. Use of this field is reserved to Digital.
IRPSL_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. EXESQIO 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 IRPSL_IOSB field with zeros so that I/O postprocessing does not write status into the IOSB.</p>
IRPSL_CHAN	Index number of process I/O channel for request. EXESQIO writes this field.
IRPSL_EXTEND	Address of first IRPE, if any, 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\$IOPOST. IRPSV_EXTEND in IRPSL_STS is set if this extension address is used.

(continued on next page)

Data Structures

3.10 IRP (I/O Request Packet)

Table 3–12 (Cont.) Contents of I/O Request Packet (IRP)

Field	Use																																
IRPSL_STS	<p>Status of I/O request. EXESQIO initializes this field to 0. EXESQIO, FDT routines, driver fork processes, or driver kernel 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 IRPSL_STS field describe the type of I/O function, as follows:</p> <table style="width: 100%; border: none;"> <tr> <td style="padding-left: 2em;">IRPSV_BUFIO</td> <td>Buffered-I/O function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_FUNC</td> <td>Read function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_PAGIO</td> <td>Paging-I/O function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_COMPLX</td> <td>Complex-buffered-I/O function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_VIRTUAL</td> <td>Virtual-I/O function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_CHAINED</td> <td>Chained-buffered-I/O function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_SWAPIO</td> <td>Swapping-I/O function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_DIAGBUF</td> <td>Diagnostic buffer is present</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_PHYSIO</td> <td>Physical-I/O function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_TERMIO</td> <td>Terminal I/O (for priority increment calculation)</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_MBXIO</td> <td>Mailbox-I/O function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_EXTEND</td> <td>An extended IRP is linked to this IRP</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_FILACP</td> <td>File ACP I/O</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_MVIRP</td> <td>Mount-verification I/O function</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_SRVIO</td> <td>Server-type I/O</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_KEY</td> <td>Encrypted function (encryption key address at IRPSL_KEYDESC)</td> </tr> </table>	IRPSV_BUFIO	Buffered-I/O function	IRPSV_FUNC	Read function	IRPSV_PAGIO	Paging-I/O function	IRPSV_COMPLX	Complex-buffered-I/O function	IRPSV_VIRTUAL	Virtual-I/O function	IRPSV_CHAINED	Chained-buffered-I/O function	IRPSV_SWAPIO	Swapping-I/O function	IRPSV_DIAGBUF	Diagnostic buffer is present	IRPSV_PHYSIO	Physical-I/O function	IRPSV_TERMIO	Terminal I/O (for priority increment calculation)	IRPSV_MBXIO	Mailbox-I/O function	IRPSV_EXTEND	An extended IRP is linked to this IRP	IRPSV_FILACP	File ACP I/O	IRPSV_MVIRP	Mount-verification I/O function	IRPSV_SRVIO	Server-type I/O	IRPSV_KEY	Encrypted function (encryption key address at IRPSL_KEYDESC)
IRPSV_BUFIO	Buffered-I/O function																																
IRPSV_FUNC	Read function																																
IRPSV_PAGIO	Paging-I/O function																																
IRPSV_COMPLX	Complex-buffered-I/O function																																
IRPSV_VIRTUAL	Virtual-I/O function																																
IRPSV_CHAINED	Chained-buffered-I/O function																																
IRPSV_SWAPIO	Swapping-I/O function																																
IRPSV_DIAGBUF	Diagnostic buffer is present																																
IRPSV_PHYSIO	Physical-I/O function																																
IRPSV_TERMIO	Terminal I/O (for priority increment calculation)																																
IRPSV_MBXIO	Mailbox-I/O function																																
IRPSV_EXTEND	An extended IRP is linked to this IRP																																
IRPSV_FILACP	File ACP I/O																																
IRPSV_MVIRP	Mount-verification I/O function																																
IRPSV_SRVIO	Server-type I/O																																
IRPSV_KEY	Encrypted function (encryption key address at IRPSL_KEYDESC)																																
IRPSL_STS2	<p>Second longword of I/O request status. EXESQIO initializes this field to 0. EXESQIO, FDT routines, and driver fork processes modify this field according to the current status of the I/O request.</p> <p>Bits in the IRPSL_STS2 field describe the type of I/O function, as follows:</p> <table style="width: 100%; border: none;"> <tr> <td style="padding-left: 2em;">IRPSV_START_PAST_HWM</td> <td>I/O starts past file highwater mark.</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_END_PAST_HWM</td> <td>I/O ends past file highwater mark.</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_ERASE</td> <td>Erase I/O function.</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_PART_HWM</td> <td>Partial file highwater mark update.</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_LCKIO</td> <td>Locked I/O request, as used by DECnet direct I/O.</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_SHDIO</td> <td>Shadowing IRP.</td> </tr> <tr> <td style="padding-left: 2em;">IRPSV_CACHEIO</td> <td>I/O using VBN cache buffers.</td> </tr> </table>	IRPSV_START_PAST_HWM	I/O starts past file highwater mark.	IRPSV_END_PAST_HWM	I/O ends past file highwater mark.	IRPSV_ERASE	Erase I/O function.	IRPSV_PART_HWM	Partial file highwater mark update.	IRPSV_LCKIO	Locked I/O request, as used by DECnet direct I/O.	IRPSV_SHDIO	Shadowing IRP.	IRPSV_CACHEIO	I/O using VBN cache buffers.																		
IRPSV_START_PAST_HWM	I/O starts past file highwater mark.																																
IRPSV_END_PAST_HWM	I/O ends past file highwater mark.																																
IRPSV_ERASE	Erase I/O function.																																
IRPSV_PART_HWM	Partial file highwater mark update.																																
IRPSV_LCKIO	Locked I/O request, as used by DECnet direct I/O.																																
IRPSV_SHDIO	Shadowing IRP.																																
IRPSV_CACHEIO	I/O using VBN cache buffers.																																

(continued on next page)

Data Structures

3.10 IRP (I/O Request Packet)

Table 3–12 (Cont.) Contents of I/O Request Packet (IRP)

Field	Use
IRPSL_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 a <i>buffered-I/O</i> transfer, address of a buffer in system address space, written here by the FDT routine allocating buffer.</p> <p>IOCSINITIATE copies this field into UCBSL_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>
IRPSL_BCNT	<p>Byte count of the I/O transfer. FDT routines calculate the count value and write the field. IOCSINITIATE copies the contents of this field into UCBSL_BCNT before calling a device driver's start-I/O routine.</p> <p>For a buffered-I/O-read function, I/O postprocessing uses IRPSL_BCNT to determine how many bytes of data to write to the user's buffer.</p>
IRPSL_BOFF	<p>Byte offset into the first (or current) page of a direct-I/O transfer. FDT routines calculate this offset and write its value into this field and IRPSL_OBOFF. For a segmented direct-I/O transfer, I/O postprocessing recalculates the value of IRPSL_BOFF before transferring each segment to account for difference between the large OpenVMS AXP memory page size and the 512-byte disk block size.</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>IOCSINITIATE copies this field into UCBSL_BOFF before calling a device driver start-I/O routine.</p> <p>I/O postprocessing uses IRPSL_BOFF in conjunction with IRPSL_BCNT and IRPSL_SVAPTE to unlock pages locked for non-segmented direct I/O transfers. For buffered I/O, I/O postprocessing adds the value of IRPSL_BOFF to the process byte count quota.</p>
IRPSPS_KPB	<p>Address of kernel process block (KPB). EXESKP_ALLOCATE_KPB, when called by EXESKP_STARTIO, returns the address of the KPB it has allocated to this field.</p>
IRPSL_IOST1	<p>First I/O status longword. IOCSREQCOM and EXESFINISHIO(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>EXESZEROPARM copies a 0 and EXESONEPARM copies p1 into this field. This field, also known as IRPSL_MEDIA, is a good place to put a \$QIO request argument. Note that, when error logging is enabled, the contents of IRPSL_MEDIA is copied into an EMB as the "disk size".</p>

(continued on next page)

Data Structures

3.10 IRP (I/O Request Packet)

Table 3–12 (Cont.) Contents of I/O Request Packet (IRP)

Field	Use						
IRP\$L_IOST2	<p>Second I/O status longword. IOCSREQCOM, EXES\$FINISHIO, and EXES\$FINISHIO(C) 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 IRPSB_CARCON. IRPSB_CARCON contains carriage control instructions to the driver. EXES\$READ and EXES\$WRITE copy the contents of p4 of the user's I/O request into this field.</p>						
IRP\$L_ABCNT	Accumulated bytes transferred in virtual I/O transfer. IOCSIOPOST reads and writes this field after a partial virtual transfer.						
IRP\$L_OBCNT	Original transfer byte count in a virtual I/O transfer. IOCSIOPOST reads this field to determine whether a virtual transfer is complete, or whether another I/O request is necessary to transfer the remaining bytes.						
IRP\$L_SEGVBN	Virtual block number of the current segment of a virtual I/O transfer. IOCSIOPOST writes this field after a partial virtual transfer.						
IRP\$L_FUNC	<p>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; EXES\$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.</p> <p>Based on this function code, EXES\$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. The upper 16 bits of this longword are reserved to Digital.</p>						
IRP\$L_DIAGBUF	<p>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, EXES\$QIO copies the buffer address into this field.</p> <p>EXES\$QIO allocates a diagnostic buffer in system address space to be filled by IOCS\$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.</p>						
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_ARB	<p>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:</p> <table style="margin-left: 2em;"> <tr> <td>ARBSQ_PRIV</td> <td>Quadword containing process privilege mask</td> </tr> <tr> <td>SPARE\$L</td> <td>Unused longword</td> </tr> <tr> <td>ARB\$L_UIC</td> <td>Longword containing process UIC</td> </tr> </table>	ARBSQ_PRIV	Quadword containing process privilege mask	SPARE\$L	Unused longword	ARB\$L_UIC	Longword containing process UIC
ARBSQ_PRIV	Quadword containing process privilege mask						
SPARE\$L	Unused longword						
ARB\$L_UIC	Longword containing process UIC						
IRP\$L_KEYDESC	Address of encryption key.						

(continued on next page)

Data Structures

3.10 IRP (I/O Request Packet)

Table 3–12 (Cont.) Contents of I/O Request Packet (IRP)

Field	Use
IRPSL_QIO_Pn	Function-specific SQIO system service arguments (p1 through p6). EXESQIO copies these arguments to the appropriate IRP fields.

3.11 IRPE (I/O Request Packet Extension)

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 64 KB. An IRPE provides space for two buffer regions, each with a 32-bit byte count.

FDT routines allocate IRPEs by calling EXESALLOCIRP. Driver routines link the IRPE to the IRP, store the IRPE's address in IRPSL_EXTEND, and set the bit field IRPE\$V_EXTEND in IRPSL_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 3–13 can store driver-dependent information.

If the IRPE specifies additional buffer regions, the FDT routine must explicitly call those buffer locking routines that call back to a driver-specified error routine if the locking procedure fails (EXES\$READLOCK_ERR, EXES\$WRITELOCK_ERR, and EXES\$MODIFYLOCK_ERR). If an error occurs during the locking procedure, the driver must unlock all previously locked regions using MMGSUNLOCK and deallocate the IRPE before returning to the buffer locking routine.

IOCSIOPOST 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. IOCSIOPOST also deallocates all the IRPEs.

The I/O request packet extension is described in Table 3–13.

Table 3–13 Contents of I/O Request Packet Extension (IRPE)

Field	Use
IRPE\$W_SIZE	Size of IRPE. EXESALLOCIRP writes the constant IRP\$K_LENGTH to this field.
IRPE\$B_TYPE	Type of data structure. EXESALLOCIRP writes the constant DYN\$C_IRP to this field.
IRPE\$L_EXTEND	Address of next IRPE, if any, for this IRP.
IRPE\$L_STS	IRPE status field. If bit IRPE\$V_EXTEND is set, it indicates that another IRPE is linked to this one.
IRPE\$L_STS2	Second longword of IRPE status field. No bits are currently defined.
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\$L_BCNT1	Size in bytes of region 1. FDT routines write this field.

(continued on next page)

Table 3–13 (Cont.) Contents of I/O Request Packet Extension (IRPE)

Field	Use
IRPE\$L_BOFF1	Byte offset 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\$L_BCNT2	Size in bytes of region 2. FDT routines write this field.
IRPE\$L_BOFF2	Byte offset of region 2. This field is set by FDT routines.

3.12 KPB (Kernel Process Block)

The kernel process block (KPB) contains the saved registers, state, and stack pointer for a kernel process.

The KPB consists of the following areas:

- **Base area.**
 The base area includes the standard OpenVMS data structure header fields, describes the kernel process stack, contains masks that describe the KPB itself and its register saveset, stores the context of a suspended KPB, and provides pointers to the other KPB areas. The KPB base area ends with offset `KPB$IS_PRM_LENGTH`.
- **Scheduling area**
 The scheduling area contains the procedure values of the routines that execute to suspend a kernel process and to resume its execution. The scheduling area can contain either a fork block or a timer queue entry. The scheduling area ends with offset `KPB$Q_FR4`.
- **Operating system special parameters area**
 The operating system special parameters area stores information required by OpenVMS device drivers, such as pointers to I/O database structures, data facilitating the selection and operation of driver macros, and driver-specific data. The OpenVMS special parameters area ends with offset `KPB$PS_DLCK`.
- **Spin lock area**
 The spin lock area is unused at present and reserved to Digital. It ends with offset `KPB$PS_SPL_RESTRT_RTN`.
- **Debugging area**
 The debugging area stores information used in the debugging of a kernel process. The KPB debugging area is contiguous with either the scheduling or spin lock KPB areas.
- **Parameter area**
 The parameter area is a variably sized area that is specified by the kernel process creator in the call to `EXE$KP_ALLOCATE_KPB`. The kernel process creator and the kernel process use this area to exchange data.

The length of each of these areas is rounded to an integral number of quadwords.

Data Structures

3.12 KPB (Kernel Process Block)

The KPB can be used in one of two general types: the OpenVMS executive software type (VEST) and the fully general type (FGT). Typically, OpenVMS software employs the VEST form of the KPB.

In a VEST KPB, the base, scheduling, OpenVMS special parameters, and spin lock areas have a fixed position relative to the starting address of the KPB. This allows you to access all fields in these areas as offsets from a single register which points to the KPB's starting address. By reducing the number of indirect reference operations, accessing VEST KPBs in this manner provides better performance than indirectly accessing the fields in the dynamic portions of a FGT KPB.

You create a VEST KPB by specifying EXE\$KP_STARTIO in the **start** argument to the DDTAB macro, or by explicitly invoking KP_ALLOCATE_KPB or calling EXE\$KP_ALLOCATE_KPB. Typically VEST KPBs do not include the debugging or parameter areas. If you require either of these areas in a VEST KPB, you must use the KPB allocation macro or routine. When present, the debugging and parameter areas are variable in size and can be located only indirectly through the pointers provided in the base KPB.

In an FGT KPB, only the base KPB and scheduling areas have a fixed position relative to the starting address of the KPB. You can reference fields in either of these areas as offsets from a KPB base pointer register. Because the other KPB areas are variably sized, you can reference them only through the pointers provided in the base KPB.

You create an FGT KPB by explicitly invoking KP_ALLOCATE_KPB or calling EXE\$KP_ALLOCATE_KPB. An FGT KPB never includes the OpenVMS special parameters area.

The base, scheduling, OpenVMS special parameters, and spin lock area are described in Table 3-14. Table 3-15 describes the debugging area.

Table 3-14 Contents of Kernel Process Block (KPB)

Field	Use
KPB\$PS_FLINK	Forward link. A driver that creates multiple kernel processes can use this field and KPB\$PS_BLINK to link together the corresponding KPBs. Doing so facilitates debugging, wherein a determined crash analysis can locate each KPB and associated kernel process stack.
KPB\$PS_BLINK	Backward link.
KPB\$IW_SIZE	Size of KPB in bytes. For VEST KPBs, EXE\$KP_ALLOCATE_KPB writes a value in this field that accounts for the presence of the base KPB, scheduling area, and spin lock area and is rounded up to a quadword multiple.
KPB\$IB_TYPE	Type of data structure. EXE\$KP_ALLOCATE_KPB writes the symbolic constant DYN\$C_MISC in this field when it creates the KPB.
KPB\$IB_SUBTYPE	Type of data structure. EXE\$KP_ALLOCATE_KPB writes the symbolic constant DYN\$C_KPB in this field when it creates the KPB.

(continued on next page)

Data Structures

3.12 KPB (Kernel Process Block)

Table 3–14 (Cont.) Contents of Kernel Process Block (KPB)

Field	Use																		
KPB\$IS_STACK_SIZE	<p>Size of kernel process stack in bytes, excluding the two guard pages. EXESKP_ALLOCATE_KPB computes the size of the kernel process stack by rounding the value of the stack_size argument up to an integral number of CPU-specific pages, converting the result to bytes, and storing it in this field.</p> <p>Note that EXESKP_STARTIO, prior to calling EXESKP_ALLOCATE_KPB, determines the size of the stack as the maximum of the value of DDT\$IS_STACK_BCNT or the symbolic constant KPBSK_MIN_IO_STACK (currently 8KB), rounded up to a multiple of CPU-specific pages.</p>																		
KPB\$IS_FLAGS	<p>The following bits are defined within KPB\$IS_FLAGS.</p> <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 30%; vertical-align: top;">KPB\$V_VALID</td> <td style="vertical-align: top;">KPB is valid. EXESKP_START sets this bit; EXESKP_END clears it.</td> </tr> <tr> <td style="vertical-align: top;">KPB\$V_ACTIVE</td> <td style="vertical-align: top;">KPB is in active use. EXESKP_START sets this bit; EXESKP_END clears it. EXESKP_STALL_GENERAL clears this bit when suspending a kernel process; EXESKP_RESTART sets it when resuming the kernel process.</td> </tr> <tr> <td style="vertical-align: top;">KPB\$V_VEST</td> <td style="vertical-align: top;">KPB is a VEST KPB. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.</td> </tr> <tr> <td style="vertical-align: top;">KPB\$V_DELETING</td> <td style="vertical-align: top;">KPB is being deleted. EXESKP_DEALLOCATE_KPB sets this bit.</td> </tr> <tr> <td style="vertical-align: top;">KPB\$V_SCHED</td> <td style="vertical-align: top;">Scheduling area is present. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.</td> </tr> <tr> <td style="vertical-align: top;">KPB\$V_SPLOCK</td> <td style="vertical-align: top;">Spin lock area is present. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.</td> </tr> <tr> <td style="vertical-align: top;">KPB\$V_DEBUG</td> <td style="vertical-align: top;">Debug area is present.</td> </tr> <tr> <td style="vertical-align: top;">KPB\$V_PARAM</td> <td style="vertical-align: top;">Parameter area is present.</td> </tr> <tr> <td style="vertical-align: top;">KPB\$V_DEALLOC_AT_END</td> <td style="vertical-align: top;">KP_END should call KP_DEALLOCATE_KPB. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.</td> </tr> </table>	KPB\$V_VALID	KPB is valid. EXESKP_START sets this bit; EXESKP_END clears it.	KPB\$V_ACTIVE	KPB is in active use. EXESKP_START sets this bit; EXESKP_END clears it. EXESKP_STALL_GENERAL clears this bit when suspending a kernel process; EXESKP_RESTART sets it when resuming the kernel process.	KPB\$V_VEST	KPB is a VEST KPB. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.	KPB\$V_DELETING	KPB is being deleted. EXESKP_DEALLOCATE_KPB sets this bit.	KPB\$V_SCHED	Scheduling area is present. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.	KPB\$V_SPLOCK	Spin lock area is present. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.	KPB\$V_DEBUG	Debug area is present.	KPB\$V_PARAM	Parameter area is present.	KPB\$V_DEALLOC_AT_END	KP_END should call KP_DEALLOCATE_KPB. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.
KPB\$V_VALID	KPB is valid. EXESKP_START sets this bit; EXESKP_END clears it.																		
KPB\$V_ACTIVE	KPB is in active use. EXESKP_START sets this bit; EXESKP_END clears it. EXESKP_STALL_GENERAL clears this bit when suspending a kernel process; EXESKP_RESTART sets it when resuming the kernel process.																		
KPB\$V_VEST	KPB is a VEST KPB. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.																		
KPB\$V_DELETING	KPB is being deleted. EXESKP_DEALLOCATE_KPB sets this bit.																		
KPB\$V_SCHED	Scheduling area is present. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.																		
KPB\$V_SPLOCK	Spin lock area is present. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.																		
KPB\$V_DEBUG	Debug area is present.																		
KPB\$V_PARAM	Parameter area is present.																		
KPB\$V_DEALLOC_AT_END	KP_END should call KP_DEALLOCATE_KPB. EXESKP_ALLOCATE_KPB sets this bit in VEST KPBs.																		
KPB\$PS_SAVED_SP	<p>Previous stack pointer. When a kernel process has been started or resumed, this field contains the value of the SP register when the executing thread is preempted (but after the registers indicated by KPB\$IS_REG_MASK have been pushed onto the stack). EXESKP_STALL_GENERAL restores this value to the SP register when the kernel process is suspended.</p>																		

(continued on next page)

Data Structures

3.12 KPB (Kernel Process Block)

Table 3–14 (Cont.) Contents of Kernel Process Block (KPB)

Field	Use
KPB\$IS_REG_MASK	<p>Kernel process register save mask. When a kernel process has been suspended, this field contains a mask of the registers that must be restored when the kernel process is resumed.</p> <p>EXESKP_STARTIO constructs this mask by merging the driver-specified register save mask (DDTSIS_REG_MASK) with the KPB minimal I/O register mask (KPREG\$K_MIN_IO_REG_MASK, which includes R2 through R5; the VAX AP, FP, SP, and PC [registers R12 through R15]; and R26, R27, and R29). Registers R0 and R1; R16 through R25; R28; and R30 and R31 (KPREG\$K_ERR_REG_MASK) cannot be saved.</p>
KPB\$PS_STACK_BASE	<p>System virtual address of the start of the no-access guard page at the base of the kernel process stack. The kernel process stack grows negatively from this address. EXESKP_ALLOCATE_KPB writes this field when it allocates the stack.</p>
KPB\$PS_STACK_SP	<p>Current kernel process SP at the time of suspension. EXESKP_STALL_GENERAL saves the current value of the SP register to this field when the kernel process is suspended, and restores to the SP register the value in KPB\$PS_SAVED_SP. When the kernel process is started, EXESKP_START initializes this field with the contents of KPB\$PS_STACK_BASE. When a kernel process is resumed, EXESKP_RESTART restores the value in this field to the SP register.</p>
KPB\$PS_SCH_PTR	<p>Address of the KPB scheduling area. EXESKP_ALLOCATE_KPB writes this field when creating the KPB. The scheduling area is contiguous with the base KPB for both VEST KPBs and FGT KPBs, and starts at offset KPB\$PS_SCH_STALL_RTN. If you reference fields in the scheduling area as offsets from the address in this field, you must use the prefix KPB\$SCH\$ in place of KPB\$ in the symbolic offsets.</p>
KPB\$PS_SPL_PTR	<p>Address of the KPB spin lock area. EXESKP_ALLOCATE_KPB writes this field when creating the KPB. The spin lock area is contiguous with the base KPB and KPB scheduling area for VEST KPBs, and starts at offset KPB\$PS_SPL_STALL_RTN. You must use the address in this field to locate the spin lock area for FGT KPBs, using the prefix KPB\$SPL\$ in place of KPB\$ in the symbolic offsets to the spin lock area's fields.</p>
KPB\$PS_DBG_PTR	<p>Address of the KPB debugging area. EXESKP_ALLOCATE_KPB writes this field when creating the KPB. See Table 3–15 for a description of the KPB debugging area. VEST KPBs do not typically include the debugging area.</p>
KPB\$PS_PRM_PTR	<p>Address of the KPB parameter area. EXESKP_ALLOCATE_KPB writes this field when creating the KPB. VEST KPBs do not typically include the parameter area.</p>

(continued on next page)

Table 3–14 (Cont.) Contents of Kernel Process Block (KPB)

Field	Use
KPB\$IS_PRM_LENGTH	Length of the KPB parameter area, as indicated in the param_length argument to EXESKP_ALLOCATE_KPB. EXESKP_ALLOCATE_KPB rounds this value up to an integral number of quadwords and writes it to this field. VEST KPBs do not typically include the parameter area.
KPB\$PS_SCH_STALL_RTN	<p>Procedure value of the routine that has been requested to suspend the kernel process described by this KPB. A kernel process scheduling stall routine preserves kernel process context not represented on the kernel process stack. It also takes steps that allow the stalled kernel process thread to be resumed at some later time (for instance, by inserting a fork block on a fork queue or by making a timer queue entry).</p> <p>A driver can implicitly specify and invoke a scheduling stall routine by calling one of the following system routines: EXESKP_FORK, EXESKP_FORK_WAIT, IOCSKP_REQCHAN, IOCSKP_WFIKPCH, or IOCSKP_WFIRLCH. (The macros KP_STALL_FORK, KP_STALL_FORK_WAIT, KP_STALL_IOFORK, KP_STALL_REQCHAN, KP_STALL_WFIKPCH, and KP_STALL_WFIRLCH may be used to call these routines.) All of these routines call EXESKP_STALL_GENERAL, which, in turn, issues a standard call to the appropriate scheduling stall routine.</p> <p>A driver can explicitly specify and invoke a scheduling stall routine by calling EXESKP_STALL_GENERAL (or invoking the KP_STALL_GENERAL macro).</p>
KPB\$PS_SCH_RESTRT_RTN	<p>Procedure value of the routine to be invoked by EXESKP_RESTART when a stalled kernel process is to be resumed.</p> <p>If the kernel process thread was suspended by EXESKP_FORK, EXESKP_FORK_WAIT, IOCSKP_REQCHAN, IOCSKP_WFIKPCH, or IOCSKP_WFIRLCH, this field contains a zero.</p> <p>A driver can explicitly specify and invoke a scheduling restart routine by calling EXESKP_STALL_GENERAL (or invoking the KP_STALL_GENERAL macro).</p>
KPB\$PS_FKBLK	Fork block address. Kernel process scheduling stall routines use this field to locate the fork block in which the kernel process thread's context is to be stored until it is resumed.
KPB\$PS_TQFL	Timer-queue forward link for embedded timer queue entry (TQE). Alternatively, as KPB\$PS_FQFL, fork-queue forward link for embedded fork block.
KPB\$PS_TQBL	Timer-queue backward link. Alternatively, as KPB\$PS_FQBL, fork-queue backward link.

(continued on next page)

Data Structures

3.12 KPB (Kernel Process Block)

Table 3–14 (Cont.) Contents of Kernel Process Block (KPB)

Field	Use
KPB\$IW_TQE_SIZE	Size of embedded TQE in bytes. Alternatively, as KPB\$IW_FKB_SIZE, size of embedded fork block in bytes. Before using this section of the KPB as a TQE or fork block, you must write the symbolic constant DYN\$C_TQE or DYN\$C_FRK, as appropriate, in this field.
KPB\$IB_FKB_TYPE	Type of data structure. Before using this section of the KPB as a TQE or fork block, you must write the symbolic constant TQESK_LENGTH or FKB\$K_LENGTH, as appropriate, in this field.
KPB\$IB_RQTYPE	Type of TQE, as described in <i>VMS for Alpha Platforms: Internals and Data Structures</i> . Before using this section of the KPB as an embedded TQE, you must indicate the TQE type in this field. Alternatively, as KPB\$IB_FLCK, this field contains the index of the fork lock that synchronizes access to the embedded fork block. Before using this section of the KPB as an embedded fork block, you must write in this field the symbolic constant (as defined by \$SPLCODDEF macro in SYSS\$LIBRARY:LIB.MLB) for the appropriate spin lock index.
KPB\$PS_FPC	Procedure value of routine at which execution resumes when the TQE becomes due or when the OpenVMS fork dispatcher dequeues the fork block. (In the latter case, EXE\$KP_FORK, EXE\$KP_IOFORK, and EXE\$KP_FORK_WAIT write this field when called to suspend driver execution.)
KPB\$Q_FR3	Value to be restored to R3 when the TQE becomes due or when the OpenVMS fork dispatcher dequeues the fork block. (In the latter case, EXE\$KP_FORK, EXE\$KP_IOFORK, and EXE\$KP_FORK_WAIT write this field when called to suspend driver execution.)
KPB\$Q_FR4	Value to be restored to R4 when the TQE becomes due or when the OpenVMS fork dispatcher dequeues the fork block. (In the latter case, EXE\$KP_FORK, EXE\$KP_IOFORK, and EXE\$KP_FORK_WAIT write this field when called to suspend driver execution.)
KPB\$IQ_TIME	Quadword system time at which a particular timer event is to occur.
KPB\$PS_UCB	UCB address. EXE\$KP_STARTIO initializes this field, which exists only in VEST KPBs. Note that this field is also known as KPB\$PS_LKB and contains the LKB address when used in lock manager operations.
KPB\$PS_IRP	IRP address. EXE\$KP_STARTIO initializes this field, which exists only in VEST KPBs.
KPB\$IS_TIMEOUT_TIME	Timeout for wait-for-interrupt operation. IOC\$KP_WFIKPC and IOC\$KP_WFIRLCH initialize this field, which is used by the corresponding scheduling stall routine when calling the appropriate basic OpenVMS suspension routine. Note that this field exists only in VEST KPBs.

(continued on next page)

Data Structures

3.12 KPB (Kernel Process Block)

Table 3–14 (Cont.) Contents of Kernel Process Block (KPB)

Field	Use								
KPB\$IS_RESTORE_IPL	IPL to be restored, and at which execution is to resume, when IOCSKP_WFIKPCH or IOCSKP_WFIRLCH returns to the initiator of the kernel process (that is, the caller of EXESKP_START or EXESKP_RESTART). IOCSKP_WFIKPCH and IOCSKP_WFIRLCH initialize this field, which is used by the corresponding scheduling stall routine when calling the appropriate basic OpenVMS suspension routine. Note that this field exists only in VEST KPBs.								
KPB\$IS_CHANNEL_DATA	Channel data passed to the request-channel scheduling stall routine (by IOCSKP_REQCHAN) and to the wait-for-interrupt scheduling stall routine (by IOCSKP_WFIKPCH or IOCSKP_WFIRLCH) to determine which basic OpenVMS suspension routine to call. Note that only VEST KPBs contain this field. VMS defines the following symbolic constants for this field: <table style="margin-left: 2em; border: none;"> <tr> <td style="padding-right: 1em;">KPB\$K_KEEP</td> <td>Keep channel as part of wait-for-interrupt operation (that is, call IOCS\$PRIMITIVE_WFIKPCH).</td> </tr> <tr> <td style="padding-right: 1em;">KPB\$K_RELEASE</td> <td>Release channel as part of wait-for-interrupt operation (that is, call IOCS\$PRIMITIVE_WFIRLCH).</td> </tr> <tr> <td style="padding-right: 1em;">KPB\$K_LOW</td> <td>Insert fork block of UCB requesting controller channel at the tail of the channel-wait queue.</td> </tr> <tr> <td style="padding-right: 1em;">KPB\$K_HIGH</td> <td>Insert fork block of UCB requesting controller channel at the head of the channel-wait queue.</td> </tr> </table>	KPB\$K_KEEP	Keep channel as part of wait-for-interrupt operation (that is, call IOCS\$PRIMITIVE_WFIKPCH).	KPB\$K_RELEASE	Release channel as part of wait-for-interrupt operation (that is, call IOCS\$PRIMITIVE_WFIRLCH).	KPB\$K_LOW	Insert fork block of UCB requesting controller channel at the tail of the channel-wait queue.	KPB\$K_HIGH	Insert fork block of UCB requesting controller channel at the head of the channel-wait queue.
KPB\$K_KEEP	Keep channel as part of wait-for-interrupt operation (that is, call IOCS\$PRIMITIVE_WFIKPCH).								
KPB\$K_RELEASE	Release channel as part of wait-for-interrupt operation (that is, call IOCS\$PRIMITIVE_WFIRLCH).								
KPB\$K_LOW	Insert fork block of UCB requesting controller channel at the tail of the channel-wait queue.								
KPB\$K_HIGH	Insert fork block of UCB requesting controller channel at the head of the channel-wait queue.								
KPB\$PS_SCSI_PTR1	Generic parameter passing field written and read by SCSI port and class drivers. Note that this field exists only in VEST KPBs.								
KPB\$PS_SCSI_PTR2	Another generic parameter passing field written and read by SCSI port and class drivers. Note that this field exists only in VEST KPBs.								
KPB\$PS_SCSI_SCDRP	Address of SCDRP used in SCSI transfers. Note that this field exists only in VEST KPBs.								
KPB\$IS_TIMEOUT	Timeout time. Note that this field exists only in VEST KPBs.								
KPB\$IS_NEWIPL	Location in which the SCSI port drivers save the current IPL when invoking the DEVICELOCK macro to synchronize access to a device's database, and from which they restore IPL when invoking the DEVICEUNLOCK macro. Note that this field exists only in VEST KPBs.								

(continued on next page)

Data Structures

3.12 KPB (Kernel Process Block)

Table 3–14 (Cont.) Contents of Kernel Process Block (KPB)

Field	Use
KPB\$PS_DLCK	Address of controller's device lock which synchronizes access to device registers and those fields in the UCB accessed at device IPL. SCSI port drivers initialize this field from SPDSL_DLCK and supply it as the lockaddr argument when invoking the DEVICELOCK and DEVICEUNLOCK macros. Note that this field exists only in VEST KPBs.
KPB\$PS_SPL_STALL_RTN	Reserved.
KPB\$PS_SPL_RESTRT_RTN	Reserved.

Table 3–15 Contents of KPB Debug Area

Field	Use
KPBDBG\$IS_START_TIME	Time at which the kernel process was started or last restarted.
KPBDBG\$IS_START_COUNT	Number of times the kernel process has been started.
KPBDBG\$IS_RESTART_COUNT	Number of times the kernel process has been restarted.
KPBDBG\$IS_VEC_INDEX	PC vector index. Indicates which longword in the PC vector index is next to be written
KPBDBG\$IS_PC_VEC	Last eight PCs which started, restarted, or suspended the kernel process.

3.13 ORB (Object Rights Block)

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

The ORB is usually allocated when the device is connected by means of a SYSMAN IO CONNECT command. The driver loading procedure also sets the address of the ORB in UCBSL_ORB at that time.

The object rights block is described in Table 3–16.

Table 3–16 Contents of Object Rights Block

Field	Use
ORB\$L_OWNER	UIC of the object's owner.
ORB\$ACL_MUTEX	Mutex for the object's access control list (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 of ORB in bytes. The driver-loading procedure writes the symbolic constant ORB\$K_LENGTH into this field when it creates an ORB.

(continued on next page)

Table 3–16 (Cont.) Contents of Object Rights Block

Field	Use
ORBSB_TYPE	Type of data structure. The driver-loading procedure writes the symbolic constant DYN\$C_ORB into this field when it creates an ORB.
ORBSB_FLAGS	Flags needed for interpreting portions of the ORB that can have alternate meanings. The following fields are defined within ORBSB_FLAGS: <div style="margin-left: 20px;"> ORBSV_PROT_16 The driver-loading procedure sets this bit to 1, signifying UIC-based protection for this object ORBSV_ACL_QUEUE This flag represents the existence of an ACL queue. The driver-loading procedure does not set this bit. ORBSV_MODE_VECTOR Use vector mode protection, not byte mode. ORBSV_NOACL This object cannot have an ACL. ORBSV_CLASS_PROT Security classification is valid. </div>
ORBSW_REFCOUNT	Reference count.
ORBSQ_MODE_PROT	Mode protection vector. The low longword of this quadword is known as ORBSL_MODE.
ORBSL_SYS_PROT	System protection field. The low word of this field is known as ORBSW_PROT and contains the standard SOGW protection.
ORBSL_OWN_PROT	Owner protection field.
ORBSL_GRP_PROT	Group protection field.
ORBSL_WOR_PROT	World protection field.
ORBSL_ACLFL	ACL queue forward link. If ORBSV_ACL_QUEUE is 0, this field should contain 0. This field is also known as ORBSL_ACL_COUNT and is cleared by the driver-loading procedure.
ORBSL_ACLBL	ACL queue backward link. If ORBSV_ACL_QUEUE is 0, this field should contain 0. This field is also known as ORBSL_ACL_DESC and is cleared by the driver-loading procedure.

3.14 UCB (Unit Control Block)

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 SYSMAN command IO 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 Step 1 driver storage.

Data Structures

3.14 UCB (Unit Control Block)

The driver-loading procedure initializes some static UCB fields when it creates the block. OpenVMS and device drivers can read and modify all nonstatic fields of the UCB. The UCB fields that are present for all devices are described in Table 3-18. The length of the basic UCB is defined by the symbol `UCBSK_LENGTH`.

UCBs are variable in length depending on the type of device and whether the driver performs error logging for the device. OpenVMS defines a number of UCB extensions in the data structure definition macro `$UCBDEF` and defines a terminal device extension in `$TTYUCBDEF`. Table 3-17 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 3-17 UCB Extensions and Sizes Defined in `$UCBDEF`

Extension	Used by	Size	Table
Base UCB	All devices	<code>UCBSK_SIZE</code>	3-18
Error log extension	All disk and tape devices	<code>UCBSK_ERL_LENGTH</code>	3-19
Dual-path extension	Reserved to Digital	<code>UCBSK_2P_LENGTH</code>	—
Local tape extension (3-20)	All tape devices	<code>UCBSK_LCL_TAPE_LENGTH</code>	value
Local disk extension (3-21)	All disk devices	<code>UCBSK_LCL_DISK_LENGTH</code>	value
Terminal extension ¹	Terminal class and port drivers	<code>UCBSK_TT_LENGTH</code>	3-22

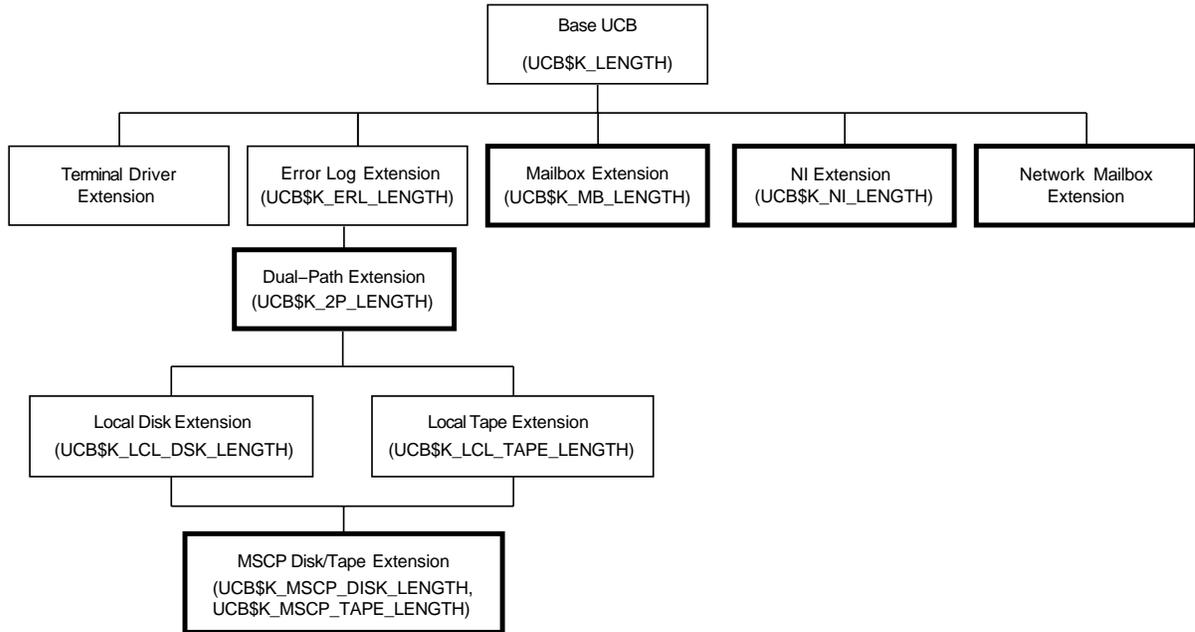
¹The terminal UCB extension is defined by the data structure definition macro, `$TTYUCBDEF`.

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

```
DPTAB  -,
      .
      .
      .
      UCBSIZE=UCBSK_LCL_TAPE_LENGTH, -
      .
      .
      .
```

As represented in Figure 3-4, each UCB extension used in a disk or tape driver builds upon the base UCB structure and any extension `$UCBDEF` defined 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 `UCBSK_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.

Figure 3-4 Composition of Extended Unit Control Blocks



Legend:
 Bold boxes indicate UCB extensions reserved for 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:

Data Structures

3.14 UCB (Unit Control Block)

```

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

```

Table 3–18 describes the contents of the unit control block.

Table 3–18 Contents of Unit Control Block

Field	Use
UCBSL_FQFL	Fork queue forward link. The link points to the next entry in the fork queue. EXE\$PRIMITIVE_FORK and OpenVMS 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.
UCBSL_FQBL	Fork queue backward link. The link points to the previous entry in the fork queue. EXE\$PRIMITIVE_FORK and OpenVMS resource management routines write this field.
UCBSW_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 Step 1 storage.
UCBSB_TYPE	Type of data structure. The driver-loading procedure writes the constant DYN\$C_UCB into this field when the procedure creates the UCB.
UCBSB_FLCK	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. When the operating system 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 multiprocessing environment. When the driver creates a fork process after an interrupt, OpenVMS inserts the fork block into a processor-specific fork queue based on this fork IPL. A 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.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use
UCBSL_FPC	<p>Procedure value of the driver fork routine. When an OpenVMS routine saves driver fork context in order to suspend driver execution, the routine stores the procedure value of the driver entry point at which execution will resume in this field. A system routine that reactivates a suspended driver transfers control to the saved PC address.</p> <p>System routines that suspend driver processing include EXESPRIMITIVE_FORK, IOCSPRIMITIVE_REQCHANL, IOCSPRIMITIVE_REQCHANH, IOCSPRIMITIVE_WFIKPCH, IOCSPRIMITIVE_WFIRLCH, EXESKP_STALL_GENERAL, EXESKP_FORK, EXESKP_FORK_WAIT, IOCSKP_REQCHAN, IOCSKP_WFIKPCH, and IOCSKP_WFIRLCH. Routines that reactivate suspended driver routines include IOCSRELCHAN, the OpenVMS fork dispatcher, 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>
UCBSQ_FR3	Value of R3 at the time that a system routine suspends a driver fork process. The value of R3 is restored just before a suspended driver regains control.
UCBSQ_FR4	Value of R4 at the time that a system routine suspends a driver fork process. The value of R4 is restored just before a suspended driver regains control.
UCBSW_BUFQUO	Buffered-I/O quota if the UCB represents a mailbox.
UCBSW_INIQUO	Initial buffered-I/O quota if the UCB represents a mailbox.
UCBSL_ORB	Address of ORB associated with the UCB. The driver-loading procedure places the address in this field.
UCBSL_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 UCBSL_DEVCHAR2).
UCB\$PS_CRAM	Header of singly linked list of CRAMs allocated to the device unit. This field contains the address of the first CRAM in the list. The field CRAM\$SL_FLINK in each CRAM points to the next CRAM in the list.
UCBSL_CRB	Address of primary CRB associated with the device. The driver-loading procedure writes this field. Driver fork processes read this field to gain access to device registers. system routines use UCBSL_CRB to locate interrupt-dispatching code and the addresses of driver unit and controller initialization routines.
UCBSL_DLCK	Address of device lock that—in a 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\$PS_DLCK) to this field as it creates a UCB for each device on a controller.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use																												
UCBSL_DDB	Address of DDB associated with device. The driver-loading procedure writes this field when the procedure creates the associated UCB. System 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.																												
UCBSL_PID	Process identification number of the process that has allocated the device. Written by the \$ALLOC system service.																												
UCBSL_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 system 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.																												
UCBSL_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.																												
UCBSL_DEVCHAR	<p>First longword of device characteristics bits. The DPT of every driver should specify symbolic constant values (defined by the \$DEVDEF macro in SYSS\$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 border="0"> <tbody> <tr> <td>DEV\$V_REC</td> <td>Record-oriented device</td> </tr> <tr> <td>DEV\$V_CCL</td> <td>Carriage control device</td> </tr> <tr> <td>DEV\$V_TRM</td> <td>Terminal device</td> </tr> <tr> <td>DEV\$V_DIR</td> <td>Directory-structured device</td> </tr> <tr> <td>DEV\$V_SDI</td> <td>Single directory-structured device</td> </tr> <tr> <td>DEV\$V_SQD</td> <td>Sequential block-oriented device (magnetic tape, for example)</td> </tr> <tr> <td>DEV\$V_SPL</td> <td>Device spooled</td> </tr> <tr> <td>DEV\$V_OPR</td> <td>Operator device</td> </tr> <tr> <td>DEV\$V_RCT</td> <td>Device contains RCT</td> </tr> <tr> <td>DEV\$V_NET</td> <td>Network device</td> </tr> <tr> <td>DEV\$V_FOD</td> <td>File-oriented device (disk and magnetic tape, for example)</td> </tr> <tr> <td>DEV\$V_DUA</td> <td>Dual-ported device</td> </tr> <tr> <td>DEV\$V_SHR</td> <td>Shareable device (used by more than one program simultaneously)</td> </tr> <tr> <td>DEV\$V_GEN</td> <td>Generic device</td> </tr> </tbody> </table>	DEV\$V_REC	Record-oriented device	DEV\$V_CCL	Carriage control device	DEV\$V_TRM	Terminal device	DEV\$V_DIR	Directory-structured device	DEV\$V_SDI	Single directory-structured device	DEV\$V_SQD	Sequential block-oriented device (magnetic tape, for example)	DEV\$V_SPL	Device spooled	DEV\$V_OPR	Operator device	DEV\$V_RCT	Device contains RCT	DEV\$V_NET	Network device	DEV\$V_FOD	File-oriented device (disk and magnetic tape, for example)	DEV\$V_DUA	Dual-ported device	DEV\$V_SHR	Shareable device (used by more than one program simultaneously)	DEV\$V_GEN	Generic device
DEV\$V_REC	Record-oriented device																												
DEV\$V_CCL	Carriage control device																												
DEV\$V_TRM	Terminal device																												
DEV\$V_DIR	Directory-structured device																												
DEV\$V_SDI	Single directory-structured device																												
DEV\$V_SQD	Sequential block-oriented device (magnetic tape, for example)																												
DEV\$V_SPL	Device spooled																												
DEV\$V_OPR	Operator device																												
DEV\$V_RCT	Device contains RCT																												
DEV\$V_NET	Network device																												
DEV\$V_FOD	File-oriented device (disk and magnetic tape, for example)																												
DEV\$V_DUA	Dual-ported device																												
DEV\$V_SHR	Shareable device (used by more than one program simultaneously)																												
DEV\$V_GEN	Generic device																												

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use
	DEV\$V_AVL Device available for use
	DEV\$V_MNT Device mounted
	DEV\$V_MBX Mailbox device
	DEV\$V_DMT Device marked for dismount
	DEV\$V_ELG Error logging enabled
	DEV\$V_ALL Device allocated
	DEV\$V_FOR Device mounted as foreign (not file structured)
	DEV\$V_SWL Device software write-locked
	DEV\$V_IDV Device capable of providing input
	DEV\$V_ODV Device capable of providing output
	DEV\$V_RND Device allowing random access
	DEV\$V_RTM Real-time device
	DEV\$V_RCK Read-checking enabled
	DEV\$V_WCK Write-checking enabled
UCBSL_DEVCHAR2	Second longword of device characteristics. The DPT of every driver should specify symbolic constant values (defined by the \$DEVDEF macro in SYSS\$LIBRARY:STARLET.MLB) for this field. The driver-loading procedure writes the field when the procedure creates the UCB.
	The system defines the following device characteristics:
	DEV\$V_CLU Device available clusterwide
	DEV\$V_DET Detached terminal
	DEV\$V_RTT Remote-terminal UCB extension
	DEV\$V_CDP Dual-pathed device with two UCBs
	DEV\$V_2P Two paths known to device
	DEV\$V_MSCP Disk or tape accessed using MSCP
	DEV\$V_SSM Shadow set member
	DEV\$V_SRV Served by MSCP server
	DEV\$V_RED Redirected terminal
	DEV\$V_NNM Device name has a prefix of the format “node\$”
	DEV\$V_WBC Device supports write-back caching
	DEV\$V_WTC Device supports write-through caching
	DEV\$V_HOC Device supports host caching
	DEV\$V_LOC Device accessible via local (non-emulated) controller
	DEV\$V_DFS Device is DFS-served
	DEV\$V_DAP Device is DAP accessed

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use
	DEV\$V_NLT Device has no bad block information on its last track
	DEV\$V_SEX Device (TAPE) supports serious exception handling
	DEV\$V_SHD Device is a member of a host based shadow set
	DEV\$V_VRT Device is a shadow set virtual unit
	DEV\$V_LDR Loader present (tapes)
	DEV\$V_NOLB Device ignores server load balancing requests
	DEV\$V_NOCLU Device will never be available clusterwide
	DEV\$V_VMEM Virtual member of a constituent set
	DEV\$V_SCSI Device is a SCSI device
	DEV\$V_WLG Device has write logging capability
	DEV\$V_NOFE Device does not support forced error
UCB\$SL_AFFINITY	Bit mask of the CPU IDs of processors in an OpenVMS 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.
UCB\$SL_XTRA	Extra longword for SMP. This field is also known as UCB\$SL_ALTLOWQ (alternate start-I/O request wait queue).
UCB\$B_DEVCLASS	Device class. The DPT of every driver should specify a symbolic constant (defined by the \$DCDEF macro in SYS\$LIBRARY:STARLET.MLB) for this field. The driver-loading procedure writes this field when it creates the UCB. 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. VMS defines the following device classes:
	DC\$_DISK Disk
	DC\$_TAPE Tape
	DC\$_SCOM Synchronous communications
	DC\$_CARD Card reader
	DC\$_TERM Terminal
	DC\$_LP Line printer

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use
	DC\$_WORKSTATION Workstation
	DC\$_REALTIME Real time. Note that the definition of a device as a real-time device (DC\$_REALTIME) is somewhat subjective; it implies no special treatment by OpenVMS.
	DC\$_BUS Bus
	DC\$_MAILBOX Mailbox
	DC\$_REMCSL_STORAGE Remote console storage
	DC\$_MISC Miscellaneous
UCB\$_DEVTYPE	Device type. The DPT of every driver should specify a symbolic constant (defined by the SDCDEF macro in SYSS\$LIBRARY:STARLET.MLB) for this field. The driver-loading procedure writes the field when it creates the UCB. 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.
UCB\$_DEVBUFSIZ	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. 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.
UCB\$_DEVDEPEND	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. 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.
UCB\$_DEVDEPND2	Second quadword for device-dependent status. This field is an extension of UCB\$_DEVDEPEND.
UCB\$_IOQFL	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. IOCSREQCOM dequeues IRPs when the device is idle. 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.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use
UCB\$SL_IOQBL	Pending-I/O queue listhead backward link. EXE\$INSERTIRP and IOCSREQCOM modify the pending-I/O queue.
UCB\$W_UNIT	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.
UCB\$W_CHARGE	Mailbox byte count quota charge, if the device is a mailbox.
UCB\$SL_IRP	Address of IRP currently being processed on the device unit by the driver fork process. IOCSINITIATE 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. The value contained in this field is not valid if the UCB\$V_BSY bit in UCB\$SL_STS is clear.
UCB\$SL_REFC	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.
UCB\$B_DIPL	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. In an OpenVMS multiprocessing environment, drivers obtain the device lock at UCB\$SL_DLCK before reading or writing device registers or accessing other fields in the UCB synchronized at device IPL, thereby also raising IPL to device IPL in the process.
UCB\$B_AMOD	Access mode at which allocation occurred, if the device is allocated. Written by the \$ALLOC and \$DALLOC system services.
UCB\$SL_AMB	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.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use
UCBSL_STS	Device unit status (formerly UCBSW_STS). Written by drivers, IOCSREQCOM, IOCSCANCELIO, IOCSINITIATE, IOCSWFIKPCCH, IOCSWFIRLCH, EXESINSIOQ, and EXESTIMEOUT. This field is read by drivers, the SQIO system service routines, IOCSREQCOM, IOCSINITIATE, and EXESTIMEOUT. This longword includes the following bits:
UCBSV_TIM	Timeout enabled.
UCBSV_INT	Interrupts expected.
UCBSV_ERLOGIP	Error log in progress.
UCBSV_CANCEL	Cancel I/O on unit.
UCBSV_ONLINE	Device is on line.
UCBSV_POWER	Power has failed while unit was busy.
UCBSV_TIMEOUT	Unit is timed out.
UCBSV_INTTYPE	Receiver interrupt.
UCBSV_BSY	Unit is busy.
UCBSV_MOUNTING	Device is being mounted.
UCBSV_DEADMO	Deallocate device at dismount.
UCBSV_VALID	Volume appears valid to software.
UCBSV_UNLOAD	Unload volume at dismount.
UCBSV_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.
UCBSV_MNTVERIP	Mount verification in progress.
UCBSV_WRONGVOL	Volume name does not match name in the VCB.
UCBSV_DELETEUCB	Delete this UCB when the value in UCBSW_REFC becomes zero.
UCBSV_LCL_VALID	The volume on this device is valid on the local node.
UCBSV_SUPMVMSG	Suppress mount-verification messages if they indicate success.
UCBSV_MNTVERPND	Mount verification is pending on the device and the device is busy.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use
	UCBSV_DISMOUNT Dismount in progress.
	UCBSV_CLUTRAN VAXcluster state transition in progress.
	UCBSV_WRTLOCKMV Write-locked mount verification in progress.
	UCBSV_SVPN_END Last byte used from page is mapped by a system virtual page number.
	UCBSV_ALTBSY Unit is busy via alternate STARTIO path.
	UCBSV_SNAPSHOT Restart validation is in progress.
UCBSL_DEVSTS	Device-dependent status. The system defines the following status bits:
	UCBSV_PRMMBX Device is a permanent mailbox. OpenVMS also defines this bitfield as UCBSV_JOB (job controller has been notified).
	UCBSV_DELMBX Mailbox is marked for deletion.
	UCBSV_SHMMBX Device is shared-memory mailbox.
	UCBSV_TEMPL_BSY Template UCB is busy.
	Disk drivers use bits in UCBSL_DEVSTS as follows:
	UCBSV_ECC ECC correction made.
	UCBSV_DIAGBUF Diagnostic buffer is specified.
	UCBSV_NOCNVRT No logical block number to media address conversion.
	UCBSV_DX_WRITE Console floppy write operation.
	UCBSV_DATACACHE Data blocks are being cached.
	Terminal class and port drivers use bits in UCBSL_DEVSTS as follows:
	UCBSV_TT_TIMO Terminal read timeout in progress.
	UCBSV_TT_NOTIF Terminal user notified of unsolicited data.
	UCBSV_TT_HANGUP Process hang up.
	UCBSV_TT_NOLOGINS No logins allowed.
UCBSL_QLEN	Number of entries in pending-I/O queue (pointed to by UCBSL_IOQFL).

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use
UCBSL_DUETIM	<p>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. IOCS\$PRIMITIVE_WFIKPCH and IOCS\$PRIMITIVE_WFIRLCH write this value when they suspend a driver to wait for an interrupt or timeout.</p> <p>EXESTIMEOUT 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, EXESTIMEOUT calls the device driver timeout handler.</p>
UCBSL_OPCNT	<p>Count of operations completed on device unit since last system bootstrap. IOCSREQCOM writes this field every time the routine inserts an IRP into the I/O postprocessing queue.</p>
UCBSL_SVPN	<p>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:</p> $(\text{index} * \text{PTESC_BYTES_PER_PTE}) + \text{MMG$GL_SPTBASE}$ <p>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 UCBSL_SVPN when the procedure creates the UCB.</p> <p>Disk drivers use this field for ECC error correction.</p>
UCBSL_SVAPTE	<p>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.</p> <p>IOCSINITIATE writes this field from IRPSL_SVAPTE before calling a driver start-I/O routine. Drivers read this value to compute the starting address of a transfer.</p>
UCBSL_BCNT	<p>Count of bytes in the I/O transfer. IOCSINITIATE copies this field from the IRP. Drivers read this field to determine how many bytes to transfer in an I/O operation.</p>
UCBSL_BOFF	<p>For a <i>direct-I/O</i> transfer, the byte offset into 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.</p> <p>IOCSINITIATE 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.</p>
UCBSL_SOFTERRCNT	<p>Reserved to Digital.</p>

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–18 (Cont.) Contents of Unit Control Block

Field	Use
UCB\$\$_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, IOCSREQCOM copies the value into the error message buffer.
UCB\$\$_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, IOCSREQCOM copies the value into the error message buffer.
UCB\$\$_ERRCNT	Number of errors that have occurred on the device since system 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\$\$_PDT	Address of port descriptor table (PDT) or SCSI port descriptor table (SPD). This field is reserved for OpenVMS SCS and SCSI port drivers.
UCB\$\$_DDT	Address of DDT for unit. The driver load procedure writes the contents of DDB\$\$_DDT for the device controller to this field when it creates the UCB.
UCB\$\$_ADP	Address of ADP. The driver-loading procedure initializes this field.
UCB\$\$_CRCTX	Address of CRCTX. A driver initializes this field when it allocates a CRCTX.
UCB\$\$_MEDIA_ID	Bit-encoded media name and type, used by MSCP devices.
UCB\$\$_DTN	Address of device-type name structure (DTN). Reserved to Digital.

Table 3–19 describes the contents of the UCB error log extension.

Table 3–19 Contents of UCB Error Log Extension

UCB\$\$_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. IOCSREQCOM writes final device status, error counters, and I/O request status into the buffer specified by this field.
UCB\$\$_FUNC	I/O function modifiers. This field is read and written by drivers that log errors.
UCB\$\$_DPC	Device-specific field. This field is reserved for driver use.

Table 3–20 describes the contents of the UCB local tape extension.

Table 3–20 Contents of UCB Local Tape Extension

Field Name	Contents
UCBSW_DIRSEQ	Directory sequence number. If the high-order bit of this word, UCBSV_AST_ARMED, is set, it indicates that the requesting process is blocking ASTs.
UCBSB_ONLCNT	Number of times the device has been placed on line since system booted.
UCBSB_PREV_RECORD	Tape position prior to the start of the last I/O operation.
UCBSL_RECORD	Current tape position or frame counter.
UCBSL_TMV_RECORD	Position following last guaranteed successful I/O operation.
UCBSW_TMV_CRC1	First CRC for mount verification's media validation.
UCBSW_TMV_CRC2	Second CRC for mount verification's media validation.
UCBSW_TMV_CRC3	Third CRC for mount verification's media validation.
UCBSW_TMV_CRC4	Fourth CRC for mount verification's media validation.

Table 3–21 describes the contents of the UCB local disk extension.

Table 3–21 Contents of UCB Local Disk Extension

Field Name	Contents
UCBSW_DIRSEQ	Directory sequence number. If the high-order bit of this word, UCBSV_AST_ARMED, is set, it indicates that the requesting process is blocking ASTs.
UCBSB_ONLCNT	Number of times device has been placed on line since OpenVMS was last bootstrapped.
UCBSL_MAXBLOCK	Maximum number of logical blocks on random-access device. This field is written by a disk driver during unit initialization and power recovery.
UCBSL_MAXBCNT	Maximum number of bytes that can be transferred. A disk driver writes this field during unit initialization and power recovery.
UCBSL_DCCB	Pointer to cache control block.
UCBSL_QLENACC	Queue length accumulator.

Table 3–22 describes the contents of the UCB terminal extension.

Data Structures

3.14 UCB (Unit Control Block)

Table 3–22 Contents of UCB Terminal Extension

Field	Use
UCB\$\$_TL_CTRLY	Listhead of CTRL/Y AST control blocks (ACBs).
UCB\$\$_TL_CTRLC	Listhead of CTRL/C ACBs.
UCB\$\$_TL_OUTBAND	Out-of-band character mask.
UCB\$\$_TL_BANDQUE	Listhead of out-of-band ACBs.
UCB\$\$_TL_PHYUCB	Address of physical UCB.
UCB\$\$_TL_CTLPID	Process ID of controlling process (used with SPAWN).
UCB\$\$_TL_BRKTHRU	Facility broadcast bit mask.
UCB\$\$_TL_POSIX_DATA	POSIX PTC pointer
UCB\$\$_TL_ASIAN_DATA	Pointer to Asian language data.
UCB\$\$_TL_A_CHARSET	Character set bitmask. The lowest byte of this field is also known as UCB\$\$_TL_A_MODE and represents the current Asian modes.
UCB\$\$_TL_A_FI_UCB	Pointer to Asian input server.
UCB\$\$_TT_RDUE	Absolute time at which a read timeout is due.
UCB\$\$_TT_RTIMOU	Address of read timeout routine.
UCB\$\$_TT_STATE1	First longword of terminal state information. The following fields are defined within UCB\$\$_TT_STATE1:
	TTY\$\$_ST_POWER Power failure
	TTY\$\$_ST_CTRLS Class output
	TTY\$\$_ST_MODEM_OFF Modem off
	TTY\$\$_ST_FILL Fill mode
	TTY\$\$_ST_CURSOR Cursor
	TTY\$\$_ST_SENDLF Forced line feed
	TTY\$\$_ST_BACKSPACE Backspace
	TTY\$\$_ST_MULTI Multi-echo
	TTY\$\$_ST_WRITE Write in progress
	TTY\$\$_ST_EOL End of line
	TTY\$\$_ST_EDITREAD Editing read in progress
	TTY\$\$_ST_RDVERIFY Read verify in progress
	TTY\$\$_ST_RECALL Command recall
	TTY\$\$_ST_READ Read in progress
	TTY\$\$_ST_POSIXREAD POSIX read
UCB\$\$_TT_STATE2	Second longword of terminal state information. The following fields are defined within UCB\$\$_TT_STATE2:
	TTY\$\$_ST_CTRLO Output enable
	TTY\$\$_ST_DEL Delete

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–22 (Cont.) Contents of UCB Terminal Extension

Field	Use
TTY\$V_ST_PASALL	Pass-all mode
TTY\$V_ST_NOECHO	No echo
TTY\$V_ST_WRTALL	Write-all mode
TTY\$V_ST_PROMPT	Prompt
TTY\$V_ST_NOFLTR	No control-character filtering
TTY\$V_ST_ESC	Escape sequence
TTY\$V_ST_BADESC	Bad escape sequence
TTY\$V_ST_NL	New line
TTY\$V_ST_REFRSH	Refresh
TTY\$V_ST_ESCAPE	Escape mode
TTY\$V_ST_TYPFUL	Type-ahead buffer full
TTY\$V_ST_SKIPLF	Skip line feed
TTY\$V_ST_ESC_O	Output escape
TTY\$V_ST_WRAP	Wrap enable
TTY\$V_ST_OVRFLO	Overflow condition
TTY\$V_ST_AUTOP	Autobaud pending
TTY\$V_ST_CTRLR	Clock prompt and data string from read buffer
TTY\$V_ST_SKIPCRLF	Skip line feed following a carriage return
TTY\$V_ST_EDITING	Editing operation
TTY\$V_ST_TABEXPAND	Expand tab characters
TTY\$V_ST_QUOTING	Quote character
TTY\$V_ST_OVERSTRIKE	Overstrike mode
TTY\$V_ST_TERMNORM	Standard terminator mask
TTY\$V_ST_ECHAES	Alternate echo string
TTY\$V_ST_PRE	Pre-type-ahead mode
TTY\$V_ST_NINTMULTI	Noninterrupt multi-echo mode
TTY\$V_ST_RECONNECT	Reconnect operation
TTY\$V_ST_CTSLOW	Clear-to-send low
TTY\$V_ST_TABRIGHT	Check for tabs to the right of the current position
UCB\$L_TT_LOGUCB	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.
UCB\$L_TT_DECHAR	First longword of default device characteristics.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–22 (Cont.) Contents of UCB Terminal Extension

Field	Use
UCB\$ <i>L</i> _TT_DECHA1	Second longword of default device characteristics.
UCB\$ <i>L</i> _TT_DECHA2	Third longword of default device characteristics.
UCB\$ <i>L</i> _TT_DECHA3	Fourth longword of default device characteristics.
UCB\$ <i>L</i> _TT_WFLINK	Write queue forward link.
UCB\$ <i>L</i> _TT_WBLINK	Write queue backward link.
UCB\$ <i>L</i> _TT_WRTBUF	Current write buffer block.
UCB\$ <i>L</i> _TT_MULTI	Address of current multi-echo buffer.
UCB\$ <i>W</i> _TT_MULTILEN	Length of multi-echo string to be written.
UCB\$ <i>W</i> _TT_SMLTLEN	Saved length of multi-echo string.
UCB\$ <i>L</i> _TT_SMLT	Saved address of multi-echo buffer.
UCB\$ <i>W</i> _TT_DESPEE	Default speed.
UCB\$ <i>B</i> _TT_DECRF	Default carriage-return fill.
UCB\$ <i>B</i> _TT_DELFF	Default line-feed fill.
UCB\$ <i>B</i> _TT_DEPARI	Default parity/character size.
UCB\$ <i>B</i> _TT_DETTYPE	Default terminal type.
UCB\$ <i>W</i> _TT_DESIZE	Default line size.
UCB\$ <i>W</i> _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\$ <i>B</i> _TT_TSPEED Transmit speed UCB\$ <i>B</i> _TT_RSPEED Receive speed
UCB\$ <i>B</i> _TT_CRFILL	Number of fill characters to be output for carriage return.
UCB\$ <i>B</i> _TT_LFFILL	Number of fill characters to be output for line feed.
UCB\$ <i>B</i> _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\$ <i>V</i> _TT_XXPARITY Reserved to Digital. UCB\$ <i>V</i> _TT_DISPARRERR Reserved to Digital. UCB\$ <i>V</i> _TT_USERFRAME Reserved to Digital. UCB\$ <i>V</i> _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\$ <i>V</i> _TT_STOP Number of stop bits: clear if one stop bit; set if two stop bits.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–22 (Cont.) Contents of UCB Terminal Extension

Field	Use
	UCBSV_TT_PARTY Parity checking. This bit is set if parity checking is enabled.
	UCBSV_TT_ODD Parity type: clear if even parity; set if odd parity.
UCBSL_TT_TYPAHD	Address of type-ahead buffer.
UCBSW_TT_CURSOR	Current cursor position.
UCBSB_TT_LINE	Current line position on page.
UCBSB_TT_LASTC	Last formatted output character.
UCBSW_TT_BSPLN	Number of back spaces to output for non-ANSI terminals.
UCBSB_TT_FILL	Current fill character count.
UCBSB_TT_ESC	Current read escape syntax state.
UCBSB_TT_ESC_O	Current write escape syntax state.
UCBSB_TT_INTCNT	Number of characters in interrupt string.
UCBSW_TT_UNITBIT	Enable and disable modem control.
UCBSW_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.
UCBSB_TT_PREMPT	Preempt character.
UCBSB_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.
UCBSL_TT_GETNXT	Address of the class driver's input routine. This field is read by the port driver.
UCBSL_TT_PUTNXT	Address of the class driver's output routine. This field is read by the port driver.
UCBSL_TT_CLASS	Address of the class driver's vector table. This field is initialized by the CLASS_CTRL_INIT macro. The port driver reads UCBSL_TT_CLASS whenever it must call the class driver at an entry point other than UCBSL_TT_GETNXT or UCBSL_TT_PUTNXT.
UCBSL_TT_PORT	Address of the port driver's vector table.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–22 (Cont.) Contents of UCB Terminal Extension

Field	Use
UCBSL_TT_OUTADR	Address of the first character of a burst of data to be written. This field is only valid when UCBSB_TT_OUTYPE contains -1. It is read and written by the port driver, and written by the class driver.
UCBSW_TT_OUTLEN	Number of characters in a burst of data to be written. This field is only valid when UCBSB_TT_OUTYPE contains -1. It is read and written by the port driver, and written by the class driver.
UCBSW_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 UCBSW_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 map registers.
TTY\$V_PC_MAPAVL	Map registers available. If set, the port driver has currently allocated map registers.
TTY\$V_PC_XOFAVL	Auto XOFF supported. If set, auto XOFF is supported for this port.
TTY\$V_PC_XOFENA	Auto XOFF enabled. If set, auto XOFF is currently enabled on this port.
TTY\$V_PC_NOCRLF	No auto line feed. If set, a line feed is not generated following a carriage return.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–22 (Cont.) Contents of UCB Terminal Extension

Field	Use
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\$L_TT_DS_ST	Current modem state.
UCB\$B_TT_DS_RCV	Current receive modem.
UCB\$B_TT_DS_TX	Current transmit modem.
UCB\$W_TT_DS_TIM	Current modem timeout.
UCB\$B_TT_MAINT	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. It contains several bits that allow the following maintenance functions:
IO\$M_LOOP	Set loopback mode.
IO\$M_UNLOOP	Reset loopback mode.

(continued on next page)

Data Structures

3.14 UCB (Unit Control Block)

Table 3–22 (Cont.) Contents of UCB Terminal Extension

Field	Use
	IOSM_AUTXOF_ENA Enable the use of auto XON/XOFF on this line. This is the default.
	IOSM_AUTXOF_DIS Disable the use of auto XON/XOFF on this line.
	IOSM_LINE_OFF Disable interrupts on this line.
	IOSM_LINE_ON Reenable interrupts on this line.
	Reference these bits by using the mask, shifted as follows:
	<pre> BITB #IOSM_LOOP@-7,- UCB\$B_TT_MAINT(R5) ;Set loopback mode </pre>
	UCB\$B_TT_MAINT also defines the bit UCB\$V_TT_DSBL that, when set, indicates that the line has been disabled.
UCB\$SL_TT_FBK	Address of fallback block.
UCB\$SL_TT_RDVERIFY	Address of read/verify table. Reserved for future use.
UCB\$SL_TT_CLASS1	First class driver longword.
UCB\$SL_TT_CLASS2	Second class driver longword.
UCB\$SL_TT_ACCPORNAM	Address of counted string.
UCB\$SL_TT_A_GCBADR	Glyph Control Block address
UCB\$W_TT_A_EDSTS	Multibyte line edit states
UCB\$B_TT_A_STATE	On-demand loading states
UCB\$B_TT_A_PARSE	ODL parse states
UCB\$B_TT_A_TRANS	JIS conversion states
UCB\$B_TT_A_XEDSTS	Extended line edit states
UCB\$SL_TT_A_DECHSET	Default char set bitmask. The lowest byte of this field is known as UCB\$B_TT_A_CHAR and represents the default Asian modes.
UCB\$SL_TP_MAP	Map registers.
UCB\$B_TP_STAT	DMA port-specific status.
	The following fields are defined within UCB\$B_TP_STAT.
	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.

3.15 VLE (Vector List Extension)

The driver loading mechanism (as directed by the SYSMAN command IO CONNECT) connects a hardware device to one or more interrupt vectors. Although most devices connected to VAX systems use preassigned vector locations, many devices on AXP systems use programmable interrupt vectors. It

Data Structures

3.15 VLE (Vector List Extension)

is the driver's responsibility to initialize such a device to use the vector or vectors to which it has been connected.

The driver loading mechanism passes this information to drivers in one of two ways:

- For devices with a single interrupt vector, the cell `IDB$L_VECTOR` contains the vector offset (into the SCB or the ADP vector table).
- For devices with multiple interrupt vectors, the cell `IDB$L_VECTOR` contains a pointer to a vector data structure which contains a list of vectors for the device.

The vector list extension is described in Table 3–23.

Table 3–23 Contents of the Vector List Extension

Field	Use
<code>VLE\$PS_IDB</code>	Address of the IDB with which the VLE is associated.
<code>VLE\$L_NUMVEC</code>	Number of vector entries in the VLE.
<code>VLE\$W_SIZE</code>	Size of VLE. The driver-loading procedure writes this field when it creates the VLE.
<code>VLE\$B_TYPE</code>	Structure type. The driver loading procedure writes the constant <code>DYN\$C_MISC</code> in this field.
<code>VLE\$B_SUBTYPE</code>	Structure subtype. The driver loading procedure writes the constant <code>DYN\$C_VLE</code> in this field.
<code>VLE\$L_VECTOR_LIST</code>	Beginning of interrupt vector list. This field is an array of unsigned longwords containing the appropriate byte offset into either the SCB or the ADP vector table.

MACRO-32 Driver Macros

This chapter describes the JSB-replacement macros, FDT completion macros, and other macros used by OpenVMS AXP device drivers.

For information about optimizing JSB-replacement macros, see the *OpenVMS AXP Device Support: Developer's Guide*.

Table 4-1 highlights some of the differences between OpenVMS VAX and OpenVMS AXP macros.

Table 4-1 New, Changed, and Unsupported OpenVMS Driver Macros

Macro	Description	Notes
ADPDISP	Causes a branch to a specified address given the existence of a selected adapter characteristic	Not supported
CLASS_UNIT_INIT	Generates the common code that must be executed by the unit initialization routine of all terminal port drivers	Changed
CPUDISP	Causes a branch to a specified address according to the CPU type of the AXP processor executing the code generated by the macro expansion	Changed
CALL_ABORTIO	Invokes FDT completion routine to abort an I/O request. Replacement for JMP EXE\$ABORTIO	New
CALL_ALTQUEPKT	Invokes FDT completion routine to queue an I/O request to the driver's alternate start I/O routine. Replacement for JSB EXE\$ALTQUEPKT	New
CALL_FINISHIO	Invokes FDT completion routine to finish an I/O request. Replacement for JMP EXE\$FINISHIO	New
CALL_FINISHIOC	Invokes FDT completion routine to finish an I/O request. Replacement for JMP EXE\$FINISHIOC	New
CALL_IORNSWAIT	Invokes FDT completion routine to wait for a resource that is required for this I/O request. Replacement for JMP EXE\$IORSNWAIT	New

(continued on next page)

MACRO-32 Driver Macros

Table 4–1 (Cont.) New, Changed, and Unsupported OpenVMS Driver Macros

Macro	Description	Notes
CALL_MODIFYLOCK_ERR	Check buffer for modify access and lock into memory. An error routine is called on any failure before the I/O request is aborted. Replacement for JSB EXESMODIFYLOCKR. See also \$DRIVER_ERRRTN_ENTRY	New
CALL_QIOACPPKT	Invokes FDT completion routine to queue an I/O request to the XQP or an ACP. Replacement for JMP EXESQIOACPPKT	New
CALL_QIODRVPKT	Invokes FDT completion routine to queue an I/O request to the driver's start I/O routine. Replacement for JMP EXESQIODRVPKT	New
CALL_READLOCK_ERR	Check buffer for read access and lock into memory. An error routine is called on any failure before the I/O request is aborted. Replacement for JSB EXESREADLOCKR. See also \$DRIVER_ERRRTN_ENTRY	New
CALL_WRITELOCK_ERR	Check buffer for read access and lock into memory. An error routine is called on any failure before the I/O request is aborted. Replacement for JSB EXESWRITELOCKR. See also \$DRIVER_ERRRTN_ENTRY	New
CRAM_ALLOC	Allocates a controller register access mailbox	New
CRAM_CMD	Calculates the COMMAND, MASK, and RBADR fields for a hardware I/O mailbox according to the requirements of a specific I/O interconnect	New
CRAM_DEALLOC	Deallocates a controller register access mailbox	New
CRAM_IO	Queues the hardware I/O mailbox defined within a controller register access mailbox (CRAM) to the mailbox pointer register (MBPR) and awaits the completion of the mailbox transaction	New
CRAM_QUEUE	Queues the hardware I/O mailbox defined within a controller register access mailbox (CRAM) to the mailbox pointer register (MBPR)	New
CRAM_WAIT	Awaits the completion of a hardware I/O mailbox transaction to a tightly coupled I/O interconnect	New
DDTAB	Generates a driver dispatch table (DDT) labeled <i>devnam\$DDT</i>	Changed
DEVICELOCK	Achieves synchronized access to a device's database as appropriate to the processing environment	Changed

(continued on next page)

Table 4–1 (Cont.) New, Changed, and Unsupported OpenVMS Driver Macros

Macro	Description	Notes
DPTAB	Generates a driver prologue table (DPT) in a program section called \$\$\$105_PROLOGUE	Changed
DPT_STORE	In the context of a DPTAB macro invocation, generates driver structure initialization and reinitialization routines which the driver loading and reloading procedures call to store values in a table or data structure	Changed
DPT_STORE_ISR	In the context of a DPTAB macro invocation, generates the addresses of the code entry point and procedure descriptor of an interrupt service routine and stores them in the interrupt transfer vector block (VEC)	New
DRIVER_CODE	Declares the program section (psect) that contains driver code	New
DRIVER_DATA	Declares the program section (psect) that contains driver data	New
SDRIVER_ALTSTART_ENTRY	Defines the driver alternate start I/O routine entry point for drivers that use the simple fork mechanism and the CALL-based fork routine environment	New
SDRIVER_CANCEL_ENTRY	Defines the driver cancel routine entry point	New
SDRIVER_CANCEL_SELECTIVE_ENTRY	Defines the driver selective cancel routine entry point	New
SDRIVER_CHANNEL_ASSIGN_ENTRY	Defines the driver channel assign routine entry point	New
SDRIVER_CLONEDUCB_ENTRY	Defines the driver cloned UCB routine entry point	New
SDRIVER_CTRLINIT_ENTRY	Defines the driver controller initialization routine entry point	New
SDRIVER_DELIVER_ENTRY	Defines the driver unit delivery routine entry point	New
SDRIVER_ERRRTN_ENTRY	Defines a driver error routine entry point. Error routines are used in conjunction with the CALL_MODIFYLOCK_ERR, CALL_READLOCK_ERR, and CALL_WRITELOCK_ERR macros	New
SDRIVER_CLONEDUCB_ENTRY	Defines the driver cloned UCB routine entry point	New
SDRIVER_FDT_ENTRY	Defines a driver upper-level FDT routine entry point	New
SDRIVER_MNTVER_ENTRY	Defines the driver mount verification routine entry point	New

(continued on next page)

MACRO-32 Driver Macros

Table 4–1 (Cont.) New, Changed, and Unsupported OpenVMS Driver Macros

Macro	Description	Notes
\$DRIVER_START_ENTRY	Defines the driver start I/O routine entry point for drivers that use the simple fork mechanism and the CALL-based fork routine environment	New
\$DRIVER_UNITINIT_ENTRY	Defines the driver unit initialization routine entry point	New
FDT_ACT	Specifies an FDT action routine for set of I/O function codes	New
FDT_BUF	Specifies the buffered functions for a function decision table	New
FDT_INI	Initializes the function decision table	New
FORK	Creates a simple fork process on the local processor	Changed
FORK_ROUTINE	Defines a fork routine entry point	New
FORK_WAIT	Inserts a fork block on the fork-and-wait queue	Changed
FORKLOCK	Achieves synchronized access to a device driver's fork database as appropriate to the processing environment	Changed
FUNCTAB	Builds a function decision table entry in an OpenVMS VAX driver	Replaced by FDT_INI, FDT_BUF, FDT_ACT
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	Replaced by TBI_ALL, TBI_DATA_64, TBI_SINGLE, and TBI_SINGLE_64 macros in OpenVMS AXP systems
IOFORK	Creates a fork process on the local processor for a device driver, disabling timeouts from the associated device	Changed
IFNORD, IFNOWRT, IFRD, IFWRT	Determines the read or write accessibility of a range of memory locations	Changed
KP_ALLOCATE_KPB	Creates a KPB and a kernel process stack, as required by the kernel process services	New
KP_DEALLOCATE_KPB	Deallocates a KPB and its associated kernel process stack	New
KP_END	Terminates the execution of a kernel process	New
KP_RESTART	Resumes the execution of a kernel process	New
KP_REQCOM	Invokes device-independent I/O postprocessing from a kernel process	New
KP_STALL_FORK, KP_STALL_IOFORK	Stall a kernel process in such a manner that it can be resumed by the fork dispatcher	New

(continued on next page)

Table 4–1 (Cont.) New, Changed, and Unsupported OpenVMS Driver Macros

Macro	Description	Notes
KP_STALL_FORK_WAIT	Stalls a kernel process in such a manner that it can be resumed by the software timer interrupt service routine's examination of the fork-and-wait queue	New
KP_STALL_GENERAL	Stalls the execution of a kernel process	New
KP_STALL_REQCHAN	Stalls a kernel process in such a manner that it can be resumed by the granting of a device controller channel	New
KP_STALL_WFIKPCH, KP_STALL_WFIRLCH	Stalls a kernel process in such a manner that it can be resumed by device interrupt processing	New
KP_START	Starts the execution of a kernel process	New
KP_SWITCH_TO_KP_STACK	Switches to kernel process context	New
LOADALT	Loads a set of Q22–bus alternate map registers	Not supported
LOADMBA	Loads MASSBUS map registers	Not supported
LOADUBA	Loads a set of UNIBUS map registers or a set of the first 496 Q22–bus map registers	Not supported
LOCK	Achieves synchronized access to a system resource as appropriate to the processing environment	Changed
RELALT	Releases a set of Q22–bus alternate map registers allocated to the driver	Not supported
RELDPR	Releases a UNIBUS adapter data path register allocated to the driver	Not supported
RELMPR	Releases a set of UNIBUS map registers or a set of the first 496 Q22–bus map registers allocated to the driver	Not supported
RELSCHAN	Releases all secondary channels allocated to the driver	Not supported
REQALT	Obtains a set of Q22–bus alternate map registers	Not supported
REQCOM	Invokes device-independent I/O postprocessing to complete an I/O request	Changed
REQCHAN	Obtains a controller's data channel	Not supported
REQDPR	Requests a UNIBUS adapter buffered data path	Not supported
REQMPR	Obtains a set of UNIBUS map registers or a set of the first 496 Q22–bus map registers	Not supported
REQPCHAN	Obtains a controller's data channel	Not supported
REQSCHAN	Obtains a secondary MASSBUS data channel	Not supported

(continued on next page)

MACRO-32 Driver Macros

Table 4–1 (Cont.) New, Changed, and Unsupported OpenVMS Driver Macros

Macro	Description	Notes
SYSDISP	Causes a branch to a specified address according to the type of AXP system executing the code in the macro expansion	New
TBI_ALL	Invalidates the data and instruction translation buffers in their entirety	New
TBI_DATA_64	Invalidates a single 64-bit virtual address in the data translation buffer	New
TBI_SINGLE	Flushes the cached contents of a single page-table entry (PTE) from the data and instruction translation buffers	New
TBI_SINGLE_64	Invalidates a single 64-bit virtual address in both the data and instruction translation buffers	New
TIMEWAIT	Waits for a specified bit to be cleared or set within a specified length of time	Not supported
TIMEDWAIT	Waits a specified interval of time for an event or condition to occur, optionally executing a series of specified instructions that test for various exit conditions	Changed
WFIKPCH, WFIRLCH	Suspends a driver fork thread and folds its context into a fork block in anticipation of a device interrupt or timeout	Changed

CALL_ABORTIO

Completes the servicing of an I/O request without returning status to the I/O status block specified in the request.

Format

CALL_ABORTIO [do_ret=YES]

Parameters

do_ret

Indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

Description

A JMP to EXE\$ABORTIO in the FDT routine of a Step 1 driver should be replaced with the CALL_ABORTIO macro. It initializes the **irp**, **pcb**, **ucb**, and **qio_status** parameters from the contents of R3, R4, R5, and R0, respectively, and calls EXE_STD\$ABORTIO. When EXE_STD\$ABORTIO returns control to the code generated by a default invocation of CALL_ABORTIO, a RET instruction returns control to the caller of CALL_ABORTIO's invoker. Status is returned in R0 and in the FDT_CONTEXT structure.

CALL_ALLOCBUF, CALL_ALLOCIRP

Allocates a buffer from nonpaged pool for a buffered-I/O operation.

Format

CALL_ALLOCBUF

CALL_ALLOCIRP

Description

A JSB to EXE\$ALLOCBUF and EXE\$ALLOCIRP in a Step 1 driver should be replaced with CALL_ALLOCBUF and CALL_ALLOCIRP, respectively. CALL_ALLOCBUF calls EXE_STD\$ALLOCBUF using the current contents of R1 as the **reqsize** argument. Both CALL_ALLOCBUF and CALL_ALLOCIRP return status in R0, the address of the allocated buffer in R2 and its size in R1. If a resource wait occurred, these macros return the address of the PCB in R4.

CALL_ALLOCEMB

Allocates an error message buffer and initializes its header.

Format

CALL_ALLOCEMB

Description

A JSB to ERL\$ALLOCEMB in a Step 1 driver should be replaced with the CALL_ALLOCEMB macro. CALL_ALLOCEMB calls ERL_STDSALLOCEMB using the current contents of R1 as the **size** argument. It returns status in R0, the address of the allocated EMB in R2 and copies the error log sequence number from EMB\$W_DV_ERRSEQ to R1.

CALL_ALTQUEPKT

Delivers an IRP to a driver's alternate start-I/O routine without regard for the status of the device.

Format

CALL_ALTQUEPKT

Description

A JSB to EXE\$ALTQUEPKT in a Step 1 driver should be replaced with the CALL_ALTQUEPKT macro. CALL_ALTQUEPKT calls EXE_STDSALTQUEPKT, using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively.

CALL_ALTREQCOM

Completes an I/O request for a device using the disk or tape class drivers.

Format

CALL_ALTREQCOM

Description

A JSB to IOC\$ALTREQCOM in a Step 1 driver should be replaced with the CALL_ALTREQCOM macro. CALL_ALTREQCOM calls IOC_STD\$ALTREQCOM, using the current contents of R0, R1, and R5 as the **iost1**, **iost2**, and **cdrp** arguments, respectively. When IOC_STD\$ALTREQCOM returns, the macro returns the address of the IRP in R3 and the address of the UCB in R4.

CALL_BROADCAST

Broadcasts the specified message to a given terminal.

Format

```
CALL_BROADCAST [save_r1]
```

Parameters

save_r1

Indicates that the macro must preserve the contents of R1 across the call to IOC_STD\$BROADCAST. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

Description

A JSB to IOC\$BROADCAST in a Step 1 driver should be replaced with the CALL_BROADCAST macro. CALL_BROADCAST calls IOC_STD\$BROADCAST, using the current contents of R1, R2, and R5 as the **msglen**, **msg_p**, and **ucb** arguments, respectively. It returns status in R0. Unless you specify **save_r1=NO**, the macro preserves the quadword register R1 across the call.

CALL_CANCELIO

Conditionally marks a UCB so that its current I/O request will be canceled.

Format

CALL_CANCELIO [save_r0r1]

Parameters

save_r0r1

Indicates that the macro must preserve the contents of R0 and R1 across the call to IOC_STD\$CANCELIO. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not preserved.)

Description

A JSB to IOC\$CANCELIO in a Step 1 driver should be replaced with the CALL_CANCELIO macro. CALL_CANCELIO calls IOC_STD\$CANCELIO, using the current contents of R2, R3, R4, and R5 as the **chan**, **irp**, **pcb**, and **ucb** arguments, respectively. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

CALL_CARRIAGE

Interprets the carriage control specifier in IRP\$B_CARCON and converts it to a generic prefix/suffix format.

Format

CALL_CARRIAGE

Description

A JSB to EXE\$CARRIAGE in a Step 1 driver should be replaced with the CALL_CARRIAGE macro. CALL_CARRIAGE calls EXE_STD\$CARRIAGE, using the current contents of R3 as the **irp** arguments.

CALL_CHKxxxACCES

Checks logical (CALL_CHKLOGACCES), physical (CALL_CHKPHYACCES), read (CALL_CHKRDACCES), write (CALL_CHKWRTACCES), execute (CALL_CHKEXEACCES), create (CALL_CHKCREACCES), or delete (CALL_CHKDELACCES) I/O function access, based on the specified protection information.

Format

```
CALL_CHKCREACCES  [save_r1]
CALL_CHKDELACCES  [save_r1]
CALL_CHKEXEACCES  [save_r1]
CALL_CHKLOGACCES  [save_r1]
CALL_CHKPHYACCES  [save_r1]
CALL_CHKRDACCES   [save_r1]
CALL_CHKWRTACCES  [save_r1]
```

Parameters

save_r1

Indicates that the macro must preserve the contents of R1 across the call to EXE_STD\$CHKPHYACCES, EXE_STD\$CHKLOGACCES, EXE_STD\$CHKWRTACCES, EXE_STD\$CHKEXEACCES, EXE_STD\$CHKCREACCES, EXE_STD\$CHKDELACCES or EXE_STD\$CHKRDACCES. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

Description

A JSB to EXE\$CHKCREACCES, EXE\$CHKDELACCES, EXE\$CHKEXEACCES, EXE\$CHKPHYACCES, EXE\$CHKLOGA, EXE\$CHKWRTACCES, or EXE\$CHKRDACCES in a Step 1 driver should be replaced with the CALL_CHKCREACCES, CALL_CHKDELACCES, CALL_CHKEXEACCES, CALL_CHKLOGACCES, CALL_CHKPHYACCES, CALL_CHKWRTACCES, or CALL_CHKRDACCES macros respectively. Each macro calls the corresponding access-checking routine, using the current contents of R0, R1, R4, and R5 as the **arb**, **orb**, **pcb**, and **ucb** arguments. Unless you specify **save_r1=NO**, the macro preserves the quadword register R1 across the call. All macros return status in R0.

CALL_CLONE_UCB

Copies a template UCB and links it to the appropriate DDB list.

Format

```
CALL_CLONE_UCB [interface_warning=YES]
```

Parameters

[interface_warning=YES]

Specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. To suppress the warning, specify **interface_warning=NO**.

Description

A JSB to IOC\$CLONE_UCB in a Step 1 driver should be replaced with CALL_CLONE_UCB. It calls IOC_STD\$CLONE_UCB using the current contents of R5 as the **tmpl_ucb** argument. CALL_CLONE_UCB returns status in R0 and the address of the newly-created UCB in R2, but does not return the address of the UCBs that precede and follow it on the DDB chain in R3 and R1, respectively.

CALL_COPY_UCB

Copies and initializes a template UCB and ORB.

Format

CALL_COPY_UCB

Description

A JSB to IOC\$COPY_UCB in a Step 1 driver should be replaced with the CALL_COPY_UCB macro. CALL_COPY_UCB calls IOC_STD\$COPY_UCB using the current contents of R5 as the **src_ucb** argument. CALL_CLONEUCB returns the address of the newly-created UCB in R2.

CALL_CREDIT_UCB

Credits the UCB charges associated with a given UCB against the process identified by the contents of UCB\$\$_CPID.

Format

CALL_CREDIT_UCB

Description

A JSB to IOC\$CREDIT_UCB in a Step 1 driver should be replaced with CALL_CREDIT_UCB. CALL_CREDIT_UCB calls IOC_STD\$CREDIT_UCB using the current contents of R5 as the **ucb** argument.

CALL_CVTLOGPHY

Conditionally converts a logical block number to a physical disk address and stores the result in the I/O request packet.

Format

CALL_CVTLOGPHY

Description

A JSB to IOC\$CVTLOGPHY in a Step 1 driver should be replaced with the CALL_CVTLOGPHY macro. CALL_CVTLOGPHY calls IOC_STD\$CVTLOGPHY, using the current contents of R0, R3, and R5 as the **lbn**, **irp** and **ucb** arguments, respectively.

CALL_CVT_DEVNAM

Converts a device name and unit number to a physical device name string.

Format

CALL_CVT_DEVNAM

Description

A JSB to IOC\$CVT_DEVNAM in a Step 1 driver should be replaced with the CALL_CVT_DEVNAM macro. CALL_CVT_DEVNAM calls IOC_STD\$CVT_DEVNAM, using the current contents of R0, R1, R4, and R5 as the **buflen**, **buf**, **form**, and **ucb** arguments, respectively.

The macro returns status in R0 and the length of the conversion string in R1.

CALL_DELATTNAST

Delivers all attention ASTs linked in the specified list.

Format

CALL_DELATTNAST [save_r0r1]

Parameters

save_r0r1

Indicates that the macro must preserve the contents of R0 and R1 across the call to COM_STD\$DELATTNAST. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

Description

A JSB to COM\$DELATTNAST in a Step 1 driver should be replaced with the CALL_DELATTNAST macro. CALL_DELATTNAST calls COM_STD\$DELATTNAST using the current contents of R4 and R5 as the **listhead** and **ucb** arguments, respectively. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

CALL_DELATTNASTP

Delivers all attention ASTs linked in the specified list for a given process.

Format

CALL_DELATTNASTP [save_r0r1]

Parameters

save_r0r1

Indicates that the macro must preserve the contents of R0 and R1 across the call to COM_STD\$DELATTNASTP. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

Description

A JSB to COM\$DELATTNASTP in a Step 1 driver should be replaced with the CALL_DELATTNASTP macro. CALL_DELATTNASTP calls COM_STD\$DELATTNASTP using the current contents of R4, R5 and R6 as the **listhead**, **ucb**, and **ipid** arguments, respectively. Unless you specify **save_r0r1=NO**, the macro preserves the quadword registers R0 and R1 across the call.

CALL_DELCTRLAST

Delivers all control ASTs, linked in the specified list, that match a given condition.

Format

CALL_DELCTRLAST [save_r0r1]

Parameters

save_r0r1

Indicates that the macro must preserve the contents of R0 and R1 across the call to COM_STD\$DELCTRLAST. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

Description

A JSB to COM\$DELCTRLAST in a Step 1 driver should be replaced with the CALL_DELCTRLAST macro. CALL_DELCTRLAST calls COM_STD\$DELCTRLAST using the current contents of R4, R5, and R3 as the **listhead**, **ucb**, and **matchchar** arguments, respectively. When COM\$DELCTRLAST returns, it moves the include character into R3. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

CALL_DELCTRLASTP

Delivers all control ASTs, linked in the specified list, that match a given condition.

Format

CALL_DELCTRLASTP [save_r0r1]

Parameters

save_r0r1

Indicates that the macro must preserve the contents of R0 and R1 across the call to COM_STD\$DELCTRLASTP. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

Description

A JSB to COM\$DELCTRLASTP in a Step 1 driver should be replaced with the CALL_DELCTRLASTP macro. CALL_DELCTRLASTP calls COM_STD\$DELCTRLASTP using the current contents of R4, R5, R6, and R3 as the **listhead**, **ucb**, **ipid**, and **matchchar** arguments, respectively. When COM\$DELCTRLASTP returns, it moves the include character into R3. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

CALL_DELETE_UCB

Deletes the specified UCB if its reference count is zero and UCBSV_DELETEUCB is set in UCBSL_STS.

Format

CALL_DELETE_UCB

Description

A JSB to IOC\$DELETE_UCB in a Step 1 driver should be replaced with the CALL_DELETE_UCB macro. CALL_DELETE_UCB calls IOC_STD\$DELETE_UCB using the current contents of R5 as the **ucb** argument.

CALL_DEVICEATTN, CALL_DEVICERR, CALL_DEVICTMO

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

Format

CALL_DEVICEATTN [save_r0r1]

CALL_DEVICERR [save_r0r1]

CALL_DEVICTMO [save_r0r1]

Parameters

save_r0r1

Indicates that the macros must preserve the contents of R0 and R1 across the call to ERL_STD\$DEVICEATTN, ERL_STD\$DEVICERR, or ERL_STD\$DEVICTMO. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

Description

JSBs to ERL\$DEVICEATTN, ERL\$DEVICERR, and ERL\$DEVICTMO in a Step 1 driver should be replaced with the CALL_DEVICEATTN, CALL_DEVICERR, and CALL_DEVICTMO macros, respectively. Each macro calls the corresponding routine using the current contents of R4 and R5 as the **driver_param** and **ucb** arguments, respectively. Unless you specify **save_r0r1=NO**, it preserves the quadword registers R0 and R1 across the call.

CALL_DIAGBUFILL

Fills a diagnostic buffer if the original \$QIO request specified such a buffer.

Format

CALL_DIAGBUFILL

Description

A JSB to IOC\$DIAGBUFILL in a Step 1 driver should be replaced with the CALL_DIAGBUFILL macro. CALL_DIAGBUFILL calls IOC_STD\$DIAGBUFILL, using the current contents of R4 and R5 as the **driver_param** and **ucb** arguments, respectively.

CALL_DRVDEALMEM

Deallocates system dynamic memory.

Format

CALL_DRVDEALMEM [save_r0r1]

Parameters

save_r0r1

Indicates that the macro must preserve the contents of R0 and R1 across the call to COM_STD\$DRVDEALMEM. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

Description

A JSB to COM\$DRVDEALMEM in a Step 1 driver should be replaced with the CALL_DRVDEALMEM macro. CALL_DRVDEALMEM calls COM_STD\$DRVDEALMEM using the current contents of R0 as the **ptr** argument. Unless you specify **save_r0r1=NO**, the macro preserves the quadword registers R0 and R1 across the call.

CALL_FILSPT

Fills a system page-table entry (PTE) with the transfer PTE of a buffer that is locked in memory so that the system PTE may be directly addressed.

Format

CALL_FILSPT

Description

A JSB to IOC\$FILSPT in a Step 1 driver should be replaced with the CALL_FILSPT macro. CALL_FILSPT calls IOC_STD\$FILSPT, passing the current contents of R5 as the **ucb** argument. It returns in R0 the system virtual address of the first byte in the page that contains the buffer.

CALL_FINISHIO, CALL_FINISHIOC, CALL_FINISHIO_NOIOST

Complete the servicing of an I/O request and return status to the I/O status block specified in the original call to the \$QIO system service.

Format

```
CALL_FINISHIO [do_ret=YES]
CALL_FINISHIOC [do_ret=YES]
CALL_FINISHIO_NOIOST [do_ret=YES]
```

Parameters

do_ret

Indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

Description

JMPs to EXE\$FINISHIO, EXE\$FINISHIOC, and EXE\$FINISHIO_NOIOST in a Step 1 driver should be replaced with the CALL_FINISHIO, CALL_FINISHIOC, and CALL_FINISHIO_NOIOST macros, respectively. CALL_FINISHIO moves the current contents of R0 and R1 into IRP\$L_IOST1 and IRP\$L_IOST2, respectively; CALL_FINISHIOC initializes IRP\$L_IOST1 from R0 and clears IRP\$L_IOST2; and CALL_FINISHIO_NOIOST fills in neither IRP field. The macros initialize the **irp** and **ucb** parameters from the contents of R3 and R5, respectively before calling EXE_STD\$FINISHIO. When EXE_STD\$FINISHIO returns control to the code generated by a default invocation of CALL_FINISHIO, CALL_FINISHIOC, or CALL_FINISHIO_NOIOST, a RET instruction returns control to the caller of the macro's invoker.

Status is returned in R0 and in the FDT_CONTEXT structure.

CALL_FLUSHATTNS

Removes specified ASTs from an attention AST list.

Format

CALL_FLUSHATTNS

Description

A JSB to COM\$FLUSHATTNS in a Step 1 driver should be replaced with the CALL_FLUSHATTNS macro. CALL_FLUSHATTNS calls COM_STD\$FLUSHATTNS using the current contents of R4, R5, R6, and R7 as the **pcb**, **ucb**, **chan**, and **acb_lh** arguments, respectively. It returns status in R0.

CALL_FLUSHCTRLS

Removes specified ASTs from a control AST list.

Format

CALL_FLUSHCTRLS

Description

A JSB to COM\$FLUSHCTRLS in a Step 1 driver should be replaced with the CALL_FLUSHCTRLS macro. CALL_FLUSHCTRLS calls COM_STD\$FLUSHCTRLS using the current contents of R2, R4, R5, R6, and R7 as the **mask**, **pcb**, **ucb**, **chan**, and **acb_lh** arguments, respectively. It returns status in R0.

CALL_GETBYTE

Fetches a single byte of data from a user buffer.

Format

CALL_GETBYTE

Description

A JSB to IOC\$GETBYTE in a Step 1 driver should be replaced with the CALL_GETBYTE macro. CALL_GETBYTE calls IOC_STD\$GETBYTE, passing the current contents of R0 and R5 as the **sva** and **ucb** arguments, respectively. It returns in R0 the byte of data (not zero-extended) returned from the user buffer. It returns in R1 the updated system virtual address. (Note that this differs from the behavior of IOC\$GETBYTE, which returns the byte of data in R1 and the updated system virtual address in R0.)

CALL_INITBUFWINDOW

Initializes a single-page window into a user buffer.

Format

CALL_INITBUFWINDOW

Description

A JSB to IOC\$INITBUFWINDOW in a Step 1 driver should be replaced with the CALL_INITBUFWINDOW macro. CALL_INITBUFWINDOW calls IOC_STD\$INITBUFWINDOW, passing the current contents of R5 as the **ucb** argument. It returns in R0 the system virtual address of the first byte in the page that contains the buffer.

CALL_INITIATE

Initiates the processing of the next I/O request for a device unit.

Format

CALL_INITIATE

Description

A JSB to IOC\$INITIATE in a Step 1 driver should be replaced with the CALL_INITIATE macro. CALL_INITIATE calls IOC_STDS\$INITIATE, using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively.

CALL_INSERT_IRP

Inserts an I/O request packet (IRP) into the specified queue of IRPs according to the base priority of the process that issued the I/O request.

Format

CALL_INSERT_IRP

Description

A JSB to EXE\$INSERT_IRP in a Step 1 driver should be replaced with the CALL_INSERT_IRP macro. CALL_INSERT_IRP calls EXE_STD\$INSERT_IRP, using the current contents of R2 and R3 as the **irp_lh** and **irp** arguments, respectively. It returns status in R0.

CALL_IOLOCK

Locks process pages in memory.

Format

CALL_IOLOCK

Description

A JSB to MMG\$ILOCK in a Step 1 driver should be replaced with the CALL_IOLOCK macro. CALL_IOLOCK calls MMG_STD\$ILOCK using the current contents of R0, R1, R2, and R4 as the **buf**, **bufsize**, **is_read**, and **pcb** arguments, respectively.

CALL_IOLOCK returns status in R0. If R0 contains SSS_NORMAL, R1 contains the system virtual address of the first page-table entry. If R0 contains zero, R1 contains the address of a page to be faulted into memory. R0 can also contain a system-level status.

CALL_IOLOCKR

Locks the I/O database mutex on behalf of its caller for read access.

Format

```
CALL_IOLOCKR  save_r1
```

Parameters

save_r1

Indicates that the macro must preserve the contents of R1 across the call to SCH_STD\$IOLOCKR. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

Description

A JSB to SCH\$IOLOCKR in a Step 1 driver should be replaced with the CALL_IOLOCKR macro. CALL_IOLOCKR calls SCH_STD\$IOLOCKR using the current contents of R4 as the **pcb** argument.

CALL_IOLOCKR returns the address of the I/O database mutex in R0. Unless you specify **save_r1=NO**, the macro preserves R1 across the call.

CALL_IOLOCKW

Locks the I/O database mutex on behalf of its caller for write access.

Format

```
CALL_IOLOCKW save_r1
```

Parameters

save_r1

Indicates that the macro must preserve the contents of R1 across the call to SCH_STDSIOLOCKW. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

Description

A JSB to SCH\$IOLOCKW in a Step 1 driver should be replaced with the CALL_IOLOCKW macro. CALL_IOLOCKW calls SCH_STDSIOLOCKW using the current contents of R4 as the **pcb** argument.

CALL_IOLOCKW returns the address of the I/O database mutex in R0. Unless you specify **save_r1=NO**, the macro preserves R1 across the call.

CALL_IORSNWAIT

Places a process in a resource wait state if it has enabled resource waits.

Format

```
CALL_IORSNWAIT [do_ret=YES]
```

Parameters

do_ret

Indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

Description

A JMP to EXE\$IORSNWAIT in a Step 1 driver should be replaced with the CALL_IORSNWAIT macro. CALL_IORSNWAIT calls EXE_STD\$IORSNWAIT using the current contents of R3, R4, R5, R6, R0, and R1 as the **irp**, **pcb**, **ucb**, **ccb**, **qio_status**, and **rsn** arguments, respectively. When EXE_STD\$IORSNWAIT returns control to the code generated by a default invocation of \$IORSNWAIT, a RET instruction returns control to the caller of the macro's invoker. Status is returned in R0 and in the FDT_CONTEXT structure.

CALL_IOUNLOCK

Releases ownership of the I/O database mutex and, if the mutex has thus become available to waiting processes, reactivates the next eligible process.

Format

CALL_IOUNLOCK

Description

A JSB to SCH\$IOUNLOCK in a Step 1 driver should be replaced with the CALL_IOUNLOCK macro. CALL_IOUNLOCK calls SCH_STD\$IOUNLOCK using the current contents of R4 as the **pcb** argument.

CALL_LINK_UCB

Searches the UCB list attached to the device data block identified by the specified UCB and links the specified UCB into the list in ascending unit number order.

Format

CALL_LINK_UCB [interface_warning=YES]

Parameters

[interface_warning=YES]

Specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. **interface_warning=NO** suppresses the warning.

Description

A JSB to IOC\$LINK_UCB in a Step 1 driver should be replaced with the CALL_LINK_UCB macro. CALL_LINK_UCB calls IOC_STD\$LINK_UCB using the current contents of R5 as the **ucb** argument. CALL_LINK_UCB returns the status in R0 and address of the newly-created UCB in R2, but does not return the address of the UCBs that precede and follow it on the DDB chain in R3 and R1, respectively.

CALL_MAPVBLK

Maps a virtual block to a logical block using a mapping window.

Format

CALL_MAPVBLK

Description

A JSB to IOC\$MAPVBLK in a Step 1 driver should be replaced with the CALL_MAPVBLK macro. CALL_MAPVBLK calls IOC_STD\$MAPVBLK, using the current contents of R0, R1, R2, R3, and R5 as the **vbn**, **numbytes**, **wcb**, **irp** and **ucb** arguments, respectively. It returns status in R0, the address of the logical block number of the first block mapped in R1, the number of unmapped bytes in R2, and the address of the updated UCB in R3. If the low bit of the status value in R0 is clear, signifying failure status, only the value in R2 is valid.

CALL_MNTVER

Assists a driver with mount verification.

Format

CALL_MNTVER

Description

A JSB to IOC\$MNTVER in a Step 1 driver should be replaced with the CALL_MNTVER macro. CALL_MNTVER calls IOC_STD\$MNTVER, using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively.

CALL_MNTVERSIO

Processes I/O functions that affect the online count and local valid status of a disk.

Format

CALL_MNTVERSIO

Description

A JSB to EXE\$MNTVERSIO in a Step 1 driver should be replaced with the CALL_MNTVERSIO macro. CALL_MNTVERSIO calls EXE_STD\$MNTVERSIO, using the current contents of R0, R3, and R5 as the **rout**, **irp**, and **ucb** arguments, respectively.

CALL_MODIFYLOCK, CALL_MODIFYLOCK_ERR

Validate and prepare a user buffer for a direct-I/O, DMA read/write operation.

Format

CALL_MODIFYLOCK

CALL_MODIFYLOCK_ERR [interface_warning=YES]

Parameters

[interface_warning=YES]

Specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. **interface_warning=NO** suppresses the warning.

Description

A JSB to EXE\$MODIFYLOCK in a Step 1 driver should be replaced with the CALL_MODIFYLOCK macro. A JSB to EXE\$MODIFYLOCK_ERR should be replaced with the CALL_MODIFYLOCK_ERR macro. CALL_MODIFYLOCK calls EXE_STD\$MODIFYLOCK, specifying 0 as the **err_rout** argument; CALL_MODIFYLOCK_ERR also calls EXE_STD\$MODIFYLOCK, using the contents of R2 as the **err_rout** argument. Both macros supply the current contents of R3, R4, R5, R6, R0, and R1 as the **irp**, **pcb**, **ucb**, **ccb**, **buf**, and **bufsize** arguments, respectively.

When EXE_STD\$MODIFYLOCK or EXE_STD\$MODIFYLOCK_ERR returns, code generated by the macro examines the return status:

- If success status (SS\$_NORMAL) is returned, the macro moves the contents of IRP\$\$_SVAPTE into R1 and writes a 5 into R2 to indicate a modify operation. Status is returned in R0 and in the FDT_CONTEXT structure.
- If failure status (SS\$_FDT_COMPL) is returned, the macro writes a 5 to R2 to indicate a modify operation and returns to FDT dispatching code in the \$QIO system service.

CALL_MOUNT_VER

During I/O postprocessing, determines whether mount verification should be initiated on a given disk or tape device on behalf of the I/O request being completed.

Format

CALL_MOUNT_VER [save_r0r1]

Parameters

save_r0r1

Indicates that the macro must preserve the contents of R0 and R1 across the call to EXE_STDSMOUNT_VER. If **save_r0r1** is blank or **save_r0r1=YES**, the 64-bit registers are saved. (In the former case, the macro generates a compile-time message. If **save_r0r1=NO**, the registers are not saved.)

Description

A JSB to EXE\$MOUNT_VER in a Step 1 driver should be replaced with the CALL_MOUNT_VER macro. CALL_MOUNT_VER calls EXE_STDSMOUNT_VER, using the current contents of R0, R1, R3, and R5 as the **iost1**, **iost2**, **irp**, and **ucb** arguments, respectively. When EXE_STDSMOUNT_VER returns, code generated by this macro copies return status from R0 to R2. Unless you specify **save_r0r1=NO**, the macro preserves the quadword registers R0 and R1 across the call.

CALL_MOVFRUSER, CALL_MOVFRUSER2

Move data from a user buffer to a device.

Format

CALL_MOVFRUSER

CALL_MOVFRUSER2

Description

JSBs to IOC\$MOVFRUSER and IOC\$MOVFRUSER2 in a Step 1 driver should be replaced with CALL_MOVFRUSER and CALL_MOVFRUSER2, respectively. CALL_MOVFRUSER calls IOC_STD\$MOVFRUSER, and CALL_MOVFRUSER2 calls IOC_STD\$MOVFRUSER2, passing the current contents of R1, R2, and R5 as the **sysbuf**, **numbytes**, and **ucb** arguments. CALL_MOVFRUSER2 also passes the current contents of R0 as the **sva** argument. Both macros return in R0 and R1, respectively, the system virtual addresses of the bytes in the internal buffer and user buffer after the last byte moved.

CALL_MOVTOUSER, CALL_MOVTOUSER2

Move data from an internal buffer to a user buffer.

Format

CALL_MOVTOUSER

CALL_MOVTOUSER2

Description

JSBs to IOC\$MOVTOUSER and IOC\$MOVTOUSER2 in a Step 1 driver should be replaced with CALL_MOVTOUSER and CALL_MOVTOUSER2, respectively. CALL_MOVTOUSER calls IOC_STD\$MOVTOUSER, and CALL_MOVTOUSER2 calls IOC_STD\$MOVTOUSER2, passing the current contents of R1, R2, and R5 as the **sysbuf**, **numbytes**, and **ucb** arguments. CALL_MOVTOUSER2 also passes the current contents of R0 as the **sva** argument. Both macros return in R0 and R1, respectively, the system virtual addresses of the bytes in the internal buffer and user buffer after the last byte moved.

CALL_PARSDEVNAM

Parses a device name string, checking its syntax and extracting the node name, allocation class number, and unit number.

Format

CALL_PARSDEVNAM

Description

A JSB to IOC\$PARSDEVNAM in a Step 1 driver should be replaced with the CALL_PARSDEVNAM macro. CALL_PARSDEVNAM calls IOC_STD\$PARSDEVNAM, using the current contents of R8, R9, and R10 as the **devnamsiz**, **devnam**, and **flags** arguments, respectively. When IOC_STD\$PARSDEVNAM returns, the macro returns status in R0; the unit number in R2; the length of the SCS node name at the beginning of the name string, allocation class number, or device type code in R3; the size of the name string in R8, the address of the name string in R9, and the flags in R10.

CALL_POST, CALL_POST_NOCNT

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

Format

CALL_POST [save_r1]

CALL_POST_NOCNT [save_r1]

Parameters

save_r1

Indicates that the macro must preserve the contents of R1 across the call to COM_STD\$POST or COM_STD\$POST_NOCNT. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

Description

A JSB to COM\$POST in a Step 1 driver should be replaced with the CALL_POST macro. CALL_POST calls COM_STD\$POST using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively. CALL_POST_NOCNT calls COM_STD\$POST_NOCNT using the current contents of R3 as the **irp** argument. Unless you specify **save_r1=NO**, the macro preserves the quadword register R1 across the call.

CALL_POST_IRP

Inserts an I/O request packet in a CPU-specific I/O postprocessing queue.

Format

CALL_POST_IRP

Description

A JSB to IOC\$POST_IRP in a Step 1 driver should be replaced with the CALL_POST_IRP macro. CALL_POST_IRP calls IOC_STD\$POST_IRP using the current contents of R3 as the **irp** argument.

CALL_PTETOPFN

Returns a page frame number (PFN) from a page-table entry (PTE) that has already been determined to be invalid.

Format

CALL_PTETOPFN

Description

A JSB to IOC\$PTETOPFN in a Step 1 driver should be replaced with the CALL_PTETOPFN macro. CALL_PTETOPFN extracts the quadword page-table entry from R3 and passes a pointer to it as the **pte** argument to IOC_STD\$PTETOPFN. It returns the page frame number in R0.

CALL_QIOACPPKT

Delivers an IRP to the appropriate ACP or XQP.

Format

```
CALL_QIOACPPKT [do_ret=YES]
```

Parameters

do_ret

Indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

Description

A JMP to EXE\$QIOACPPKT in a Step 1 driver should be replaced with the CALL_QIOACPPKT macro. CALL_QIOACPPKT calls EXE_STDSQIOACPPKT using the current contents of R3, R4, and R5 as the **irp**, **pcb**, and **ucb** arguments, respectively. When EXE_STDSQIOACPPKT returns control to the code generated by a default invocation of \$QIOACPPKT, a RET instruction returns control to the caller of the macro's invoker. Status is returned in R0 and in the FDT_CONTEXT structure.

CALL_QIODRVPKT

Delivers an IRP to the driver's start-I/O routine or pending-I/O queue.

Format

```
CALL_QIODRVPKT [do_ret=YES]
```

Parameters

do_ret

Indicates that the macro generates a RET instruction at the end of its expansion, thus returning control to the caller of the routine that invokes it.

Description

A JMP to EXE\$QIODRVPKT in a Step 1 driver should be replaced with the CALL_QIODRVPKT macro. CALL_QIODRVPKT clears IRP\$PS_FDT_CONTEXT and calls EXE_STD\$INSIOQ, using the current contents of R3 and R5 as the **irp** and **ucb** arguments, respectively. When EXE_STD\$INSIOQ returns control to the code generated by a default invocation of CALL_QIODRVPKT, a RET instruction returns control to the caller of the macro's invoker. Status is returned in R0.

CALL_QNXTSEG1

Queues the next segment of a virtual I/O request that did not map to a single contiguous I/O request.

Format

CALL_QNXTSEG1

Description

A JSB to IOC\$QNXTSEG1 in a Step 1 driver should be replaced with the CALL_QNXTSEG1 macro. CALL_QNXTSEG1 calls IOC_STD\$QNXTSEG1 using the current contents of R0, R1, R2, R3, R4, and R5 as the **vbn**, **bcnt**, **wcb**, **irp**, **pcb**, and **ucb** arguments. It returns the address of the updated UCB in R5.

CALL_QXQPPKT

Inserts an IRP on the end of the XQP work queue and initiates its processing if it is the only request on the queue.

Format

CALL_QXQPPKT

Description

A JMP to EXESQXQPPKT in a Step 1 driver should be replaced with the CALL_QXQPPKT macro. CALL_QXQPPKT calls EXE_STDSQXQPPKT using the current contents of R4 and R5 as the **pcb** and **acb** arguments, respectively. Status is returned in R0 and in the FDT_CONTEXT structure.

CALL_READCHK, CALL_READCHKR

Verifies that a process has write access to the pages in the buffer specified in a \$QIO request.

Format

CALL_READCHK
CALL_READCHKR

Description

A JSB to EXE\$READCHK in a Step 1 driver should be replaced with the CALL_READCHK macro. A JSB to EXE\$READCHKR should be replaced with the CALL_READCHKR macro. Both macros call EXE_STDS\$READCHK using the current contents of R3, R4, R5, R0, and R1 as the **irp**, **pcb**, **ucb**, **buf**, and **bufsize** arguments, respectively.

When EXE_STDS\$READCHK returns, CALL_READCHK and CALL_READCHKR move 1 into R2 to indicate a read operation and examines the return status:

- If success status (SS\$NORMAL) is returned, CALL_READCHK and CALL_READCHKR copy the contents of IRP\$BCNT into R1. CALL_READCHK writes the starting address of the I/O buffer in R0; CALL_READCHKR preserves the return status value in R0.
- If failure status (SS\$FDT_COMPL) is returned, CALL_READCHK returns to FDT dispatching code in the \$QIO system service. CALL_READCHKR does not return control to \$QIO.

CALL_READLOCK, CALL_READLOCK_ERR

Validate and prepare a user buffer for a direct-I/O, DMA write operation.

Format

CALL_READLOCK

CALL_READLOCK_ERR [interface_warning=YES]

Parameters

[interface_warning=YES]

Specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. **interface_warning=NO** suppresses the warning.

Description

A JSB to EXE\$READLOCK in a Step 1 driver should be replaced with the CALL_READLOCK macro. A JSB to EXE\$READLOCK_ERR in a Step 1 driver should be replaced with CALL_READLOCK_ERR. CALL_READLOCK calls EXE_STD\$READLOCK, specifying 0 as the **err_rout** argument; CALL_READLOCK_ERR also calls EXE_STD\$READLOCK, using the contents of R2 as the **err_rout** argument. Both macros supply the current contents of R3, R4, R5, R6, R0, and R1 as the **irp**, **pcb**, **ucb**, **ccb**, **buf**, and **bufsize** arguments, respectively.

When EXE_STD\$READLOCK or EXE_STD\$READLOCK_ERR returns, code generated by the macro examines the return status:

- If success status (SS\$_NORMAL) is returned, the macro copies the contents of IRP\$\$_SVAPTE into R1 and writes a 1 to R2 to indicate a read operation. Status is returned in R0 and in the FDT_CONTEXT structure.
- If failure status (SS\$_FDT_COMPL) is returned, the macro writes a 1 to R2 to indicate a read operation and returns to FDT dispatching code in the \$QIO system service.

CALL_RELCHAN

Releases device ownership of all controller data channels.

Format

CALL_RELCHAN

Description

A JSB to IOC\$RELCHAN in a Step 1 driver should be replaced with the CALL_RELCHAN macro. CALL_RELCHAN calls IOC_STD\$RELCHAN using the current contents of R5 as the **ucb** argument.

CALL_RELEASEMB

Releases an error message buffer to the error-logging process.

Format

CALL_RELEASEMB

Description

A JSB to ERL\$RELEASEMB in a Step 1 driver should be replaced with the CALL_RELEASEMB macro. CALL_RELEASEMB calls ERL_STD\$RELEASEMB using the current contents of R2 as the **embdv** argument.

CALL_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.

Format

CALL_REQCOM

Description

A JSB to IOC\$REQCOM in a Step 1 driver should be replaced with the CALL_REQCOM macro. CALL_REQCOM calls IOC_STD\$REQCOM, using the current contents of R0, R1, and R5 as the **iost1**, **iost2**, and **ucb** arguments, respectively.

CALL_SEARCHDEV

Searches the I/O database for a specific physical device.

Format

CALL_SEARCHDEV

Description

A JSB to IOC\$SEARCHDEV in a Step 1 driver should be replaced with the CALL_SEARCHDEV macro. CALL_SEARCHDEV calls IOC_STD\$SEARCHDEV, using the current contents of R1 as the **descr_p** argument. When IOC_STD\$SEARCHDEV returns, the macro returns status in R0, the UCB address in R1, the DDB address in R2, and the SB address in R3.

CALL_SEARCHINT

Searches the I/O database for the specified device, using specified search rules.

Format

CALL_SEARCHINT

Description

A JSB to IOC\$SEARCHINT in a Step 1 driver should be replaced with the the CALL_SEARCHINT macro. CALL_SEARCHINT calls IOC_STD\$SEARCHINT, using the current contents of R2, R3, R8, R9 and R10 as the **unit**, **sclen**, **devnamlen**, **devnam**, and **flags** arguments, respectively. When IOC_STD\$SEARCHINT returns, the macro returns status in R0, the UCB address in R5, the DDB address in R6, and the SB address in R7.

CALL_SETATTNAST

Enables or disables attention ASTs.

Format

CALL_SETATTNAST

Description

A JSB to COM\$SETATTNAST in a Step 1 driver should be replaced with the CALL_SETATTNAST macro. CALL_SETATTNAST calls COM_\$STD\$SETATTNAST using the current contents of R3, R4, R5, R6, and R7, as the **irp**, **pcb**, **ucb**, **ccb**, and **acb_lh** arguments, respectively. It returns status in R0 and in the FDT_CONTEXT structure.

CALL_SETCTRLAST

Enables or disables control ASTs.

Format

CALL_SETCTRLAST

Description

A JSB to COM\$SETCTRLAST in a Step 1 driver should be replaced with the CALL_SETCTRLAST macro. CALL_SETCTRLAST calls COM_STD\$SETCTRLAST using the current contents of R3, R4, R5, R7, and R2, as the **irp**, **pcb**, **ucb**, **acb_lh**, and **mask** arguments, respectively. It returns the TAST block in R2. It returns status in R0 and in the FDT_CONTEXT structure.

CALL_SEVER_UCB

Removes the specified UCB from the UCB list of the device data block identified within the specified UCB.

Format

CALL_SEVER_UCB

Description

A JSB to IOCSSEVER_UCB in a Step 1 driver should be replaced with the CALL_SEVER_UCB macro. CALL_SEVER_UCB calls IOC_STD\$SEVER_UCB using the current contents of R5 as the **ucb** argument.

CALL_SIMREQCOM

Completes an I/O operation by setting an event flag, modifying an I/O status block (IOSB), setting an event flag, or queuing an AST to the process requesting the I/O. The caller of this routine is responsible for checking quotas and updating the I/O count.

Format

CALL_SIMREQCOM

Description

A JSB to IOC\$SIMREQCOM in a Step 1 driver should be replaced with the CALL_SIMREQCOM macro. CALL_SIMREQCOM calls IOC_STD\$SIMREQCOM, using the current contents of R1, R2, R3, R4, R5, and R6 as the **iosb**, **pri**, **efn**, **iost**, **acb**, and **acmode** arguments, respectively.

CALL_SNDEVMSG

Builds and sends a device-specific message to the mailbox of a system process, such as the job controller or OPCOM.

Format

CALL_SNDEVMSG [save_r1]

Parameters

save_r1

Indicates that the macro must preserve the contents of R1 across the call to COM_STD\$POST. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

Description

A JSB to EXE\$SNDEVMSG in a Step 1 driver should be replaced with the the CALL_SNDEVMSG macro. CALL_SNDEVMSG calls EXE_STD\$SNDEVMSG, using the current contents of R3, R4, and R5 as the **mb_uch**, **msgtyp**, and **uch** arguments, respectively. It returns status in R0. Unless you specify **save_r1=NO**, the macro preserves the R1 across the call.

CALL_THREADCRB

Threads a controller request block (CRB) onto the due-time chain headed by IOC\$GL_CRBTMOUT.

Format

```
CALL_THREADCRB [save_r0]
```

Parameters

save_r0

Indicates that the macro must preserve the contents of R0 across the call to IOC_STD\$THREADCRB. If **save_r0** is blank or **save_r0=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r0=NO**, R0 is not saved.)

Description

A JSB to IOC\$THREADCRB in a Step 1 driver should be replaced with the CALL_THREADCRB macro. CALL_THREADCRB calls IOC_STD\$THREADCRB using the current contents of R3 as the **crb** argument. Unless you specify **save_r1=NO**, the macro preserves the quadword register R1 across the call.

CALL_UNLOCK

Unlocks process pages previously locked for a direct-I/O operation.

Format

CALL_UNLOCK

Description

A JSB to MMG\$UNLOCK in a Step 1 driver should be replaced with the CALL_UNLOCK macro. CALL_UNLOCK calls MMG_STD\$UNLOCK using the current contents of R1 and R3 as the **npages** and **svapte** arguments, respectively.

CALL_WRITECHK, CALL_WRITECHKR

Verify that a process has read access to the pages in the buffer specified in a \$QIO request.

Format

CALL_WRITECHK
CALL_WRITECHKR

Description

A JSB to EXE\$WRITECHK in a Step 1 driver should be replaced with the CALL_WRITECHK macro. A JSB to EXE\$READCHKR in a Step 1 driver should be replaced with the CALL_READCHKR macro. Both macros call EXE_STD\$READCHK using the current contents of R3, R4, R5, R0, and R1 as the **irp**, **pcb**, **ucb**, **buf**, and **bufsize** arguments, respectively.

When EXE_STD\$WRITECHK returns, CALL_WRITECHK and CALL_WRITECHKR clear R2 to indicate a write operation and examines the return status:

- If success status (SS\$NORMAL) is returned, CALL_WRITECHK and CALL_WRITECHKR copy the contents of IRPSL_BCNT into R1. CALL_WRITECHK writes the starting address of the I/O buffer in R0; CALL_WRITECHKR preserves the return status value in R0.
- If failure status (SS\$FDT_COMPL) is returned, CALL_WRITECHK returns to FDT dispatching code in the \$QIO system service. CALL_WRITECHKR does not return control to \$QIO.

CALL_WRITELOCK, CALL_WRITELOCK_ERR

Validate and prepare a user buffer for a direct-I/O, DMA read operation.

Format

CALL_WRITELOCK

CALL_WRITELOCK_ERR [interface_warning=YES]

Parameters

[interface_warning=YES]

Specifies that the macro generate a compile-time warning indicating how the behavior of the macro differs from the Step 1 version of the corresponding system routine. **interface_warning=NO** suppresses the warning.

Description

A JSB to EXE\$WRITELOCK in a Step 1 driver should be replaced with the CALL_WRITELOCK macro. A JSB to EXE\$WRITELOCK_ERR in a Step 1 driver should be replaced with the CALL_WRITELOCK_ERR macro. CALL_WRITELOCK calls EXE_STD\$WRITELOCK, specifying 0 as the **err_rout** argument; CALL_WRITELOCK_ERR also calls EXE_STD\$WRITELOCK, using the contents of R2 as the **err_rout** argument. Both macros supply the current contents of R3, R4, R5, R6, R0, and R1 as the **irp**, **pcb**, **ucb**, **ccb**, **buf**, and **bufsize** arguments, respectively.

When EXE_STD\$WRITELOCK or EXE_STD\$WRITELOCK_ERR returns, code generated by the macro examines the return status:

- If success status (SS\$NORMAL) is returned, the macro moves the contents of IRP\$L_SVAPTE into R1 and clears R2 to indicate a write operation. Status is returned in R0 and in the FDT_CONTEXT structure.
- If failure status (SS\$FDT_COMPL) is returned, the macro clears R2 to indicate a write operation and returns to FDT dispatching code in the \$QIO system service.

CALL_WRTMAILBOX

Sends a message to a mailbox.

Format

CALL_WRTMAILBOX [save_r1]

Parameters

save_r1

Indicates that the macro must preserve the contents of R1 across the call to COM_STD\$POST. If **save_r1** is blank or **save_r1=YES**, the 64-bit register is saved. (In the former case, the macro generates a compile-time message. If **save_r1=NO**, R1 is not saved.)

Description

A JSB to EXE\$WRTMAILBOX in a Step 1 driver should be replaced with the CALL_WRTMAILBOX macro. CALL_WRTMAILBOX calls EXE_STD\$WRTMAILBOX, using the current contents of R5, R3, and R4 as the **mb_uch**, **msgsiz**, and **msg** arguments, respectively. It returns status in R0. Unless you specify **save_r1=NO**, the macro preserves the R1 across the call.

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 [ucb=R5] [,port_vector=R0]

Parameters

[ucb=R5]

Address of UCB.

[port_vector=R0]

Address of port driver vector table.

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 fields as indicated:

Field	Contents
UCB\$\$_TT_CLASS	Class driver vector table address
UCB\$\$_TT_PORT	Port driver vector table address
UCB\$\$_TT_GETNXT	Procedure value of the class driver's get-next-character routine (CLASS_GETNXT)
UCB\$\$_TT_PUTNXT	Procedure value of the class driver's put-next-character routine (CLASS_PUTNXT)
UCB\$\$_TT_PARITY	Current parity, frame, and stop bit information (from TTY\$\$_PARITY)
UCB\$\$_TT_DEPARI	Default parity, frame, and stop bit information (from TTY\$\$_PARITY)
DDT\$\$_START	Procedure value of the class driver's start-I/O routine
DDT\$\$_FDT	Address of the class driver's function-decision table
DDT\$\$_CANCEL	Procedure value of the class driver's cancel-I/O routine
DDT\$\$_ALTSTART	Procedure value of the class driver's alternate start-I/O routine

OpenVMS Macros Used by OpenVMS AXP Device Drivers

CLASS_UNIT_INIT

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

Because an OpenVMS AXP terminal port driver cannot share a single DDT with the OpenVMS AXP terminal class driver, the CLASS_UNIT_INIT macro does not write the address of the class_driver's DDT into UCB\$\$_DDT. Rather, it assumes that the port driver has created its own DDT with entries for its controller initialization routine (DDT\$\$_CTRLINIT) and unit initialization routine (DDT\$\$_UNITINIT). CLASS_UNIT_INIT further initializes the port driver's DDT (the address of which it obtains from UCB\$\$_DDT) by copying to it from the class driver's DDT the procedure values of the class driver's start-I/O routine, function-decision table, cancel-I/O routine, and alternate start-I/O routine.

CPUDISP

Causes a branch to a specified address according to the CPU type of the AXP processor executing the code generated by the macro expansion.

Format

CPUDISP list [,continue=YES]

Parameters

list

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

<CPU-type, destination>

The **CPU-type** parameter identifies the type of an AXP processor for which the macro is to generate a case table entry.

The CPUDISP macro identifies the following AXP systems:

EV3	Reduced functionality AXP system
EV4	Fully functional AXP system
MANNEQUIN	AXP simulator

continue=YES

Specifies whether execution should continue at the line immediately after the CPUDISP macro if the value at EXESGQ_CPUATYPE does not correspond to any of the values specified as the **CPU-type** in the **list** argument. A fatal bugcheck of UNSUPRTCPU occurs if the dispatching code does not find the executing processor identified in the **list** 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.

CPUDISP constructs appropriate symbolic constants for each **CPU-type** listed in **list**, and compares them against the contents of EXESGQ_CPUATYPE. These constants have the form HWRPBS_CPU_TYPE\$K_CPU-type.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- With the presence of the new SYSDISP macro, the operation of the CPUDISP macro becomes less complex. The OpenVMS AXP version of CPUDISP provides a means for transferring control to a routine entry point based solely on the type of processor chip employed in the AXP system. The ability to dispatch specifically on the AXP system type (or subtype, as this parameter is called in descriptions of the OpenVMS VAX version of CPUDISP) is provided on OpenVMS AXP systems by the SYSDISP macro.
- The default value of the **continue** argument on OpenVMS AXP systems is **YES**. In other words, CPUDISP does not request the UNSUPRTCPU bugcheck by default, should you not specify the executing CPU-type in the **list** argument.

CRAM_ALLOC

Allocates a controller register access mailbox.

Format

```
CRAM_ALLOC  cram [,idb] [,ucb] [,adp]
```

Parameters

cram

Location to which the address of the allocated CRAM is returned.

[idb]

Address of IDB for device.

[ucb]

Address of UCB for device.

[adp]

Address of ADP for device.

Description

CRAM_ALLOC allocates a controller register access mailbox (CRAM) by calling IOC\$ALLOCATE_CRAM. Code must be executing at or below IPL\$SYNCH and not be holding spin locks ranked higher than IO_MISC when invoking the CRAM_ALLOC macro. For example:

```
CRAM_ALLOC      CRAM=PDT$L_R_XBE(R4), -  
                IDB=R3, -  
                UCB=R5, -  
                ADP=R2
```

CRAM_CMD

Calculates the COMMAND, MASK, and RBADR fields for a hardware I/O mailbox according to the requirements of a specific I/O interconnect.

Format

```
CRAM_CMD index ,offset ,adp [,cram] [,command]
```

Parameters

index

Command index. IOC\$CRAM_CMD uses this index to generate a mailbox command that is specific to the tightly-coupled interconnect that is to be the target of a request using this CRAM. You can specify any of the following values (defined by the \$SCRAMDEF macro), although which of these I/O operations is supported depends on the I/O interconnect that is to be the object of the mailbox operation.

Command Index	Description
CRAMCMD\$K_RDQUAD32	Quadword read in 32-bit space
CRAMCMD\$K_RDLONG32	Longword read in 32-bit space
CRAMCMD\$K_RDWORD32	Word read in 32-bit space
CRAMCMD\$K_RDBYTE32	Byte read in 32-bit space
CRAMCMD\$K_WTQUAD32	Quadword write in 32-bit space
CRAMCMD\$K_WTLONG32	Longword write in 32-bit space
CRAMCMD\$K_WTWORD32	Word write in 32-bit space
CRAMCMD\$K_WTBYTE32	Byte write in 32-bit space
CRAMCMD\$K_RDQUAD64	Quadword read in 64 bit space
CRAMCMD\$K_RDLONG64	Longword read in 64 bit space
CRAMCMD\$K_RDWORD64	Word read in 64 bit space
CRAMCMD\$K_RDBYTE64	Byte read in 64 bit space
CRAMCMD\$K_WTQUAD64	Quadword write in 64 bit space
CRAMCMD\$K_WTLONG64	Longword write in 64 bit space
CRAMCMD\$K_WTWORD64	Word write in 64 bit space
CRAMCMD\$K_WTBYTE64	Byte write in 64 bit space

offset

Byte offset of the field to be written or read from the base of device interface register (CSR) space. Calculation of the RBADR and MASK fields of the hardware mailbox depends on the addressing and masking mechanisms provided by the remote bus. The **byte_offset** parameter is used by IOC\$CRAM_CMD to calculate the RBADR, and for write operations, is used to calculate the MASK as well.

adp

Address of ADP associated with this command. IOC\$CRAM_CMD uses this parameter to determine which tightly-coupled I/O interconnect is the object of the mailbox transaction and to construct the mailbox command accordingly.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

CRAM_CMD

[cram]

Address of CRAM. If this parameter is specified, IOC\$CRAM_CMD returns the command, mask, and remote bus address values in the corresponding fields of the hardware I/O mailbox. You must specify the **cram** argument, **command** argument, or both.

[command]

Address of buffer, two quadwords in length. If this parameter is specified, IOC\$CRAM_CMD returns the command, mask, and remote bus address values in the specified buffer. You must specify the **cram** argument, **command** argument, or both.

Description

CRAM_CMD calls IOC\$CRAM_CMD to generate bus-specific values for the command, mask, and remote bus fields of the hardware I/O mailbox that is the target of the mailbox operation, inserting these values into the indicated mailbox, buffer, or both.

```
CRAM_CMD      INDEX=#CRAMCMD$K_RDLONG32,-  
              OFFSET=#XMI$L_XDEV,-  
              ADP=R2,-  
              CRAM=PDT$L_R_XDEV(R4)
```

CRAM_DEALLOC

Deallocates a controller register access mailbox.

Format

```
CRAM_DEALLOC  cram
```

Parameters

cram

Address of CRAM to be deallocated by IOC\$DEALLOCATE_CRAM.

Description

CRAM_DEALLOC deallocates a controller register access mailbox. When invoking the CRAM_DEALLOC macro, a device driver must be executing at or below IPL 8 and not be holding spin locks ranked higher than IO_MISC.

CRAM_IO

Queues the hardware I/O mailbox defined within a controller register access mailbox (CRAM) to the mailbox pointer register (MBPR) and awaits the completion of the mailbox transaction.

Format

CRAM_IO cram

Parameters

cram

Address of CRAM associated with the hardware I/O mailbox transaction.

Description

The CRAM_IO macro calls IOC\$CRAM_IO to perform an entire hardware I/O mailbox transaction from the queuing of the hardware I/O mailbox to the MBPR to the transaction's completion. Invoking the CRAM_IO macro is the equivalent to successive invocations of the CRAM_QUEUE and CRAM_WAIT macros. Prior to invoking CRAM_IO, a driver typically invokes CRAM_CMD to insert a command, mask, and remote interconnect address into the hardware I/O mailbox portion of the CRAM. For CRAMs involved in writes to device interface registers, the driver must also insert the data to be written into CRAM\$Q_WDATA.

CRAM_QUEUE

Queues the hardware I/O mailbox defined within a controller register access mailbox (CRAM) to the mailbox pointer register (MBPR).

Format

CRAM_QUEUE cram

Parameters

cram
Address of CRAM to be queued.

Description

The CRAM_QUEUE macro calls IOC\$CRAM_QUEUE to initiate an I/O operation to a device in remote I/O space by writing the physical address of the hardware I/O mailbox portion of a CRAM to the MBPR. Prior to invoking CRAM_QUEUE, a driver typically invokes CRAM_CMD to insert a command, mask, and remote interconnect address into the hardware I/O mailbox portion of the CRAM. For CRAMs involved in writes to device interface registers, the driver must also insert the data to be written into CRAM\$Q_WDATA,

It is expected that the driver will eventually invoke CRAM_WAIT to await completion of the request.

CRAM_WAIT

Awaits the completion of a hardware I/O mailbox transaction to a tightly-coupled I/O interconnect.

Format

CRAM_WAIT cram

Parameters

cram

Address of CRAM associated with a previously queued hardware I/O mailbox transaction.

Description

The CRAM_WAIT macro calls IOC\$CRAM_WAIT to check the done bit in the hardware I/O mailbox (CRAM\$V_MBX_DONE in CRAM\$W_MBX_FLAGS) and return status. It is expected that the caller has previously called IOC\$CRAM_QUEUE to post to the MBPR the hardware I/O mailbox defined within the specified CRAM for an I/O operation.

DDTAB

Generates a driver dispatch table (DDT) labeled *devnam\$DDT*.

Format

```
DDTAB devnam [,start=IOC$RETURN_SUCCESS]
      [,ctrlinit=IOC$RETURN_SUCCESS] ,functb
      [,cancel=IOC$RETURN_SUCCESS] [,regdmp=IOC$RETURN_SUCCESS]
      [,diagbf=0] [,erlgbf=0] [,unitinit=IOC$RETURN_SUCCESS]
      [,altstart=IOC$RETURN_SUCCESS] [,mntver=IOC_STD$MNTVER]
      [,cloneducb=IOC$RETURN_SUCCESS]
      [,mntv_sssc=IOC$RETURN_SUCCESS]
      [,mntv_for=IOC$RETURN_SUCCESS]
      [,mntv_sqd=IOC$RETURN_SUCCESS]
      [,channel_assign=IOC$RETURN_SUCCESS]
      [,cancel_selective=IOC$RETURN_SUCCESS] [,kp_stack_size=0]
      [,kp_reg_mask=0] [,kp_startio=IOC$RETURN_SUCCESS] [,aux_storage=0]
      [,aux_routine=IOC$RETURN_SUCCESS] [,step]
```

Parameters

devnam

Generic name of the device.

[start=IOC\$RETURN_SUCCESS]

Address of the driver's start-I/O routine. For drivers that use the kernel process services, this is the address of the kernel process start-I/O routine (EXE_STD\$KP_STARTIO).

[ctrlinit=IOC\$RETURN_SUCCESS]

Address of the controller initialization routine.

functb

Address of the driver's function decision table (FDT).

[cancel=IOC\$RETURN_SUCCESS]

Address of the cancel-I/O routine. Many drivers specify the address of the system cancel-I/O routine (IOC_STD\$CANCELIO) in this argument.

[regdmp=IOC\$RETURN_SUCCESS]

Address of the routine that dumps the device registers to an error message buffer or to a diagnostic buffer.

[diagbf=0]

Length in bytes of the diagnostic buffer.

[erlgbf=0]

Length in bytes of the error message buffer.

[unitinit=IOC\$RETURN_SUCCESS]

Address of the unit initialization routine.

[altstart=IOC\$RETURN_SUCCESS]

Address of the alternate start-I/O routine.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

DDTAB

[mntver=IOC_STD\$MNTVER]

Address of the system-provided routine that is called at the beginning and end of a mount verification operation. The default, IOC_STD\$MNTVER, is suitable for all single-stream disk drives. This argument is reserved to Digital.

[cloneducb=IOC\$RETURN_SUCCESS]

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

[mntv_sssc=IOC\$RETURN_SUCCESS]

Address of the routine called when the system performs mount verification for a shadow set state change. This argument is reserved to Digital.

[mntv_for=IOC\$RETURN_SUCCESS]

Address of the routine called when the system performs mount verification for a foreign device. This argument is reserved to Digital.

[mntv_sqd=IOC\$RETURN_SUCCESS]

Address of the routine called when the system performs mount verification for a sequential device. This argument is reserved to Digital.

[channel_assign=IOC\$RETURN_SUCCESS]

Address of the routine, called by SYSS\$ASSIGN, to complete channel assignment in a device-specific manner. This argument is reserved to Digital. (Channel-assignment routines are not yet implemented on OpenVMS AXP systems.)

[cancel_selective=IOC\$RETURN_SUCCESS]

Address of the routine that cancels a list of I/O requests from the specified channel, including both waiting and active requests. This argument is reserved to Digital. (Cancel selective routines are not yet implemented on OpenVMS AXP systems.)

[kp_stack_size=0]

Size in bytes of the kernel process stack. EXE_STD\$KP_STARTIO uses this value, or KPBSK_MIN_IO_STACK (currently 8KB), whichever is larger, to determine the size of the stack created for the driver's start I/O kernel process thread.

[kp_reg_mask=0]

Kernel process register save mask.

This mask represents those registers used by a kernel process that must be preserved across kernel process context switches. R12 through R15, R26, R27, and R29 (KPREG\$K_MIN_REG_MASK) are always preserved across kernel process context switches, and that EXE\$KP_STARTIO additionally includes R2 through R5 in this register set (KPREG\$K_MIN_IO_REG_MASK). R0, R1, R16 through R25, R27, R28, R30, and R31 (KPREG\$K_ERR_REG_MASK) are never preserved and are illegal in a register save mask.

[kp_startio=IOC\$RETURN_SUCCESS]

Address of the start-I/O routine of a driver that uses the kernel process services. Such a driver typically specifies the system routine EXE_STD\$KP_STARTIO in the **start** argument to the DDTAB macro. EXE_STD\$KP_STARTIO calls the start-I/O routine specified in this argument after setting up the kernel process environment.

OpenVMS Macros Used by OpenVMS AXP Device Drivers DDTAB

[aux_storage=0]

Address of auxiliary storage area. This argument is reserved to Digital. (Auxiliary storage areas are not yet implemented on OpenVMS AXP systems.)

[aux_routine=IOC\$RETURN_SUCCESS]

Address of an auxiliary routine in the OpenVMS VAX mailbox driver that is called by SYSS\$ASSIGN. This argument is reserved to Digital. (Auxiliary routines are not yet implemented on OpenVMS AXP systems.)

[step]

OpenVMS AXP driver step number. You may indicate that a given driver conforms to the coding practices for an Step 1 OpenVMS AXP device driver by supplying **step=2** in the DDTAB macro invocation. If you previously specified the **step** argument to the DPTAB macro, you need not repeat it here.

If you supply the **step** argument, but specify a value other than 1 or 2, the DPTAB macro generates the following message:

```
%MACRO-E-GENERR, Generated ERROR: DDTAB must declare driver STEP=1 or STEP=2
```

Step 2 drivers typically supply a value for the **step** argument of the DPTAB macro. If the step values given the DPTAB and DDTAB macros conflict, the DDTAB macro generates the error:

```
%MACRO-E-GENERR, Generated ERROR: DDTAB STEP=x conflicts with prior declaration.
```

Description

The DDTAB macro creates a driver dispatch table (DDT), using the DRIVER_DATA macro to place it within the driver's data program section (\$\$\$110_DATA). The macro assigns the table a label in the form of **devnam\$DDT**.

The DDTAB macro writes the address of the universal executive routine vector IOC\$RETURN_SUCCESS into routine address fields of the DDT that are not supplied in the macro invocation (with the exception of the **mntver** argument). IOC\$RETURN_SUCCESS places success status in R0 and issues an RSB instruction.

Example

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

This code excerpt uses the DDTAB macro to create a driver dispatch table for the XX device type.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

DDTAB

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- The OpenVMS AXP version of the DDTAB macro does not automatically define the code psect \$\$\$115_DRIVER; rather, it invokes the DRIVER_DATA macro to include the DDT in data psect \$\$\$110_DATA. On OpenVMS AXP systems, you must explicitly invoke the DRIVER_CODE macro to define the \$\$\$115_DRIVER code psect prior to the first line of executable code.
- The number and order of the arguments to the DDTAB macro are different on OpenVMS AXP systems than on OpenVMS VAX systems.
- On OpenVMS AXP systems, you do not distinguish (by use of the plus sign (+)) entry points of OpenVMS routines that are at absolute addresses from entry points at relative locations within the driver. For instance, an OpenVMS AXP device driver could specify the following argument to the DDTAB macro:

```
CANCEL=IOC_STD$CANCELIO,-
```

It is the equivalent of the following argument specification in an OpenVMS VAX device driver:

```
CANCEL=+IOC$CANCELIO,-
```

- An OpenVMS AXP device driver that uses the kernel process services specifies the name of EXE_STD\$KP_STARTIO in **start** argument, and the procedure value of the driver's start-I/O routine in the **kp_startio** argument.
- An OpenVMS AXP device driver that uses the kernel process services indicates the size of the kernel mode stack in the **kp_stack_size**, and specifies a mask of registers to be preserved across kernel process context switches in the **kp_reg_mask** argument.
- Because the procedure value of the controller initialization routine is stored in the DDT (DDT\$PS_CTRLINIT) in OpenVMS AXP systems, you specify its location by using the new **ctrlinit** argument to the DDTAB macro. (On OpenVMS VAX systems, you specify the location of the controller initialization by issuing a DPT_STORE macro to VEC\$L_INITIAL.)
- The OpenVMS AXP version of the DDTAB macro does not provide the **unsolic** argument.
- Although the **channel_assign**, **cancel_selective**, **aux_storage**, and **aux_routine** arguments are allowed in the macro invocation, the functionality they represent has not yet been implemented in OpenVMS AXP systems.
- An OpenVMS AXP terminal port driver cannot share a single DDT with the OpenVMS AXP terminal class driver. The terminal port driver must invoke the DDTAB macro specifying the **ctrlinit** and **unitinit** arguments. The CLASS_UNIT_INIT macro, when invoked by the port driver, initializes the remainder of the port driver's DDT from the class driver's DDT.

DEVICELOCK

Achieves synchronized access to a device's database as appropriate to the processing environment.

Format

DEVICELOCK [lockaddr] [,lockipl] [,savipl] [,condition] [,preserve=YES]

Parameters

[lockaddr]

Address of the device lock to be obtained. If **lockaddr** is not present, DEVICELOCK presumes that R5 contains the address of the UCB and uses the value at UCBSL_DLCK(R5) as the lock address.

[lockipl]

Synchronization IPL. OpenVMS AXP always obtains this IPL from the device lock's data structure and, thus, ignores this argument.

[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. In some instances, setting IPL is undesirable or unnecessary when a driver obtains a device lock. For example, when an interrupt service routine issues the DEVICELOCK macro, the dispatching of the device interrupt has already raised IPL to device 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 DEVICELOCK macro raises IPL to the IPL indicated by the device lock's data structure (if **condition=NOSETIPL** is not specified).

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

- Preserves R0 through the macro call (if **preserve=YES** is specified).
- Stores the address of the device lock in R0.
- 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).

OpenVMS Macros Used by OpenVMS AXP Device Drivers

DEVICELock

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$L_DLCK(R5),- ;Lock device access
    SAVIPL=-(SP),- ;Save current IPL
    PRESERVE=YES ;Save R0
    SETIPL #31 ;Disable all interrupts
    BBC #UCB$V_POWER,- ;If clear - no power failure
    UCB$L_STS(R5),L1 ;...
    ;Service power failure!
.
.
.
DEVICEUNLOCK -
    LOCKADDR=UCB$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
```

This start-I/O routine 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 restores the saved IPL.

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] ,[smp=NO] ,[decode] ,step=0,
      [,idb_crams=0] [,ucb_crams=0] [,bt_order] [,ddt=DDT$BASE]
      [,struc_init=DRIVER$STRUC_INIT] [,struc_reinit=DRIVER$STRUC_REINIT]
      [,psect=$$$105_PROLOGUE] [,dpt=DRIVER$DPT]
```

Parameters

end

Unused in OpenVMS AXP device drivers.

adapter

Type of adapter. You can supply any name that, when appended to the string "AT\$_", results in a symbol defined by the \$DCDEF macro in SYSS\$LIBRARY:STARLET.MLB. Of these symbols, the driver-loading procedure takes special action only when the keyword **NULL** is present. The driver-loading procedure creates no ADP for a null adapter (AT\$_NULL) and clears the VEC\$PS_ADP and IDB\$SL_ADP fields.

[flags=0]

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

DPT\$M_SVP	<p>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_STD\$MOVFRUSER and IOC_STD\$MOVTOUSER.</p> <p>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\$SL_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:</p> $SVA = SEXT((LEFT_SHIFT(ucb\$l_svpn, mmg\$gl_vpn_to_va)) \text{ OR } va\$m_system)$
DPT\$M_NOUNLOAD	<p>Indicates that the driver cannot be reloaded. When this bit is set, the driver can be unloaded only by rebooting the system. Driver unloading and reloading are not supported on OpenVMS AXP systems.</p>

OpenVMS Macros Used by OpenVMS AXP Device Drivers

DPTAB

DPT\$M_SMPMOD	Indicates that the driver has been designed to execute within an OpenVMS multiprocessing environment. Use of any of the 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.
DPT\$M_DECW_ DECODE	Indicates that the driver is a DECwindows class input (decoding) driver
DPT\$M_NO_IDB_ DISPATCH	Tells the driver-loading procedure not to create a list of UCB addresses at the end of the IDB (at IDBSL_UCBLST), regardless of the value of the maxunits argument or the maximum units specified in the /MAX_UNITS qualifier of the System Management (SYSMAN) utility command IO CONNECT.

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.

[unload]

Address of the driver routine invoked by the driver-loading procedure before it unloads an old version of the driver to load a new version.

Note

The OpenVMS AXP operating system does not yet permit driver reloading and does not support driver-unloading routines.

[maxunits=8]

Maximum number of units that this driver supports on a controller. If you omit the **maxunits** argument, the default is eight units. You can override the value specified in the DPT at driver-loading time by using the /MAX_UNITS qualifier to the SYSMAN command IO CONNECT. If DPT\$M_NO_IDB_DISPATCH is not specified in the **flags** argument to the DPTAB macro, these values affect the size of the UCB list the driver-loading procedure generates at the end of the IDB.

[defunits=1]

Maximum number of UCBs to be created by the autoconfiguration facility (one for each device unit to be configured). The unit numbers assigned are zero to **defunits**-1.

If you do not specify the **deliver** argument, the autoconfiguration facility creates the number of units specified by **defunits**. If you specify the address of a unit delivery routine in the **deliver** argument, the autoconfiguration facility 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 autoconfiguration facility should configure automatically. If you omit the **deliver** argument, the autoconfiguration facility creates the number of units specified by the **defunits** argument.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

DPTAB

[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. Because the OpenVMS AXP driver-loading procedure automatically generates a driver name and writes it to the DPT, it effectively ignores this argument.

[smp=NO]

Indication of whether the driver is suitably synchronized to execute in an OpenVMS multiprocessing system. Use of any of the spin lock synchronization macros in a device driver causes the DPTAB macro to indicate multiprocessing synchronization. All OpenVMS AXP drivers must specify **smp=YES**.

[decode]

Address of counted ASCII string that identifies a DECwindows class input (decoding) driver to serial-line switching code.

step

OpenVMS AXP driver step number. You must indicate that a given driver conforms to the coding practices for an Step 1 OpenVMS AXP device driver by supplying **step=2** in the DPTAB macro invocation. If you specify **step=1**, the macro generates the following message:

```
%MACRO-E-GENERR, Generated ERROR: *CAUTION* Step 1 drivers will be obsolete in V2.0
```

If you omit the **step** argument entirely, or specify a value other than 1 or 2, the DPTAB macro generates the message:

```
%MACRO-E-GENERR, Generated ERROR: DPTAB must declare driver STEP=1 or STEP=2
```

Step 2 drivers may also optionally supply a value for the **step** argument of the DDTAB macro. If the step values given the DPTAB and DDTAB macros conflict, the DPTAB macro generates an error of the form:

```
%MACRO-E-GENERR, Generated ERROR: DPTAB STEP=x conflicts with prior declaration.
```

idb_cramps

Number of CRAMS to be allocated and associated with the IDB. The driver-loading procedure allocates the number of CRAMS specified in **idb_cramps** argument to the DPTAB macro and inserts them in the linked list headed by IDB\$PS_CRAM. These CRAMS are therefore available to the driver's controller and unit initialization routine.

ucb_cramps

Number of CRAMS to be allocated and associated with the UCB. The driver-loading procedure allocates the number of CRAMS specified in **ucb_cramps** argument to the DPTAB macro and inserts them in the linked list headed by UCB\$PS_CRAM. These CRAMS are therefore available to the driver's unit initialization routine.

[bt_order]

Ordering number for call to the runtime drivers for boot devices.

[ddt=DDT\$BASE]

Address of DDT. The default is required for all devices not supplied by Digital drivers.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

DPTAB

[struc_init=DRIVER\$STRUC_INIT]

Address of the driver I/O database initialization routine automatically generated by an invocation of the DPT_STORE macro with the **INIT** label. This routine initializes those data structure fields indicated by the invocations of the DPT_STORE macro that follow the DPT_STORE **INIT** and precede the DPT_STORE **REINIT**. The driver-loading procedure calls this initialization routine when it creates the structures and loads the driver, prior to calling the driver's controller and unit initialization routines.

The default value of this argument is required for all Step 2 OpenVMS AXP device drivers.

[struc_reinit=DRIVER\$STRUC_REINIT]

Address of the driver I/O database reinitialization routine automatically generated by an invocation of the DPT_STORE macro with the **REINIT** label. This routine initializes those data structure fields indicated by the invocations of the DPT_STORE macro that follow the DPT_STORE **INIT** and precede the DPT_STORE **END**. The driver-loading procedure calls this reinitialization routine when the driver is first loaded into the system, and whenever the driver is reloaded, prior to calling the driver's controller and unit initialization routines.

The default value of this argument is required for all Step 2 OpenVMS AXP device drivers.

Note that driver unloading and reloading are not supported on OpenVMS AXP systems.

[psect=\$\$105_PROLOGUE]

Program section in which the DPT is created. The default value of this argument is required for all devices not supplied by Digital.

[,dpt=DRIVER\$DPT]

Global symbol for DPT location. The default value of this argument is required for all non-Digital-supplied device drivers.

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

OpenVMS Macros Used by OpenVMS AXP Device Drivers

DPTAB

```

DPTAB   STEP=2,-           ;OpenVMS AXP Step 2 driver
        ADAPTER=CI,-       ;Adapter type
        UCBSIZE=UCB$C_PASIZE,- ;UCB size
        NAME=PNDRIVER,-    ;Driver name
        SMP=YES,-         ;SMP capable
        FLAGS=<DPT$M_SCS!-  ;Driver requires SCS load,
        DPT$M_NOUNLOAD> ; cannot be reloaded

DPT_STORE INIT
DPT_STORE   UCB,UCB$B_FLCK,B,SPL$C_SCS ;SCS spinlock
DPT_STORE   UCB,UCB$L_DEVCHAR,L,<-    ;Device characteristics:
        DEV$M_SHR!-           ; Sharable
        DEV$M_AVL!-           ; Available
        DEV$M_ELG!-           ; Error logging device
        DEV$M_IDV!-           ; Input device
        DEV$M_ODV>           ; Output device
DPT_STORE   UCB,UCB$B_DIPL,B,PN_BR_LEVEL+16 ;Device interrupt IPL
DPT_STORE   UCB,UCB$B_DEVCLASS,B,-    ;Device class =
        DC$BUS                 ; bus
DPT_STORE   UCB,UCB$L_ERTMAX,L,50     ;Retry count is 50 times
DPT_STORE   UCB,UCB$L_ERTCNT,L,50     ; without reboot of system
DPT_STORE REINIT
DPT_STORE   DDB,DDB$L_DDT,D,PN$DDT   ;DDT address
DPT_STORE_ISR CRB$L_INTD,PN$MISC_INTERRUPT ; ISR address
DPT_STORE_ISR CRB$L_INTD+<CRB$S_INTD>,PN$RSP_INTERRUPT
DPT_STORE   END

```

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

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- You must indicate that a given driver conforms to the coding practices for a Step 2 OpenVMS AXP device driver by supplying **step=2** in the **step** argument.
- A driver can request the driver-loading procedure to allocate CRAMs and associate them with the IDB or UCB by specifying the **idb_crams** and **ucb_crams** arguments.
- The OpenVMS AXP driver-loading procedure does not support the reloading of Step 1 OpenVMS AXP device drivers. It therefore ignores the **unload** argument to the DPTAB macro.
- Because the OpenVMS AXP driver-loading procedure automatically generates a driver name and writes it to the DPT, it effectively ignores the **name** argument.
- OpenVMS AXP ignores the **end** argument.
- The DPTAB macro, in conjunction with invocations of the DPT_STORE macro which specify the **INIT**, **REINIT**, and **END** labels, automatically generates driver structure initialization and reinitialization routines, storing their procedure values in the DPT. The default global symbol name for the DPT location is now DRIVERSDPT instead of EVMS\$DRIVER_DPT.

DPT_STORE

In the context of a DPTAB macro invocation, generates driver structure initialization and reinitialization routines which the driver loading and reloading procedures call 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 in the bit insertion operation.

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$!_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 provides a mechanism for a driver to initialize specific data structure fields when the driver is first loaded and when the driver is reloaded. A driver typically contains a series of DPT_STORE invocations which, together, automatically create a driver I/O database initialization routine and a driver I/O database reinitialization routine. The DPTAB macro writes the locations of these routines in the DPT. The driver-loading routine calls the initialization routine when a driver is first loaded; it calls the reinitialization routine both when the driver is first loaded and when the driver is reloaded. Step 2 OpenVMS AXP device drivers cannot be reloaded.

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\$C_IOLOCK8
9	SPL\$C_IOLOCK9
10	SPL\$C_IOLOCK10
11	SPL\$C_IOLOCK11

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 and DPT_STORE_ISR invocations that supply initialization and reinitialization data. The following fields are declared with the DPT_STORE_ISR macro:

CRB\$L_INTD Interrupt service routine
 CRB\$L_INTD2 Interrupt service routine for second interrupt vector

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

OpenVMS Macros Used by OpenVMS AXP Device Drivers

DPT_STORE

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

- Because the OpenVMS AXP driver-loading procedure automatically stores the address of the DDT in the DDB, an OpenVMS AXP device driver does not invoke the DPT_STORE macro to write this address. For instance, the following line should be removed from an existing OpenVMS VAX driver that is to be moved to OpenVMS AXP:

```
DPT_STORE DDB,DDB$DDB_DDT,D,XA$DDT ;Address of DDT
```

- Because the procedure value of the controller and unit initialization routines are stored in the DDT (DDT\$PS_CTRLINIT and DDT\$SL_UNITINIT, respectively) in OpenVMS AXP systems, you specify their location by using the **ctrlinit** and **unitinit** arguments to the DDTAB macro. The following uses of the DPT_STORE macro do not work on OpenVMS AXP systems:

```
DPT_STORE CRB,VEC$SL_INITIAL,D,XA$CTRL_INIT ;Address of controller init routine
DPT_STORE CRB,VEC$SL_UNITINIT,D,XA$UNIT_INIT ;Address of unit init routine
```

- Because the interrupt dispatcher requires the addresses of both the code entry point and the procedure descriptor of an interrupt service routine, you must use the new DPT_STORE_ISR macro (which generates both) to declare the routine. For instance, you should use:

```
DPT_STORE_ISR CRB$SL_INTD, XA_INTERRUPT
;Address of interrupt service routine
```

instead of:

```
DPT_STORE CRB,CRB$SL_INTD+VEC$SL_ISR,D,-
XA_INTERRUPT ;Address of interrupt service routine
```

- The DPTAB macro, in conjunction with invocations of the DPT_STORE macro which specify the **INIT**, **REINIT**, and **END** labels, automatically generates driver structure initialization and reinitialization routines, storing their procedure values in the DPT.
- Be aware that certain data structure fields (such as UCB\$B_FIPL) have been made obsolete in OpenVMS AXP. The names of other fields may have changed, typically to reflect a change in size of the datum.

DPT_STORE_ISR

In the context of a DPTAB macro invocation, generates the addresses of the code entry point and procedure descriptor of an interrupt service routine and stores them in the interrupt transfer vector block (VEC).

Format

```
DPT_STORE_ISR  vec_off ,entry
```

Parameters

vec_off

Symbolic offset to interrupt transfer vector within the CRB. These offsets are of the following form:

Symbolic Offset	Description
CRB\$L_INTD	First interrupt transfer vector
CRB\$L_INTD2	Second interrupt transfer vector
CRB\$L_INTD+<2*VEC\$K_LENGTH>	Third interrupt transfer vector

entry

Procedure value of an interrupt service routine.

Description

The DPT_STORE_ISR macro provides a mechanism for a driver to initialize the VEC\$PS_ISR_PD and VEC\$PS_ISR_CODE fields of an interrupt transfer vector block (VEC) with the addresses of an interrupt service routine's procedure descriptor and code entry point, respectively. Like invocations of the DPT_STORE macro, you invoke the DPT_STORE_ISR macro within the context of the DPTAB macro.

Typically, you use DPT_STORE_ISR within the reinitialization section of the DPT (following DPT_STORE REINIT), so that the VEC fields are initialized at both driver loading and reloading.

Example

```
DPT_STORE_ISR  CRB$L_INTD, XA_INTERRUPT
```

This invocation of the DPT_STORE_ISR macro locates the first interrupt transfer vector associated with the device controller, and places the address of XA_INTERRUPT's procedure descriptor in VEC\$PS_ISR_PD and the address of its code entry point in VEC\$PS_ISR_CODE.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

\$DRIVER_ALTSTART_ENTRY

\$DRIVER_ALTSTART_ENTRY

Format

```
$DRIVER_ALTSTART_ENTRY PRESERVE=<R2,R3,R4,R5>,FETCH=YES
```

```
$OFFDEF ALTARG, < -  
  irp, -  
  ucb >
```

Parameter offsets:

```
MOVL ALTARG$_IRP(AP), R3  
MOVL ALTARG$_UCB(AP), R5
```

\$DRIVER_CANCEL_ENTRY

Format

\$DRIVER_CANCEL_ENTRY PRESERVE=<R2,R3,R4>, FETCH=YES

```
$OFFDEF CANARG, < -  
  chan,-  
  irp,-  
  pcb,-  
  ucb,-  
  reason >
```

Parameter offsets:

```
MOVL CANARG$_CHAN(AP), R2  
MOVL CANARG$_IRP(AP), R3  
MOVL CANARG$_PCB(AP), R4  
MOVL CANARG$_UCB(AP), R5  
MOVL CANARG$_REASON(AP), R8
```

OpenVMS Macros Used by OpenVMS AXP Device Drivers

\$DRIVER_CANCEL_SELECTIVE

\$DRIVER_CANCEL_SELECTIVE

Format

```
$DRIVER_CANCEL_SELECTIVE_ENTRY PRESERVE, FETCH=YES
```

```
$OFFDEF CANSARG, < -  
pcb, -  
ucb, -  
chan, -  
iosb_vector, -  
iosb_count >
```

Parameter offsets:

```
MOVL #SS$_UNSUPPORTED, R0  
MOVL CANSARG$_PCB(AP), R4  
MOVL CANSARG$_UCB(AP), R5  
MOVL CANSARG$_CHAN(AP), R6  
MOVL CANSARG$_IOSB_VECTOR(AP), R7  
MOVL CANSARG$_IOSB_COUNT(AP), R8
```

\$DRIVER_CHANNEL_ASSIGN

Format

```
$DRIVER_CHANNEL_ASSIGN_ENTRY PRESERVE, FETCH=YES
```

```
  $OFFDEF CHANARG, < -  
    ucb, -  
    ccb >
```

Parameter offsets:

```
  MOVL CHANARG$_UCB(AP), R5  
  MOVL CHANARG$_CCB(AP), R8
```

OpenVMS Macros Used by OpenVMS AXP Device Drivers

\$DRIVER_CLONEDUCB

\$DRIVER_CLONEDUCB

Format

```
$DRIVER_CLONEDUCB PRESERVE=R3, FETCH=YES
```

```
$OFFDEF CLONEARG, < -  
  cloned_uch,-  
  ddt,-  
  pcb,-  
  template_uch >  
Parameter offsets:  
          MOVL    #SS$_NORMAL, R0  
MOVL CLONEARG$_CLONED_UCB(AP), R2  
MOVL CLONEARG$_DDT(AP), R3  
MOVL CLONEARG$_PCB(AP), R4  
MOVL CLONEARG$_TEMPLATE_UCB(AP), R5
```

DRIVER_CODE

Declares the program section (psect) that contains driver code.

Format

```
DRIVER_CODE [pname=$$115_DRIVER]
```

Parameters

[pname=\$115_DRIVER]

Name of driver psect that contains driver code. The default psect name, \$115_DRIVER, is suitable for most temporary OpenVMS AXP drivers, although you can specify an alternative name.

Description

The DRIVER_CODE macro generates a psect for driver code, with attributes that allow the Linker utility (linker) to properly and compatibly collect driver image sections into a loadable executive image.

You must precede the first line of executable code in a Step 1 OpenVMS AXP device driver with an invocation of the DRIVER_CODE macro. If the driver consists of multiple source modules, you should replace each explicit setting of the \$\$115_DRIVER psect with an invocation of this macro to ensure that the correct standard psect for driver code sections is always used.

OpenVMS driver macros that construct driver code automatically invoke the DRIVER_CODE macro prior to creating the code. For instance, the DPT_STORE macro automatically invokes the DRIVER_CODE macro prior to constructing the driver initialization and reinitialization routines.

Note

Use of the DRIVER_CODE macro requires that you define the symbol "EVAX".

OpenVMS Macros Used by OpenVMS AXP Device Drivers

\$DRIVER_CRTLINIT

\$DRIVER_CRTLINIT

Format

```
$DRIVER_CTRLINIT_ENTRY PRESERVE=R2, FETCH=YES
```

```
$OFFDEF CTRLARG, < -  
  idb, -  
  ddb, -  
  crb >
```

Parameter offsets:

```
          MOVL    #SS$NORMAL, R0  
MOVL CTRLARG$_IDB(AP), R4  
MOVL CTRLARG$_IDB(AP), R5  
MOVL CTRLARG$_DDB(AP), R6  
MOVL CTRLARG$_CRB(AP), R8
```

\$DRIVER_DELIVER_ENTRY

Format

```
$DRIVER_DELIVER_ENTRY PRESERVE=<R2>, FETCH=YES
```

```
$OFFDEF DLVRARG, < -  
  idb, -  
  unit_number, -  
  scratch_area, -  
  adp >
```

Parameter offsets:

```
MOVL DLVRARG$_IDB(AP), R3  
MOVL DLVRARG$_IDB(AP), R4  
MOVL DLVRARG$_UNIT_NUMBER(AP), R5  
MOVL DLVRARG$_SCRATCH_AREA(AP), R7  
MOVL DLVRARG$_ADP(AP), R8
```

OpenVMS Macros Used by OpenVMS AXP Device Drivers

\$DRIVER_ERRRTN

\$DRIVER_ERRRTN

Format

```
$DRIVER_ERRRTN_ENTRY PRESERVE, FETCH=YES
```

```
$OFFDEF ERRARG, < -  
  irp, -  
  pcb, -  
  ucb, -  
  ccb, -  
  status>
```

```
Parameter offsets:
```

```
MOVL ERRARG$_IRP(AP),R3  
MOVL ERRARG$_PCB(AP),R4  
MOVL ERRARG$_UCB(AP),R5  
MOVL ERRARG$_CCB(AP),R6  
MOVL ERRARG$_STATUS(AP),R0
```

\$DRIVER_FDT_ENTRY

Format

```
$DRIVER_FDT_ENTRY PRESERVE=<R2,R3,R4,R5,R6,R7,R8,R9,R10,R11,R12,  
R13,R14,R15>, FETCH=YES
```

Parameter offsets:

```
MOVL FDTARG$_IRP(AP),R3  
MOVL FDTARG$_PCB(AP),R4  
MOVL FDTARG$_UCB(AP),R5  
MOVL FDTARG$_CCB(AP),R6
```

OpenVMS Macros Used by OpenVMS AXP Device Drivers

\$DRIVER_MNTVER

\$DRIVER_MNTVER

Format

```
$DRIVER_MNTVER_ENTRY PRESERVE, FETCH=YES
```

```
$OFFDEF MNTARG, < -  
  irp, -  
  ucb >
```

Parameter offsets:

```
MOVL MNTARG$_IRP(AP), R3  
MOVL MNTARG$_UCB(AP), R5
```

\$DRIVER_REGDUMP

Format

```
$DRIVER_REGDUMP_ENTRY PRESERVE=<R2>, FETCH=YES
```

```
$OFFDEF REGARG, < -  
  buffer, -  
  cram, -  
  ucb >
```

Parameter offsets:

```
MOVL REGARG$_BUFFER(AP), R0  
MOVL REGARG$_CRAM(AP), R4  
MOVL REGARG$_UCB(AP), R5
```

OpenVMS Macros Used by OpenVMS AXP Device Drivers

\$DRIVER_START_ENTRY

\$DRIVER_START_ENTRY

Format

```
$DRIVER_START_ENTRY PRESERVE=<R2,R4>, FETCH=YES
```

```
  $OFFDEF STARTARG, < -  
    irp, -  
    ucb >
```

Parameter offsets:

```
  MOVL STARTARG$_IRP(AP), R3  
  MOVL STARTARG$_UCb(AP), R5
```

\$DRIVER_UNITINIT

Format

```
$DRIVER_UNITINIT_ENTRY PRESERVE=<R2>, FETCH=YES
```

```
  $OFFDEF UNITARG, < -  
    idb, -  
    ucb >
```

```
Parameter offsets:
```

```
      MOVL    #SS$_NORMAL, R0  
    MOVL UNITARG$_IDB(AP), R4  
    MOVL UNITARG$_UCB(AP), R5
```

DRIVER_DATA

Declares the program section (psect) that contains driver data.

Format

```
DRIVER_DATA [pname=$$110_DATA]
```

Parameters

[pname=\$110_DATA]

Name of driver psect that contains driver data. The default psect name, \$110_DATA, is suitable for most temporary OpenVMS AXP drivers, although you can specify an alternative name.

Description

The DRIVER_DATA macro generates a psect for driver data, with attributes that allow the Linker to properly and compatibly collect driver image sections into a loadable executive image. You must precede any driver data by an invocation of this macro.

OpenVMS driver macros that construct data, such as DDTAB and FUNCTAB, automatically invoke the DRIVER_DATA macro prior to creating the data.

\$FDTARGDEF

Format

\$FDTARGDEF

\$OFFDEF FDTARG, <IRP, PCB, UCB, CCB>

FDT_ACT

Initializes the FDT action routine vector slot corresponding to one or more specified I/O function codes with the procedure value of the specified upper-level FDT action routine.

Format

FDT_ACT action, codes

Parameters

action

Action routine that services the I/O function codes identified by the **codes** argument.

codes

List of codes (enclosed within angle brackets and separated by commas) for I/O functions serviced by the specified upper-level FDT action routine. The macro expansion prefixes each code with the string IO\$_; for example, READVBLK expands to IO\$_READVBLK.

Description

The FDT_ACT macro identifies the upper-level FDT action routine that processes one or more specified I/O function codes. If, at the time it invokes FDT_ACT, the driver has not yet invoked the FDT_INI macro, FDT_ACT invokes it on the driver's behalf, creating an FDT with the label DRIVERSFDT.

An OpenVMS AXP device driver specifies one or more legal I/O functions by supplying the address of an upper-level FDT action routine for that function to the FDT_ACT macro. The FDT_ACT macro initializes the slot in the FDT action routine vector corresponding to each supplied function code with the procedure value of the specified routine.

Multiple invocations of the FDT_ACT macro, in sum, define the full set of I/O functions serviced by the driver. An illegal I/O function is one that the driver does not list in any FDT_ACT macro invocations. Its vector slot contains the procedure value of the illegal I/O function processing routine (EXE\$ILLIOFUNC).

Note, however, only one upper-level FDT action routine can service any given I/O function. If you reuse an I/O function code in an FDT_ACT invocation, the compiler generates an error of the form:

```
%MACRO-E-GENERR, Generated ERROR: Multiple actions defined for function IO$_xxxxxx
```

A consequence of this limitation is that, if the preprocessing of a given function requires that several routines be executed, the upper-level FDT action routine must set up the appropriate call chain.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

FDT_ACT

Example

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

This function decision table (FDT) 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. System I/O preprocessing will call routines `EXE_STD$SETMODE` and `EXE_STD$SENSEMODE` for the device's set-characteristics and sense-mode functions.

FDT_BUF

Builds the buffered function mask within a driver's function decision table (FDT) from the specified list of I/O functions.

Format

FDT_BUF [codes]

Parameters

[codes]

List of codes (enclosed within angle brackets and separated by commas) for I/O functions supported by the driver that require an intermediate system buffer. The macro expansion prefixes each code with the string IOS_; for example, READVBLK expands to IOS_READVBLK.

Description

The FDT_BUF macro builds the buffered function mask within an FDT from the specified list of I/O functions.

An OpenVMS AXP device driver invokes the FDT_BUF macro to indicate which of the I/O functions it supports require a system buffer. If the driver has not yet invoked the FDT_INI macro, FDT_BUF invokes it on the driver's behalf, creating an FDT with the label DRIVER\$FDT.

A driver specifies a legal I/O function by supplying the address of an upper-level FDT action routine for that function to the FDT_ACT macro. Beware of specifying a function code in an FDT_BUF invocation that you do not also specify in an FDT_ACT invocation. The FDT action routine vector slot for such a function contains a pointer to the illegal I/O function processing routine (EXESILLIOFUNC).

An example of the use of FDT_BUF appears in the description of the FDT_ACT macro.

FDT_INI

Creates, labels, and initializes a function decision table (FDT).

Format

```
FDT_INI [fdt=DRIVER$FDT]
```

Parameters

[fdt=DRIVER\$FDT]

Label of the start of the FDT.

Description

The FDT_INI macro creates an FDT, using the DRIVER_DATA macro to place it within the driver's data program section (\$\$\$110_DATA). The macro properly aligns the FDT in memory, assigning it the label specified by the **fdt** argument.

FDT_INI initializes the FDT by clearing the buffered function mask and entering the address of the illegal I/O function processing routine (EXE\$ILLIOFUNC) in all FDT action routine vector slots.

An OpenVMS AXP device driver invokes the FDT_BUF macro to indicate which of the I/O functions it supports require a system buffer. A driver specifies a legal I/O function by supplying the address of an upper-level FDT action routine for that function to the FDT_ACT macro.

An example of the use of FDT_INI appears in the description of the FDT_ACT macro.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

FORK

FORK

Creates a simple fork process on the local processor.

Format

FORK [routine] [,continue] [environment=JSB|CALL]

Parameters

[routine]

Name of the routine to be executed in fork context. If you omit this argument, the FORK macro assumes that the fork routine immediately follows the invocation.

[,continue]

Label where execution continues after the fork block has been inserted on the fork queue. If you omit this argument, control returns to the caller of the routine that invoked the FORK macro.

[environment]

Keyword that specifies the fork routine environment as either JSB or CALL. The default is JSB. If specified as JSB, then EXE\$PRIMITIVE_FORK is called and a .JSB_ENTRY directive is used to generate the fork routine. If specified as CALL, then EXE\$STD\$PRIMITIVE_FORK is called, a .CALL_ENTRY directive is used to generate the fork routine, the FR3, FR4, and FKB parameters in the fork routine are copied into R3, R4, and R5.

Description

The FORK macro creates a fork process. When the FORK macro is invoked, the following registers must contain the values listed:

Register	Contents
R3	Contains the 64-bit value to pass to the fork routine via FKBSQ_FR3(R5)
R4	Contains the 64-bit value to pass to the fork routine via FKBSQ_FR4(R5)
R5	Contains a pointer to the fork block

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

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

Implicit outputs to caller:

ENVIRONMENT=CALL

R0,R1 are scratched.

ENVIRONMENT=JSB

R3,R4 outputs from EXE\$PRIMITIVE_FORK.

R0,R1 are preserved.

Implicit outputs to fork routine, i.e. entry conditions:

ROUTINE=routine_name

If the routine name is specified then the fork entry point is assumed to be at the named location and no fork entry point is defined here. The named fork routine can use either the new standard call interface or the traditional JSB interface as described in section 4.2 regardless of the setting of the ENVIRONMENT keyword.

ROUTINE=<not specified>,ENVIRONMENT=CALL

A fork routine entry point is generated for a routine using the new standard call interface as described in section 4.2.

R3,R4,R5 contain traditional fork routine parameter values copied from the standard call interface actual parameters,

R0,R1 can be scratched.

ROUTINE=<not specified>,ENVIRONMENT=JSB

A fork routine entry point is generated for a routine using the traditional JSB interface as described in section 4.2.

R3,R4,R5 contain traditional fork routine parameters,

R0-R4 can be scratched.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

FORK_ROUTINE

FORK_ROUTINE

Defines the entry point of a fork routine.

Format

```
FORK_ROUTINE [name=fork_routine_name] [,symbol=LOCAL|GLOBAL]
              [,environment=JSB|CALL] [,fetch=YES|NO]
```

Parameters

[name]

Name of the fork routine.

[,symbol]

Specifies if the routine name should be declared as a local or global symbol. The default is for a local symbol.

[,environment]

Specifies the fork routine environment as either JSB or CALL. If specified as JSB, then a .JSB_ENTRY directive is used to define the fork routine entry point. If specified as CALL, then a .CALL_ENTRY directive is used to define the fork routine entry point. The default is JSB.

[,fetch]

Specifies if the fork routine parameters for an ENVIRONMENT=CALL fork routine should be copied into the traditional R3, R4, and R5 register. The default is YES.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

Implicit inputs:

None.

Implicit outputs, i.e. fork routine entry conditions:

ENVIRONMENT=CALL

FORKARG\$_FR3(AP), FORKARG\$_FR4(AP), FORKARG\$_FKB(AP)
the symbolic parameter offsets are defined and
can be used to access the fork routine
parameters,

R3,R4,R5 contain traditional fork routine parameters if
FETCH=YES,

R0,R1 can be scratched.

ENVIRONMENT=JSB

R3,R4,R5 contain traditional fork routine parameters,

R0-R4 can be scratched.

FORK_WAIT

Inserts a fork block on the fork-and-wait queue.

Format

FORK_WAIT [routine] [,continue] [,environment=JSB | CALL]

Parameters

[routine]

Name of the routine to be executed in fork context. If you omit this argument, the FORK_WAIT macro assumes that the fork routine immediately follows the invocation.

[,continue]

Label where execution continues after the fork block has been inserted on the fork-and-wait queue. If you omit this argument, control returns to the caller of the routine that invoked the FORK_WAIT macro.

[,environment]

Specifies the fork routine environment as either JSB or CALL. The default is JSB. If specified as JSB, then EXE\$PRIMITIVE_FORK_WAIT is called and a .JSB_ENTRY directive is used to generate the fork routine. If specified as CALL, then EXE_STD\$PRIMITIVE_FORK_WAIT is called, a .CALL_ENTRY directive is used to generate the fork routine, the FR3, FR4, and FKB parameters in the fork routine are copied into R3, R4, and R5.

Description

The FORK_WAIT macro inserts a fork block on the system fork-and-wait queue. When the FORK_WAIT macro is invoked, the following registers must contain the values listed:

Register	Contents
R3	Contains the 64-bit value to pass to the fork routine via FKBSQ_FR3(R5)
R4	Contains the 64-bit value to pass to the fork routine via FKBSQ_FR4(R5)
R5	Contains a pointer to the fork block

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

Implicit outputs to caller:

ENVIRONMENT=CALL

R0,R1 are scratched.

ENVIRONMENT=JSB

R0,R1 are preserved.

Implicit outputs to fork routine, i.e. entry conditions:

ROUTINE=routine_name

OpenVMS Macros Used by OpenVMS AXP Device Drivers

FORK_WAIT

If the routine name is specified then the fork entry point is assumed to be at the named location and no fork entry point is defined here. The named fork routine can use either the new standard call interface or the traditional JSB interface as described in section 4.2 regardless of the setting of the ENVIRONMENT keyword.

ROUTINE=<not specified>,ENVIRONMENT=CALL

A fork routine entry point is generated for a routine using the new standard call interface as described in section 4.2.

R3,R4,R5 contain traditional fork routine parameter values copied from the standard call interface actual parameters,

R0,R1 can be scratched.

ROUTINE=<not specified>,ENVIRONMENT=JSB

A fork routine entry point is generated for a routine using the traditional JSB interface as described in section 4.2.

R3,R4,R5 contain traditional fork routine parameters,

R0-R4 can be scratched.

FORKLOCK

Achieves synchronized access to a device driver's fork database as appropriate to the processing environment.

Format

FORKLOCK [lock] [,lockipl] [,savipl] [,preserve=YES]

Parameters

[lock]

Index of the fork lock to be obtained. If the **lock** argument is not present in the macro invocation, FORKLOCK presumes that R5 contains the address of the fork block and uses the value at FKB\$B_FLCK(R5) as the lock index.

[lockipl]

Synchronization IPL. OpenVMS AXP obtains this IPL from the spin lock data structure or spin lock IPL vector and ignores this argument.

[savipl]

Location at which to save the current 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 FORKLOCK macro raises IPL to the IPL indicated by the entry in the spin lock IPL vector (SMP\$AL_IPLVEC) that corresponds to the fork lock index.

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 **savi pl** is specified)
- Sets the SMP-modified bit in the driver prologue table (DPT\$V_SMPMOD in DPT\$L_FLAGS)

OpenVMS Macros Used by OpenVMS AXP Device Drivers FORKLOCK

Notes for Converting VAX Drivers

If you are converting an OpenVMS VAX driver to a Step 2 driver, note the following:

- Because OpenVMS AXP obtains this IPL from the spin lock IPL vector or the spin lock data structure that corresponds to the fork lock index, it ignores the **lockipl** argument, if specified.
- Because OpenVMS AXP drivers must use multiprocessing synchronization semantics, the **fipl** argument to the FORKLOCK macro has been removed.

IOFORK

Creates a fork process on the local processor for a device driver, disabling timeouts from the associated device.

Format

```
IOFORK [routine] [,continue] [,ENVIRONMENT=JSB|CALL ]
```

Parameters

[routine]

Name of the routine to be executed in fork context. If you omit this argument, the IOFORK macro assumes that the fork routine immediately follows the invocation.

[,continue]

Label where execution continues after the fork block has been inserted on the fork queue. If you omit this argument, control returns to the caller of the routine that invoked the IOFORK macro.

[,environment]

Keyword that specifies the fork routine environment as either JSB or CALL. The default is JSB. If specified as JSB, then EXE\$PRIMITIVE_FORK is called and a .JSB_ENTRY directive is used to generate the fork routine. If specified as CALL, then EXE_STD\$PRIMITIVE_FORK is called, a .CALL_ENTRY directive is used to generate the fork routine, the FR3, FR4, and FKB parameters in the fork routine are copied into R3, R4, and R5.

Description

The IOFORK macro disables device timeouts by clearing the UCBSV_TIM bit in the field UCBSL_STS and creates a fork process. 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 (64 bits)
R4	Contents to be placed in R4 of the fork process (64 bits)
R5	Address of fork block

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

Implicit outputs to caller:

```
ENVIRONMENT=CALL
```

```
R0,R1      are scratched.
```

```
ENVIRONMENT=JSB
```

```
R3,R4      outputs from EXE$PRIMITIVE_FORK.
```

```
R0,R1      are preserved.
```

Implicit outputs to fork routine, i.e. entry conditions:

```
ROUTINE=routine_name
```

OpenVMS Macros Used by OpenVMS AXP Device Drivers

IOFORK

If the routine name is specified then the fork entry point is assumed to be at the named location and no fork entry point is defined here. The named fork routine can use either the new standard call interface or the traditional JSB interface as described in section 4.2 regardless of the setting of the ENVIRONMENT keyword.

ROUTINE=<not specified>,ENVIRONMENT=CALL

A fork routine entry point is generated for a routine using the new standard call interface as described in section 4.2.

R3,R4,R5 contain traditional fork routine parameter values copied from the standard call interface actual parameters,

R0,R1 can be scratched.

ROUTINE=<not specified>,ENVIRONMENT=JSB

A fork routine entry point is generated for a routine using the traditional JSB interface as described in section 4.2.

R3,R4,R5 contain traditional fork routine parameters,

R0-R4 can be scratched.

IFNORD, IFNOWRT, IFRD, IFWRT

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

Format

$$\left. \begin{array}{l} \text{IFNORD} \\ \text{IFNOWRT} \\ \text{IFRD} \\ \text{IFWRT} \end{array} \right\} \text{ siz ,adr ,dest ,[mode=#0] [,prvmod] [,page] [,page_store]}$$

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.

[prvmod]

Known previous mode of the processor, extracted from the processor status (PS).

[page]

Shifted base address of a known accessible page. The value you specify for the **page** argument can either be zero, or the value returned in the buffer specified as the **page_store** argument in a previous invocation of the macro.

[page_store]

Address of location in which the macro returns the shifted base address of the last page probed.

OpenVMS Macros Used by OpenVMS AXP Device Drivers IFNORD, IFNOWRT, IFRD, IFWRT

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 inline 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 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.

On OpenVMS AXP systems each VAX PROBE instruction generates two PALcode calls—one to read the processor status (PS) to obtain the previous processor mode and one to perform the actual probe. In modules that perform many probes—for instance, code that verifies the accessibility of an item list—these macros provide the following optimizations:

- Because the previous PS does not change in single-threaded kernel mode code, such code can store the previous mode value and reuse it for each probe operation. The **prvmod** argument is available for this purpose.
- Because all of the user's buffers are within the same CPU-specific page, particularly when processing item lists, modules that store the base address of a known accessible page, can compare buffer addresses against this base and avoid any PALcode calls. The **page** and **page_store** arguments are available for this purpose.

When processing an item list, specify the same storage location for both the **page** and **page_store** arguments in each probe macro invocation. This keeps the known accessible page base updated. If the item list does cross a page boundary, the probe operation will be performed only the one item that actually crosses the boundary; subsequent items will share the updated page base value and do not require probing.

When probing a buffer specified by an item descriptor, use the **page** argument with the known probed page, but do not use the **page_store** argument. Other items in the item list are likely to reside within the last known probed page. If the buffer is not, omitting the **page_store** argument allows you to avoid overwriting the last known probed page and issuing an AXP PALcode call when you process subsequent items in the item list.

If you specify zero in the **page** argument as the page base address, these macros skip page base comparison. This is useful in routines that probe a number of input parameters which may or may not be present.

If a routine is probing a number of input parameters which may or may not be present, it should specify a zero in the **page** argument and clear the location pointed to by the **page_store** argument. When the **page** argument is zero, the macros skip page base comparison. In the event an argument is missing, the cleared **page_store** location allows subsequent probe macro invocations to forego checking that location before using its value in the **page** argument.

OpenVMS Macros Used by OpenVMS AXP Device Drivers IFNORD, IFNOWRT, IFRD, IFWRT

CAUTION

These macros expect you to keep known readable pages separate from known writable pages.

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.

Notes for Converting VAX Drivers

If you are converting an OpenVMS VAX driver to a Step 2 driver, note that the OpenVMS AXP versions of these macros provide optional arguments (**prvmod**, **page**, and **page_store**) that allow you to optimize code that performs many probes.

KP_ALLOCATE_KPB

Creates a KPB and a kernel process stack, as required by the Open VMS kernel process services.

Format

```
KP_ALLOCATE_KPB kpb [,stack=#1024] [,flags] [,param]
```

Parameters

kpb

Address of KPB.

[stack=#1024]

Requested size (in bytes) of kernel process stack.

[flags]

Flags indicating the type, size, and configuration of the KPB to be created. KP_ALLOCATE_KPB accepts only the following flags:

KPBSV_VEST	KPB is a VEST KPB. (See Chapter 3 for a description of VEST KPBs.)
KPBSV_SPLOCK	Spin lock area is present. (EXE\$KP_ALLOCATE_KPB automatically sets this bit when KPBSV_VEST is set.)
KPBSV_DEBUG	Debug area is present.
KPBSV_DEALLOC_AT_END	KP_END should call KP_DEALLOCATE.

[param]

Size in bytes of KPB parameter area, if any.

Description

The KP_ALLOCATE_KPB macro calls EXE\$KP_ALLOCATE_KPB to create the KPB and the kernel process stack needed by a kernel process. When a driver invokes KP_ALLOCATE_KPB it cannot be executing above IPL\$_SYNCH or be holding any spin locks that have higher rank than the MMG spin lock.

KP_DEALLOCATE_KPB

Deallocates a KPB and its associated kernel process stack.

Format

KP_DEALLOCATE_KPB kpb

Parameters

kpb
Address of KPB.

Description

The KP_DEALLOCATE_KPB macro calls EXE\$KP_DEALLOCATE_KPB to deallocate the KPB and the associated kernel process stack. When a driver invokes KP_DEALLOCATE_KPB, it cannot be executing above IPL\$ SYNCH or be holding any spin locks of higher rank than MMG.

KP_END

Terminates the execution of a kernel process.

Format

KP_END kpb

Parameters

kpb
Address of KPB.

Description

The KP_END macro calls EXE\$KP_END to terminate the execution of a kernel process and, if KPBSV_DEALLOC_AT_END in KPBSIS_FLAGS is set, to deallocate its KPB. When a driver invokes the KP_END macro, it must be executing at IPL\$RESCHED or above.

KP_RESTART

Resumes the execution of a kernel process.

Format

KP_RESTART kpb

Parameters

kpb
Address of KPB.

Description

The KP_RESTART macro calls EXE\$KP_RESTART to restart a kernel process. The caller of EXE\$KP_RESTART, usually a kernel process scheduling stall routine, must be executing at IPL\$_RESCHED or above.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

KP_REQCOM

KP_REQCOM

Invokes OpenVMS device-independent I/O postprocessing from a kernel process.

Format

KP_REQCOM

Description

The KP_REQCOM macro issues a JSB instruction to IOC\$REQCOM to complete the processing of an I/O request after a kernel process within a driver has finished its portion of I/O postprocessing. (The REQCOM macro cannot be used within the context of a kernel process.)

When the KP_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 KP_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.

KP_STALL_FORK, KP_STALL_IOFORK

Stall a kernel process in such a manner that it can be resumed by the OpenVMS fork dispatcher.

Format

KP_STALL_FORK kpb [,fkb=KPB\$PS_FQFL]

KP_STALL_IOFORK [kpb=IRP\$PS_KPB] [,fkb=UCB\$L_FQFL]

Parameters

kpb

Address of KPB (which must be a VEST KPB). KPB\$PS_UCB must contain the address of a UCB and KPB\$PS_IRP must contain the address of an IRP.

KP_STALL_FORK requires a value for this argument.

[fkb]

Address of a fork block.

Description

The KP_STALL_FORK and KP_STALL_IOFORK macros stall a kernel process by calling EXESKP_FORK.

Prior to calling IOC\$KP_FORK, the KP_STALL_IOFORK macro disable timeouts from the device represented by the UCB associated with the kernel process by clearing UCBSV_TIM in UCB\$L_STS.

The macros can only be called by a kernel process.

KP_STALL_FORK_WAIT

Stalls a kernel process so that it can be resumed by the software timer interrupt service routine's examination of the fork-and-wait queue.

Format

KP_STALL_FORK_WAIT kpb [,fkb]

Parameters

kpb

Address of the caller's KPB.

[fkb]

Address of a fork block. If this argument is omitted, EXESKP_FORK_WAIT uses the fork block within the KPB (KPB\$PS_FKBLK).

Description

The KP_STALL_FORK_WAIT macro stalls a kernel process by calling EXESKP_FORK_WAIT. Only a kernel process executing at or above IPL\$ SYNCH can invoke KP_STALL_FORK_WAIT.

KP_STALL_GENERAL

Stalls the execution of a kernel process.

Format

```
KP_STALL_GENERAL kpb ,stall_routine [,resume_routine]
```

Parameters

kpb

Register containing address of the caller's KPB.

stall_routine

Procedure value of the routine to be called requested to suspend the kernel process described by the specified **kpb**.

A kernel process scheduling stall routine preserves kernel process context not represented on the kernel process stack and takes steps that allow the stalled kernel process thread to be resumed at some later time (for instance, by inserting a fork block on a fork queue or by making a timer queue entry).

At the time a kernel process scheduling stall routine is called, kernel process context has been stored in the KPB and on the kernel process stack. The stall routine can thus immediately resume the kernel process thread.

[resume_routine]

Procedure value of the routine to be invoked by EXESKP_RESTART when a stalled kernel process is to be resumed.

Description

The KP_STALL_GENERAL macro calls EXESKP_STALL_GENERAL to suspend execution of the current kernel process. A kernel process invokes KP_STALL_GENERAL directly — instead of KP_STALL_FORK, KP_STALL_FORK_WAIT, KP_STALL_IOFORK, KP_STALL_REQCHAN, KP_STALL_WFIKPCH, or KP_STALL_WFIRLCH — when it requires a specialized scheduling stall routine or scheduling restart routine.

Only a kernel process can invoke the KP_STALL_GENERAL macro.

KP_STALL_REQCHAN

Stalls a kernel process in such a manner that it can be resumed by the granting of a device controller channel.

Format

KP_STALL_REQCHAN [pri=LOW] [,kpb=IRP\$PS_KPB] [,idb=YES]

Parameters

[pri=LOW]

Priority of the request for the controller channel. You can specify one of the following keywords:

Keyword	Meaning
LOW	Insert fork block of UCB requesting controller channel at the tail of the channel-wait queue.
HIGH	Insert fork block of UCB requesting controller channel at the head of the channel-wait queue.

[kpb=IRP\$PS_KPB]

Address of the caller's KPB (which must be a VEST KPB). KPB\$PS_UCB must contain the address of a UCB and KPB\$PS_IRP must contain the address of an IRP.

[idb=YES]

Flag requesting the return of the IDB address in R4. **idb=YES**, the default, assuming that the address of the UCB is in R5 at the time the macro is invoked, causes the address of the IDB to be placed in R4 after the channel request has been granted.

R4. If your driver does not require KP_STALL_REQCHAN to emulate the IDB returned in R4 behavior of REQCHAN (or REQPCCHAN), you can save two inline MACRO-32 instructions by adding IDB=NO to the KP_STALL_REQCHAN invocation.

Description

The KP_STALL_REQCHAN macro calls IOCSKP_REQCHAN to request ownership of the controller channel. If the channel is not busy, the kernel process acquires the channel immediately and does not stall. If the channel is busy, the kernel process is placed in the channel-wait-queue to be later resumed by IOCSRELCHAN when it grants the channel request.

Only a kernel process executing at fork IPL and holding the appropriate fork lock can invoke the KP_STALL_REQCHAN macro.

KP_STALL_WFIKPCH, KP_STALL_WFIRLCH

Stall a kernel process in such a manner that it can be resumed by device interrupt processing.

Format

KP_STALL_WFIKPCH `excpt ,time=65536 [,newipl=(SP)+] [,kpb=IRP$PS_KPB]`

KP_STALL_WFIRLCH `excpt ,time=65536 [,newipl=(SP)+] [,kpb=IRP$PS_KPB]`

Parameters

excpt

Label of the timeout handling code. When the **excpt** argument is present, the macro expands to use a BLBC to transfer to that routine in the event that SSS_TIMEOUT status is returned. A driver writer may choose to omit the **excpt** argument and decode the R0 status directly.

[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.

[newipl=(SP)+]

IPL to which to lower before returning to caller. Typically this is the fork IPL associated with device processing that was pushed on the stack by a prior invocation of the DEVICELock macro.

[kpb=IRP\$PS_KPB]

Address of the caller's KPB (which must be a VEST KPB). KPB\$PS_UCB must contain the address of a UCB and KPB\$PS_IRP must contain the address of an IRP.

Description

The KP_STALL_WFIKPCH and KP_STALL_WFIRLCH macros call IOC\$KP_WFIKPCH and IOC\$KP_WFIRLCH respectively to initiate a stall of the kernel process: These macros can only be invoked by a kernel process.

When invoked, KP_STALL_WFIKPCH or KP_STALL_WFIRLCH assumes that the local processor has obtained the appropriate synchronization with the device database by securing the appropriate device lock, as recorded in the unit control block (UCB\$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.

KP_START

Starts the execution of a kernel process.

Format

```
KP_START kpb ,routine [,registers]
```

Parameters

kpb

Address of KPB.

routine

Procedure value of the routine to be started as the top-level routine in the kernel process.

[registers]

Optional register save mask, indicating which registers must be preserved across kernel process context switches. Registers R0, R1, R16 through R25, R27, R28, R30, and R31 are never preserved across context switches; a **reg-mask** that indicates any of these registers is illegal. Registers R12 through R15, R26, and R29 are always saved and need not be specified.

Description

The KP_START macro calls EXE\$KP_START to create a kernel process and start its execution.

When invoking the KP_START macro, code must be executing at IPL\$_RESCHED or above.

KP_SWITCH_TO_KP_STACK

Switch to kernel process context.

Format

```
KP_SWITCH_TO_KP_STACK [kpb=R6] [,return=RSB]  
                        [,registers=<R0,R1,R2,R3,R4,R5,R6>]
```

Parameters

[kpb=R6]

Address of KPB.

[return=RSB]

Return semantic to be used when the kernel process stalls or completes. Valid keywords are **RSB** and **RET**.

[,registers=<R0,R1,R2,R3,R4,R5,R6>]

Register save mask, indicating which registers are to be preserved across context switches between the kernel process and the main thread.

Description

The KP_SWITCH_TO_KP_STACK macro creates a kernel process by calling EXESKP_START, supplying a routine embedded in the macro as the top-level kernel process routine. Execution proceeds in kernel process context, with the address of the KPB in the location indicated by the **kpb** parameter and the kernel process stack active.

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]

Synchronization IPL. OpenVMS AXP obtains this IPL from the spin lock data structure corresponding to the **lockname** and, thus, ignores this argument.

[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 indicated by the entry in the spin lock IPL vector (SMP\$AL_IPLVEC) that corresponds to the spin lock index SPL\$C_**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).

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)

Notes for Converting VAX Drivers

If you are converting an OpenVMS VAX driver to a Step 2 driver, note that because OpenVMS AXP obtains the synchronization IPL from the spin lock data structure corresponding to the **lockname**, it ignores the **lockipl** argument, if specified.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

RELCHAN

RELCHAN

Releases all controller data channels allocated to a device.

Format

RELCHAN

Description

The RELCHAN macro releases 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 R1.

REQCHAN

Requests exclusive use of the CRB and defines the channel grant routine entry point.

Format

```
REQCHAN [pri=LOW | HIGH] [,ENVIRONMENT=JSB | CALL]
```

Parameters

[pri=LOW]

Priority of request. If the priority is **HIGH**, REQCHAN calls IOC_STDS\$PRIMITIVE_REQCHANH; otherwise it calls IOC_STDS\$PRIMITIVE_REQCHANL.

[,environment]

Specifies the callers and grant routine environments as either JSB or CALL. The default is JSB. If specified as JSB, then an RSB is used to return from the current routine if the channel is not granted immediately and a .JSB_ENTRY directive is used to generate the grant routine. If specified as CALL, then an RET is used to return from the current routine if the channel is not granted immediately, a .CALL_ENTRY directive is used to generate the grant routine, and the grant routine parameters are copied into R3, R4, and R5.

Description

The REQCHAN macro obtains a controller's data channel.

If the channel is granted immediately, execution continues at the line of code that immediately follows the macro invocation. If no channel is available, the UCB is placed in a channel-wait queue, and the macro returns control to its caller's caller. When the channel request is granted, execution resumes at the line of code following the macro execution.

When the REQCHAN macro is invoked, R5 must contain the address of the UCB.

The REQCHAN macro returns the address of the IDB in R4 and destroys the contents of R0 through R2.

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

Implicit inputs:

R3 contains a pointer to the IRP which if necessary is passed to the grant routine via UCB\$Q_FR3(R5),
R5 contains a pointer to the UCB,

Implicit outputs to caller:

R4 contains the IDB address,
R0-R2 are scratched.

Implicit outputs to grant routine, that is, entry conditions:

ENVIRONMENT=CALL

OpenVMS Macros Used by OpenVMS AXP Device Drivers

REQCHAN

A driver channel grant routine entry point is generated for a routine using the new standard call interface as described in section 4.3.

R3,R4,R5 contain traditional channel grant routine parameter values copied from the standard call interface actual parameters,

R0,R1 can be scratched.

ENVIRONMENT=JSB

A driver channel grant routine entry point is generated for a routine using the traditional JSB interface.

R3,R4,R5 contain traditional channel grant routine parameters,

R0-R5 can be scratched.

REQCOM

Places the current IRP on the post processing queue and to logically end the driver fork thread that began on entry into the start I/O or alternate start I/O routines.

Format

REQCOM [,environment=JSB | CALL]

Parameters

[,environment]

Specifies the fork routine environment as either JSB or CALL. The default is JSB. If specified as JSB, then an RSB is used to return from the current routine. If specified as CALL, then an RET is used to return from the current routine.

Description

The REQCOM macro completes 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 R1. All other registers are also destroyed if the action of the macro initiates the processing of a waiting I/O request for the device.

REQPCHAN

Obtains a controller's data channel.

Format

REQPCHAN [pri=LOW] [,environment=JSB | CALL]

Parameters

[pri=LOW]

Priority of request. If the priority is **HIGH**, REQPCHAN calls IOC\$PRIMITIVE_REQCHANH; otherwise it calls IOC\$PRIMITIVE_REQCHANL.

[,environment=JSB | CALL]

Specifies the fork routine environment as either JSB or CALL. The default is JSB. If specified as JSB, then an RSB is used to return from the current routine. If specified as CALL, then an RET is used to return from the current routine.

Description

The REQPCHAN macro calls obtains a controller's data channel.

If the channel is granted immediately, execution continues at the line of code that immediately follows the macro invocation. If no channel is available, the UCB is placed in a channel-wait queue, and the macro returns control to its caller's caller. When the channel request is granted, execution resumes at the line of code following the macro execution.

When the REQPCHAN macro is invoked, R5 must contain the address of the UCB.

The REQPCHAN macro returns the address of the IDB in R4 and destroys the contents of R0 through R2.

SYSDISP

Causes a branch to a specified address according to the type of AXP system executing the code in the macro expansion.

Format

SYSDISP list [,continue=YES]

Parameters

list

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

<**system-type**, **destination**>

The **system-type** parameter identifies the type of AXP system for which the macro is to generate a case table entry. The SYSDISP macro identifies the following AXP systems:

ADU	Prototype AXP system
DEC 4000-600	AXP deskside system
LASER	AXP mid-range system
DEC 3000-300	AXP workstation
MANNEQUIN	AXP simulator

[continue=YES]

Specifies whether execution should continue at the line immediately after the SYSDISP macro if the value at EXESGQ_SYSTYPE does not correspond to any of the values specified as the **system-type** in the **list** argument. A fatal bugcheck of UNSUPRTCPU occurs if the dispatching code does not find the executing system identified in the **list** and the value of **continue** is NO.

Description

The SYSDISP macro provides a means for transferring control to a specified destination depending on the type of the executing system.

SYSDISP constructs appropriate symbolic constants for each **system-type** listed in **list**, and compares them against the contents of EXESGQ_SYSTYPE. These constants have the form HWRPBS_SYSTYPE\$K_*system-type*.

TBI_ALL

Invalidates the data and instruction translation buffers in their entirety.

Format

TBI_ALL [environ=MP]

Parameters

[environ=MP]

Context of translation buffer invalidation. When **environ=LOCAL**, the macro invalidates the translation buffer only in the context of the local processor. When **environment** is not specified or does not equal **LOCAL**, the macro extends to all system components (that is, processors and device controllers) that may have cached PTEs.

Description

The TBI_ALL macro flushes the entire contents of the data and instruction translation buffers.

The AXP architecture specifies that whenever a PTE is modified in a way that results in a reduction of access to a virtual address, software must ensure that any cached copies of the previous PTE contents are flushed from the translation buffer before the new PTE contents can be accessed. For example, code must invalidate a translation buffer cache entry if it clears the valid bit of the associated PTE, increases its page protection, or sets one of its memory management fault bits.

If the fault-on-execute bit in the modified PTE is set, the page was never used in execution. For such pages, code can achieve better performance by avoiding an instruction translation buffer flush and invoke the TBI_DATA_64 macro to flush only the data translation buffer.

TBI_DATA_64

Invalidates a single 64-bit virtual address in the data translation buffer.

Format

```
TBI_DATA_64  addr [,environ=MP]
```

Parameters

addr

64-bit virtual address described by the translation buffer entry to be invalidated.

The TBI_DATA_64 macro assumes that the virtual address supplied in the **addr** argument will normally be in a register. Although it also accepts a memory address, you should quadword align it to avoid the performance degradation caused by the servicing of an alignment fault.

[environ=MP]

Context of translation buffer invalidation. When **environ=LOCAL**, the macro invalidates the translation buffer only in the context of the local processor. When **environment** is not specified or does not equal **LOCAL**, the macro extends to all system components (that is, processors and device controllers) that may have cached PTEs.

Description

As specified by the AXP architecture, the TBI_DATA_64 macro invalidates a single 64-bit virtual address in the data translation buffer only. The instruction translation buffer is not affected.

The AXP architecture specifies that whenever a PTE is modified in a way that results in a reduction of access to a virtual address, software must ensure that any cached copies of the previous PTE contents are flushed from the translation buffer before the new PTE contents can be accessed. For example, code must invalidate a translation buffer cache entry if it clears the valid bit of the associated PTE, increases its page protection, or sets one of its memory management fault bits.

If R2 is not specified as the **addr** argument, it is preserved across the macro call.

The TBI_DATA_64 macro and its callers depend on the MTPR instruction to save and restore the registers it destroys. If a MACRO compiler built-in is ever used again, those registers must be specifically preserved.

TBI_SINGLE

Flushes the cached contents of a single page-table entry (PTE) from the data and instruction translation buffers.

Format

```
TBI_SINGLE  addr [,environ=MP]
```

Parameters

addr

32-bit virtual address to be invalidated.

[environ=MP]

Context of translation buffer invalidation. When **environ=LOCAL**, the macro invalidates the translation buffer only in the context of the local processor. When **environment** is not specified or does not equal **LOCAL**, the macro extends to all system components (that is, processors and device controllers) that may have cached PTEs.

Description

As specified by the AXP architecture, the TBI_SINGLE macro flushes the cached contents of a single page-table entry (PTE) from both the data and instruction translation buffers.

The AXP architecture specifies that whenever a PTE is modified in a way that results in a reduction of access to a virtual address, software must ensure that any cached copies of the previous PTE contents are flushed from the translation buffer before the new PTE contents can be accessed. For example, code must invalidate a translation buffer cache entry if it clears the valid bit of the associated PTE, increases its page protection, or sets one of its memory management fault bits.

If the fault-on-execute bit in the modified PTE is set, the page was never used in execution. For such pages, code can achieve better performance by avoiding an instruction translation buffer flush and invoke the TBI_DATA_64 macro to flush only the data translation buffer.

TBI_SINGLE_64

Invalidates a single 64-bit virtual address in both the data and instruction translation buffers.

Format

TBI_SINGLE_64 *addr* [,*environ*=MP]

Parameters

addr

64-bit virtual address to be invalidated.

The TBI_SINGLE_64 macro assumes that the virtual address supplied in the **addr** argument will normally be in a register. Although the TBI_DATA_64 macro also accepts a memory address, you should quadword align it to avoid the performance degradation caused by the servicing of an alignment fault.

[environ=MP]

Context of translation buffer invalidation. When **environ=LOCAL**, the macro invalidates the translation buffer only in the context of the local processor. When **environment** is not specified or does not equal **LOCAL**, the macro extends to all system components (that is, processors and device controllers) that may have cached PTEs.

Description

As specified by the AXP architecture, the TBI_SINGLE_64 macro invalidates a single 64-bit virtual address in both the data and instruction translation buffers.

The AXP architecture specifies that whenever a PTE is modified in a way that results in a reduction of access to a virtual address, software must ensure that any cached copies of the previous PTE contents are flushed from the translation buffer before the new PTE contents can be accessed. For example, code must invalidate a translation buffer cache entry if it clears the valid bit of the associated PTE, increases its page protection, or sets one of its memory management fault bits.

If R2 is not specified as the **addr** argument, it is preserved across the macro call.

TIMEDWAIT

Waits a specified interval of time for an event or condition to occur, optionally executing a series of specified instructions that test for various exit conditions.

Format

```
TIMEDWAIT  time [,ins1] [,ins2] [,ins3] [,ins4] [,ins5] [,ins6] [,donelbl] [,imbedlbl]  
           [,ublbl] [,nsec] [,bus] [,userins]
```

Parameters

time

Delay time specified in 10-microsecond intervals. Actual delay time depends on number and type of loop instructions, clock frequency, and other variables.

Note that the **time** and **nsec** arguments are mutually exclusive.

[ins1]

First instruction to be executed in the delay loop.

[ins2]

Second instruction to be executed in the delay loop.

[ins3]

Third instruction to be executed in the delay loop.

[ins4]

Fourth instruction to be executed in the delay loop.

[ins5]

Fifth instruction to be executed in the delay loop.

[ins6]

Sixth instruction to be executed in the delay loop.

[donelbl]

Label placed after the instruction at the end of the TIMEDWAIT loop; embedded instructions can pass control to this label in order to pass control to the instruction following the invocation of the TIMEDWAIT macro.

[imbedlbl]

Label placed at the first of the embedded instructions; after executing a processor-specific delay, the TIMEDWAIT macro passes control here to retest for the condition.

[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.

[nsec]

Delay time in nanoseconds. Actual delay time depends on number and type of loop instructions, clock frequency, and other variables.

Note that the **nsec** and **time** arguments are mutually exclusive.

OpenVMS Macros Used by OpenVMS AXP Device Drivers

TIMEDWAIT

[bus]

Address of ADP of the bus, if a bus-specific delay should be added to the delay loop. You would add a bus-specific delay, for instance, to avoid saturating a bus with CSR references in the instruction loop.

[userins]

Additional instructions to be executed in the delay loop. This list can be of indefinite length and is executed after (or in place of) the instructions specified in the **ins1** through **ins6** arguments.

Description

The TIMEDWAIT macro provides the ability to write code that is based on specific time intervals and is independent of system-specific timer implementations. You can use the TIMEDWAIT macro for the following tasks:

- **Timeout handling.** The macro generates a time delay in which a number of instructions tests for the occurrence of a specific event or condition. In this case, either the specified time delay is completed or an exit condition is met. A device driver uses this mechanism to establish time bounds for the execution of a given instruction sequence.
- **Simple delay.** The macro generates a time delay with no embedded instructions. When the delay has completed, the code thread continues.
- **Optimistic polling.** If a device is known to respond quickly, a driver might invoke the TIMEDWAIT macro with a short time delay to check for device completion and potentially avoid the overhead of device interrupt servicing. If the instructions supplied to the delay loop determine that the operation has completed, an interrupt has been saved. Otherwise, the delay completes and the suspended code thread continues.

The TIMEDWAIT macro returns a status code (SS\$_NORMAL or SS\$_TIMEOUT) in R0. Note that the embedded instructions can overwrite SS\$_NORMAL status, although SS\$_TIMEOUT status cannot be overwritten. The macro destroys the contents of R1, and preserves all other registers.

Examples

1.

```
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
```
2.

```
TIMEDWAIT NSEC=#<100*1000>,-
          DONELBL=10$,-               ;Label to exit wait loop
          USERINS=<<BITB #1,CSR(R4)>,<BNEQ 10$>> ;Check CSR
```

OpenVMS Macros Used by OpenVMS AXP Device Drivers

TIMEDWAIT

Notes for Converting VAX Drivers

If you are converting an OpenVMS VAX driver to a Step 2 driver, note the following:

- The OpenVMS AXP TIMEDWAIT macro, unlike the OpenVMS VAX version, *does* read a processor register. As such, the interval it waits corresponds very closely with the delay specified in the **time** or **nsec** argument, and is not affected by the number or complexity of the imbedded instructions that may be specified as arguments.
- The OpenVMS AXP TIMEDWAIT macro does not automatically adjust the **time** (or **nsec**) argument to accommodate a bus-specific or CPU-specific delay factor. If a bus-specific delay is needed, you can request one by specifying the address of the bus's ADP in the **bus**.
- The OpenVMS AXP TIMEDWAIT macro allows you to specify the delay time in either 10-microsecond intervals (using the **time** argument) or in nanosecond units (using the **nsec** argument).
- The OpenVMS AXP TIMEDWAIT macro allows you to specify, in the **userins** argument, instructions to execute within the delay loop and test for exit conditions. These instructions execute within the loop after (or in place of) the instructions specified in the **ins1** through **ins6** arguments.

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 }  excpt [,time=65536] [,newipl] [environment=JSB | CALL] [,toutrout]
```

Parameters

excpt

Label of the timeout handling code within the driver.

[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 1 second.

[newipl=(SP)+]

IPL to which to lower before returning to caller. Typically this is the fork IPL associated with device processing that was pushed on the stack by a prior invocation of the DEVICELock macro.

[,environment]

Specifies the current and interrupt resume routine environments are either JSB or CALL. The default is JSB. If specified as JSB, then the return from the current procedure is via a RSB, and a .JSB_ENTRY directive is used to generate the resume routine entry point. If specified as CALL, then the return from the current procedure is via a RET, a .CALL_ENTRY directive is used to generate the resume routine entry point, and the IRP, FR4, and UCB parameters in the resume routine are copied into R3, R4, and R5.

[,toutrout]

Specifies the timeout routine entry point. The timeout routine is a fork routine that can either use the traditional or the new standard call interface. If not specified then the resume routine entry point is also used as the timeout routine entry point. The timeout routine procedure value is loaded into UCB\$PS_TOUTROUT(R5) for potential use by EXESTIMEOUT. This parameter cannot be specified together with the EXCPT parameter.

Description

The WFIKPCH and WFIRLCH macros construct an inline entry point for the code that follows the macro invocation (normally called the fork routine). They insert an instruction at the beginning of the fork routine that tests UCB\$V_TIMEOUT in UCB\$L_STS and branches to the label of the timeout code (specified in the **excpt** argument) if it is set.

OpenVMS Macros Used by OpenVMS AXP Device Drivers WFIKPCH, WFIRLCH

Finally, WFIKPCH and WFIRLCH place the procedure value of the fork routine (at the instruction following the macro invocation) in UCB\$L_FPC, insert the **time** value in R1 and **newipl** value in R2, and call the appropriate wait-for-interrupt routine (either IOC\$PRIMITIVE_WFIKPCH or IOC\$PRIMITIVE_WFIRLCH).

When the wait-for-interrupt routine returns control, the WFIKPCH or WFIRLCH macro issues an RSB instruction to the caller of the routine which invoked it (that is, the caller of the start-I/O routine).

Either the device interrupt servicing routine or the software timer interrupt servicing routine will eventually issue a JSB instruction to the fork routine. In both instances, code can assume that only R3 and R4 have been preserved across the suspension.

IOC\$WFIKPCH and IOC\$WFIRLCH assume that, prior to the invocation of the macro, a DEVICELOCK macro has been issued to synchronize with other device activity.

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 (if the newipl argument is not specified)

Notes for Converting Step 1 Drivers

If you are converting a Step 1 driver to a Step 2 driver, note the following:

Implicit inputs:

R3 contains a pointer to the IRP which is passed to the interrupt resume routine via UCB\$Q_FR3(R5),

R4 contains the 64-bit value to pass to the interrupt resume routine via UCB\$Q_FR4(R5),

R5 contains a pointer to the UCB,

Implicit outputs to caller:

R0,R1,R2 are scratched.

Implicit outputs to resume routine, i.e. entry conditions:

ENVIRONMENT=CALL

An entry point is generated that conforms to both the new standard call interface for a driver resume from interrupt routine as described in section 4.7 and the new standard call interface for a fork routine as described in section 4.2.

R3,R4,R5 contain traditional resume from interrupt routine parameter values (64-bit value for R4) copied from the standard call interface actual parameters,

R0,R1 can be scratched.

ENVIRONMENT=JSB

OpenVMS Macros Used by OpenVMS AXP Device Drivers WFIKPCH, WFIRLCH

An entry point is generated that conforms to both the traditional JSB interface for a driver resume from interrupt routine and the traditional JSB interface for a fork routine..

R3,R4,R5 contain traditional resume from interrupt routine parameters (64-bit value for R4),

R0-R4 can be scratched.

C Driver Macros

This chapter describes the C macros that support the simple fork mechanism and are used by device drivers on the start I/O code path. These macros are all defined in the `vms_drivers.h` header file. Any definitions of data types, function prototypes, and other macros that are needed by these macros are also included by the `vms_drivers.h` file.

OpenVMS C Macros Used by OpenVMS AXP Device Drivers

DEVICE_LOCK

DEVICE_LOCK

Use to acquire a device spinlock and to optionally save the original IPL.

Format

DEVICE_LOCK lockaddr, raise_ipl, savipl_p

Description

Inputs:

lockaddr is a pointer to the device spinlock structure of type SPL.

raise_ipl is either the integer value RAISE_IPL or NORAISE_IPL. The symbol RAISE_IPL is defined to be 1 and the symbol NORAISE_IPL is defined to be 0 by the vms_drivers.h file. If raise_ipl is equal to RAISE_IPL, then the IPL is set using the value in the spinlock structure.

Outputs:

savipl_p is a 32-bit integer passed by reference in which the original IPL is returned. If the address of this parameter is NOSAVE_IPL, then the current IPL is not returned. The symbol NOSAVE_IPL is defined to be a null pointer, i.e. ((int *) 0), by the vms_drivers.h file.

For example, in a driver fork routine one could save the current IPL and take the device lock by:

```
device_lock(ucb->ucb$l_dlck, RAISE_IPL, &orig_ipl);
```

However, in an interrupt service routine, where the IPL is already known to be at device IPL one could take the device lock and leave the IPL unchanged by:

```
device_lock(ucb->ucb$l_dlck, NORAISE_IPL, NOSAVE_IPL);
```

The definition of the device_lock macro is:

```
device_lock(lockaddr, raise_ipl, savipl_p)
extern SMP smp$gl_flags;
if(savipl_p != NOSAVE_IPL)
    *savipl_p = __PAL_MFPR_IPL();
if(raise_ipl == NORAISE_IPL) {
    if(smp$gl_flags.smp$sv_enabled)
        smp_std$acqnoipl( lockaddr );
    }
    else {
        if(smp$gl_flags.smp$sv_enabled)
            smp_std$acquirel( lockaddr );
        else
            __PAL_MTPR_IPL( lockaddr->spl$l_ipl );
    }
}
```

OpenVMS C Macros Used by OpenVMS AXP Device Drivers DEVICE_LOCK

Notes:

1. The use of either the RAISE_IPL or NORaise_IPL constant for the raise_ipl parameter allows the C compiler to generate code for each case exclusively without the need for a runtime test.
2. Caution should be exercised not to use the constant NOLOWER_IPL for the raise_ipl parameter. If NOLOWER_IPL is used erroneously the effect is that IPL will be raised. However, the effect of using values other than RAISE_IPL or NORaise_IPL for the raise_ipl parameter should be considered as unpredictable. Remember that IPL is raised (or held) when acquiring a spin lock and lowered (or held) on release.
3. If the savipl_p parameter is specified as NOSAVE_IPL, then all the code for handling the savipl_p parameter may be optimized away by the C compiler. This convention is similar to that used by the fork_lock and sys_lock macros.
4. Compile time type checking will detect an erroneous use of an ordinary constant other than NOSAVE_IPL for the savipl_p parameter.
5. This macro is similar to the MACRO-32 DEVICELOCK macro. However, the C version does not have a parameter corresponding to the optional LOCKIPL parameter. In DEVICELOCK, if the runtime check finds SMP\$V_ENABLED set, then the new IPL is always obtained from the spin lock structure. However, if SMP\$V_ENABLED is clear and LOCKIPL was specified at compiletime, then the LOCKIPL value is used.

OpenVMS C Macros Used by OpenVMS AXP Device Drivers

DEVICE_UNLOCK

DEVICE_UNLOCK

Use to either release or restore (that is, conditionally release) a device spinlock and to optionally set a new IPL.

Format

DEVICE_UNLOCK lockaddr, newipl, restore

Description

Inputs:

lockaddr is a pointer to the device spinlock structure of type SPL.

newipl is the integer value of the desired new IPL or the value NOLOWER_IPL if the IPL should be left unchanged. The symbol NOLOWER_IPL is defined to be -1 by the vms_drivers.h file.

restore is the either the integer value SMP_RESTORE or SMP_RELEASE. If SMP_RELEASE is specified then the spinlock is unconditionally released by calling SMP_STD\$RELEASEL, otherwise the spinlock is conditionally released by calling SMP_STD\$RESTOREL. The symbol SMP_RESTORE is defined to be 1 and the symbol SMP_RELEASE is defined to be 0 by the vms_drivers.h file.

For example, in a driver fork routine one could release the device lock and restore the previously saved IPL by:

```
device_unlock (ucb->ucb$l_dlck, orig_ipl, SMP_RELEASE);
```

However, in an interrupt service routine, one could release the device lock and stay at the current IPL by:

```
device_unlock (ucb->ucb$l_dlck, NOLOWER_IPL, SMP_RELEASE);
```

The definition of the device_unlock macro is:

```
device_unlock(lockaddr, newipl, restore)
{
    extern SMP smp$gl_flags;
    if(smp$gl_flags.smp$v_enabled) {
        if(restore == SMP_RELEASE)
            smp_std$releasel( lockaddr );
        else
            smp_std$restorel( lockaddr );
    }
    if(newipl >= 0)
        __PAL_MTPR_IPL( newipl );
}
```

Notes:

1. If the newipl parameter is specified as NOLOWER_IPL, then all the code for handling the newipl parameter may be optimized away by the C compiler. This convention is similar to that used by the fork_unlock and sys_unlock macros.

OpenVMS C Macros Used by OpenVMS AXP Device Drivers DEVICE_UNLOCK

2. Using either the `SMP_RELEASE` or `SMP_RESTORE` constant for the restore parameter allows the C compiler to generate code for either a call to `SMP_STD$RELEASEL` or `SMP_STD$RESTOREL` without the need for a runtime test. This convention is similar to that used by the `fork_unlock` and `sys_unlock` macros.
3. Caution should be exercised not to use the constants `RAISE_IPL` nor `NORAISE_IPL` instead of `NOLOWER_IPL` for the `newipl` parameter. If either of these are used erroneously in place of `NOLOWER_IPL` the effect is that IPL will be set to either 0 or 1. Remember that IPL is raised (or held) when acquiring a spin lock and lowered (or held) on release.

FORK

Use to queue a specified fork routine with specified fork routine parameters. After the fork routine is queued, execution continues with the next statement following the fork macro.

Format

FORK fork_routine, fr3, fr4, fkb

Description

Inputs:

fork_routine is the procedure value of the routine that is to be executed in a fork thread. This value is passed to the fork dispatcher via fkb->fkb\$l_fpc.

fr3 is the 64-bit value to pass to the fork routine via fkb->fkb\$q_fr3. This parameter is cast as a 64-bit integer.

fr4 is the 64-bit value to pass to the fork routine via fkb->fkb\$q_fr4. This parameter is cast as a 64-bit integer.

fkb is a pointer to the fork block. This parameter is cast as a pointer to an FKB.

The definition of the fork macro is:

```
fork(fork_routine, fr3, fr4, fkb)
{
  ((FKB *) fkb)->fkb$l_fpc = fork_routine;
  ((FKB *) fkb)->fkb$q_fr3 = (__int64) fr3;
  ((FKB *) fkb)->fkb$q_fr4 = (__int64) fr4;
  exe_std$queue_fork( (FKB *) fkb );
}
```

FORK_LOCK

Use to acquire a fork spin lock and to optionally save the original IPL.

Format

FORK_LOCK lockidx, savipl_p

Description

Inputs:

lockidx is the integer value of the spin lock index.

Outputs:

savipl_p is a 32-bit integer passed by reference in which the original IPL is returned. If the address of this parameter is NOSAVE_IPL, then the original IPL is not returned. The symbol NOSAVE_IPL is defined to be a null pointer, i.e. ((int *) 0), by the vms_drivers.h file.

For example, one can take the UCB fork lock and store the original IPL by:

```
fork_lock (ucb->ucb$b_flck, &orig_ipl);
```

If there is no need to save the original IPL, then one can take the fork lock by:

```
fork_lock (ucb->ucb$b_flck, NOSAVE_IPL);
```

The definition of the fork_lock macro is:

```
fork_lock(lockidx, savipl_p)
{
    extern SMP smp$gl_flags;
    extern int smp$al_iplvec[];

    if(savipl_p != NOSAVE_IPL)
        *savipl_p = __PAL_MFPR_IPL();

    if(smp$gl_flags.smp$v_enabled)
        smp_std$acquire( lockidx );
    else
        __PAL_MTPR_IPL( smp$al_iplvec[lockidx] );
}
```

Notes:

1. If the savipl_p parameter is specified as NOSAVE_IPL, then all the code for handling the savipl_p parameter may be optimized away by the C compiler. This convention is similar to that used by the device_lock and sys_lock macros.
2. Compiletime type checking will detect an erroneous use of an ordinary constant other than NOSAVE_IPL for the savipl_p parameter.

FORK_UNLOCK

Use to either release or restore (i.e. conditionally release) a fork spin lock and to optionally set a new IPL.

Format

`FORK_UNLOCK lockidx, newipl, restore`

Description

Inputs:

<code>lockidx</code>	is the integer value of the spin lock index.
<code>newipl</code>	is the integer value of the desired new IPL or the value <code>NOLOWER_IPL</code> if the IPL should be left unchanged. The symbol <code>NOLOWER_IPL</code> is defined to be -1 by the <code>vms_drivers.h</code> file.
<code>restore</code>	is the either the integer value <code>SMP_RESTORE</code> or <code>SMP_RELEASE</code> . If <code>SMP_RELEASE</code> is specified then the spin lock is unconditionally released by calling <code>SMP_STD\$RELEASE</code> , otherwise the spin lock is conditionally released by calling <code>SMP_STD\$RESTORE</code> . The symbol <code>SMP_RESTORE</code> is defined to be 1 and the symbol <code>SMP_RELEASE</code> is defined to be 0 by the <code>vms_drivers.h</code> file.

For example, one can conditionally release the UCB fork lock and set IPL by:

```
fork_unlock (ucb->ucb$b_flck, orig_ipl, SMP_RESTORE);
```

If there is no need to change IPL, then one can conditionally release the fork lock by:

```
fork_unlock (ucb->ucb$b_flck, NOLOWER_IPL, SMP_RESTORE);
```

The definition of the `fork_unlock` macro is:

```
fork_unlock(lockidx, newipl, restore)
{
    extern SMP smp$gl_flags;
    if(smp$gl_flags.smp$v_enabled) {
        if(restore == SMP_RELEASE)
            smp_std$release( lockidx );
        else
            smp_std$restore( lockidx );
    }
    if(newipl >= 0)
        __PAL_MTPR_IPL( newipl );
}
```

Notes:

1. If the `newipl` parameter is specified as `NOLOWER_IPL`, then all the code for handling the `newipl` parameter may be optimized away by the C compiler. This convention is similar to that used by the `sys_unlock` macro.

OpenVMS C Macros Used by OpenVMS AXP Device Drivers FORK_UNLOCK

2. Using either the `SMP_RELEASE` or `SMP_RESTORE` constant for the restore parameter allows the C compiler to generate code for either a call to `SMP_STD$RELEASE` or `SMP_STD$RESTORE` without the need for a runtime test. This convention is similar to that used by the `fork_unlock` and `sys_unlock` macros.
3. Caution should be exercised not to use the constants `RAISE_IPL` nor `NORAISE_IPL` instead of `NOLOWER_IPL` for the `newipl` parameter. If either of these are used erroneously in place of `NOLOWER_IPL` the effect is that IPL will be set to either 0 or 1. Remember that IPL is raised (or held) when acquiring a spin lock and lowered (or held) on release.

FORK_WAIT

Use to queue a specified fork routine with specified fork routine parameters for delayed execution. After the fork routine is queued, execution continues with the next statement following the fork macro.

Format

`FORK_WAIT fork_routine, fr3, fr4, fkb`

Description

Inputs:

<code>fork_routine</code>	is the procedure value of the routine that is to be executed in a fork thread. This value is passed to the fork dispatcher via <code>fkb->fkb\$l_fpc</code> .
<code>fr3</code>	is the 64-bit value to pass to the fork routine via <code>fkb->fkb\$q_fr3</code> . This parameter is cast as a 64-bit integer.
<code>fr4</code>	is the 64-bit value to pass to the fork routine via <code>fkb->fkb\$q_fr4</code> . This parameter is cast as a 64-bit integer.
<code>fkb</code>	is a pointer to the fork block. This parameter is cast as a pointer to an FKB.

The definition of the `fork_wait` macro is:

```
fork_wait(fork_routine, fr3, fr4, fkb)
{
    ((FKB *) fkb)->fkb$l_fpc = fork_routine;
    exe_std$primitive_fork_wait( (__int64) fr3, (__int64) fr4,
                                (FKB *) fkb );
}
```

IOFORK

Use to queue a fork routine with specified fork routine parameters. This macro is very similar to fork, except that the fork block is assumed to be a UCB and the ucb\$*v_tim* bit is cleared before the fork routine is queued.

Format

IOFORK fork_routine, fr3, fr4, ucb

Description

Inputs:

fork_routine	is the procedure value of the routine that is to be executed in a fork thread. This value is passed to the fork dispatcher via ucb->ucb\$l_fpc.
fr3	is the 64-bit value to pass to the fork routine via ucb->ucb\$q_fr3. This parameter is cast as a 64-bit integer.
fr4	is the 64-bit value to pass to the fork routine via ucb->ucb\$q_fr4. This parameter is cast as a 64-bit integer.
ucb	is a pointer to the unit control block. This parameter is cast as a pointer to a UCB.

The definition of the iofork macro is:

```
iofork(fork_routine, fr3, fr4, ucb)
{
    ((UCB *) ucb)->ucb$v_tim = 0;
    fork( fork_routine, fr3, fr4, ucb);
}
```

OpenVMS C Macros Used by OpenVMS AXP Device Drivers

RFI

RFI

Use in an interrupt service routine to invoke the resume from interrupt routine that has been set up by either the `wfipch` or `wfirlch` macros.

Format

```
RFI irp, fr4, ucb
```

Description

Inputs:

<code>irp</code>	is a usually a pointer to an IRP type, but can be any value which is expected as the first parameter of the resume from interrupt routine.
<code>fr4</code>	is any value which is expected as the second parameter of the resume from interrupt routine.
<code>ucb</code>	is a pointer to a Unit Control Block and is the third parameter of the resume from interrupt routine. This parameter is cast as a pointer to an UCB.

The definition of the `rfi` macro is:

```
rfi(irp, fr4, ucb)
( *((UCB *) ucb)->ucb$l_fpc ) (irp, fr4, ucb)
```

Note that it may be possible to eliminate the driver resume from interrupt routine (and thus the need to use the `rfi` macro) by moving some processing directly into the interrupt service routine and by resuming the driver in a fork routine. The driver fork routine would then be resumed from the interrupt service routine by:

```
ucb->ucb$v_tim = 0;
exe_std$queue_fork ( (FKB *) ucb );
```

WFIKPCH

Use to set up an interrupt resume routine and a device interrupt timeout routine without releasing the channel, that is, CRB.

Format

WFIKPCH resume_rout, tout_rout, irp, fr4, ucb, tmo, restore_ipl

Description

Inputs:

resume_rout	is the procedure value of the resume from interrupt routine that is to be called by the interrupt service routine. This value is passed to the interrupt service routine via ucb->ucb\$l_fpc.
tout_rout	is the procedure value of the device interrupt timeout routine that may be called by EXE\$TIMEOUT. This value is passed via ucb->ucb\$ps_toutrout.
irp	is a pointer to an IRP type which is passed to the interrupt resume or timeout routine via ucb->ucb\$q_fr3.
fr4	is a 64-bit value to pass to the resume from interrupt or timeout routine via ucb->ucb\$q_fr4. This parameter is cast as a 64-bit integer.
ucb	is a pointer to a Unit Control Block. This parameter is cast as a pointer to an UCB.
tmo	is an integer specifying the timeout value in seconds.
restore_ipl	is an integer specifying the IPL to lower to prior to returning.

For example, a driver start I/O routine might use the wfikpch macro in the following fashion, where start_device_xfer is a device specific routine implemented in the driver:

```
{
  device_lock (ucb->ucb$l_dlck, RAISE_IPL, &orig_ipl);
  start_device_xfer ( ucb, irp );
  wfikpch (resume_rout, tout_rout, irp, 0, ucb, 3, orig_ipl);
  return;
}
```

The definition of the wfikpch macro is:

```
wfikpch(resume_rout, tout_rout, irp, fr4, ucb, tmo, restore_ipl)
{
  ((UCB *) ucb)->ucb$l_fpc = resume_rout;
  ((UCB *) ucb)->ucb$ps_toutrout = tout_rout;
  ioc_std$primitive_wfikpch (irp, (__int64) fr4, (UCB *) ucb,
                             tmo, restore_ipl );
}
```

WFIRLCH

Use to set up an interrupt resume routine and a device interrupt timeout routine, and to release the channel, that is, CRB.

Format

WFIRLCH resume_rout, tout_rout, irp, fr4, ucb, tmo, restore_ipl

Description

The wfirch macro works like the wfkpch macro with the exception that wfirch calls routine IOC_STDSPRIMITIVE_WFIRLCH instead of IOC_STDSPRIMITIVE_WFIKPCH. In all other respects the description of the parameters of the wfkpch macro applies to the wfirch macro.

A

ACBSV_QUOTA, 2-37
ACB (AST control block), 2-20 to 2-21, 2-22 to 2-23, 2-30 to 2-31
 contents, 2-36 to 2-38
ACP_STDSACCESSNET routine, 2-8 to 2-9
ACP_STDSACCESS routine, 2-6 to 2-7
ACP_STDSDEACCESS routine, 2-10 to 2-11
ACP_STDSMODIFY routine, 2-12 to 2-13
ACP_STDSMOUNT routine, 2-14 to 2-15
ACP_STDSREADBLK routine, 2-16 to 2-17
ACP_STDSWRITEBLK routine, 2-18 to 2-19
ADP (adapter control block), 3-3 to 3-12
 child, 3-4
 parent, 3-4
 peer, 3-5
ADP list, 3-4 to 3-5
Alternate start I/O routine, 2-82 to 2-83
Alternate start-I/O routines, 1-2 to 1-3
AST (asynchronous system trap), 2-36 to 2-38, 2-39 to 2-41
 delivering, 2-20 to 2-21, 2-22 to 2-23, 2-77
 for aborted I/O request, 2-77
 process-requested, 2-37
Attention AST
 delivering, 2-20 to 2-21, 2-22 to 2-23
 disabling, 2-36 to 2-38
 enabling, 2-36 to 2-38
 flushing, 2-30 to 2-31

B

Buffer
 allocating, 2-79 to 2-81
 deallocating, 2-28 to 2-29
 locking, 2-103 to 2-106, 2-107 to 2-112, 2-129 to 2-132, 2-137 to 2-142, 2-150 to 2-153, 2-158 to 2-163, 2-278 to 2-279
 moving data to from system to user, 2-246 to 2-248, 4-49
 moving data to from user to system, 2-243 to 2-245, 4-48
 testing accessibility of, 2-103 to 2-106, 2-107 to 2-112, 2-129 to 2-132, 2-133 to 2-136, 2-137 to 2-142, 2-150 to 2-153, 2-154 to 2-157, 2-158 to 2-163

Buffer (cont'd)

 unlocking, 2-280 to 2-281
BUSARRAY, 3-10 to 3-12
Bus array entry, 3-11 to 3-12
BYTCNT (byte count) quota
 debiting, 2-80
BYTLM (byte limit) quota
 debiting, 2-80

C

CALL_ABORTIO macro, 4-7
CALL_ALLOCBUF macro, 4-8
CALL_ALLOCEMB macro, 4-9
CALL_ALLOCIRP macro, 4-8
CALL_ALTQUEPKT macro, 4-10
CALL_ALTREQCOM macro, 4-11
CALL_BROADCAST macro, 4-12
CALL_CANCELIO macro, 4-13
CALL_CARRIAGE macro, 4-14
CALL_CHKCREACCES macro, 4-15
CALL_CHKDELACCES macro, 4-15
CALL_CHKEXEACCES macro, 4-15
CALL_CHKLOGACCES macro, 4-15
CALL_CHKPHYACCES macro, 4-15
CALL_CHKRDACCES macro, 4-15
CALL_CHKWRTACCES macro, 4-15
CALL_CLONE_UCB macro, 4-16
CALL_COPY_UCB macro, 4-17
CALL_CREDIT_UCB macro, 4-18
CALL_CVTLOGPHY macro, 4-19
CALL_CVT_DEVNAM macro, 4-20
CALL_DELATTNAST macro, 4-21
CALL_DELATTNASTP macro, 4-22
CALL_DELCTRLAST macro, 4-23
CALL_DELCTRLASTP macro, 4-24
CALL_DELETE_UCB macro, 4-25
CALL_DEVICEATTN macro, 4-26
CALL_DEVICERR macro, 4-26
CALL_DEVICTMO macro, 4-26
CALL_DIAGBUFILL macro, 4-27
CALL_DRVDEALMEM macro, 4-28
CALL_FILSPT macro, 4-29
CALL_FINISHIOC macro, 4-30
CALL_FINISHIO macro, 4-30

CALL_FINISHIO_NOIOST macro, 4-30
 CALL_FLUSHATTNS macro, 4-31
 CALL_FLUSHCTRLS macro, 4-32
 CALL_GETBYTE macro, 4-33
 CALL_INITBUFWIND macro, 4-34
 CALL_INITIATE macro, 4-35
 CALL_INSERT_IRP macro, 4-36
 CALL_IOLOCK macro, 4-37
 CALL_IOLOCKR macro, 4-38
 CALL_IOLOCKW macro, 4-39
 CALL_IORSNWAIT macro, 4-40
 CALL_IOUNLOCK macro, 4-41
 CALL_LINK_UCB macro, 4-42
 CALL_MAPVBLK macro, 4-43
 CALL_MNTVER macro, 4-44
 CALL_MNTVERSIO macro, 4-45
 CALL_MODIFYLOCK macro, 4-46
 CALL_MODIFYLOCK_ERR macro, 4-46
 CALL_MOUNT_VER macro, 4-47
 CALL_MOVFRUSER2 macro, 4-48
 CALL_MOVFRUSER macro, 4-48
 CALL_MOVTOUSER2 macro, 4-49
 CALL_MOVTOUSER macro, 4-49
 CALL_PARSDEVNAM macro, 4-50
 CALL_POST macro, 4-51
 CALL_POST_IRP macro, 4-52
 CALL_POST_NOCNT macro, 4-51
 CALL_PTETOPFN macro, 4-53
 CALL_QIOACPPKT macro, 4-54
 CALL_QIODRVPKT macro, 4-55
 CALL_QNXTSEG1 macro, 4-56
 CALL_QXQPPKT macro, 4-57
 CALL_READCHK macro, 4-58
 CALL_READCHKR macro, 4-58
 CALL_READLOCK macro, 4-59
 CALL_READLOCK_ERR macro, 4-59
 CALL_RELCHAN macro, 4-60
 CALL_RELEASEMB macro, 4-61
 CALL_REQCOM macro, 4-62
 CALL_SEARCHDEV macro, 4-63
 CALL_SEARCHINT macro, 4-64
 CALL_SETATTNAST macro, 4-65
 CALL_SETCTRLAST macro, 4-66
 CALL_SEVER_UCB macro, 4-67
 CALL_SIMREQCOM macro, 4-68
 CALL_SNDEVMSG macro, 4-69
 CALL_THREADCRB macro, 4-70
 CALL_UNLOCK macro, 4-71
 CALL_WRITECHK macro, 4-72
 CALL_WRITECHKR macro, 4-72
 CALL_WRITELOCK macro, 4-73
 CALL_WRITELOCK_ERR macro, 4-73
 CALL_WRTMAILBOX macro, 4-74
 Cancel I/O routine
 flushing ASTs in, 2-30 to 2-31
 Cancel-I/O routines, 1-4 to 1-6

Cancel selective routines, 1-7 to 1-8
 CCB (channel control block), 3-12 to 3-13
 Channel assign routines, 1-9 to 1-10
 Channel index number, 2-216
 CLASS_UNIT_INIT macro, 4-75 to 4-76
 Cloned UCB routines, 1-11 to 1-13
 COM_STDSDELATTNASTP routine, 2-22 to 2-23
 COM_STDSDELATTNAST routine, 2-20 to 2-21
 COM_STDSDELCTRLASTP routine, 2-26 to 2-27
 COM_STDSDELCTRLAST routine, 2-24 to 2-25
 COM_STSDRVDEALMEM routine, 2-28 to 2-29
 COM_STDSFLUSHATTNS, 2-37
 COM_STDSFLUSHATTNS routine, 2-30 to 2-31
 COM_STDSFLUSHCTRLS routine, 2-32 to 2-33
 COM_STDS\$POST routine, 2-34 to 2-35
 COM_STDS\$POST_NOCNT routine, 2-34 to 2-35
 COM_STDS\$SETATTNAST routine, 2-36 to 2-38
 COM_STDS\$SETCTRLAST routine, 2-39 to 2-41
 Control AST
 disabling, 2-39 to 2-41
 enabling, 2-39 to 2-41
 Controller initialization routines, 1-14 to 1-16
 Counted resource
 defined, 2-175
 CPUDISP macro, 4-77
 CRAM (controller register access mailbox), 3-13
 to 3-17
 CRAM_ALLOC macro, 4-78
 CRAM_CMD macro, 4-79 to 4-80
 CRAM_DEALLOC macro, 4-81
 CRAM_IO macro, 4-82
 CRAM_QUEUE macro, 4-83
 CRAM_WAIT macro, 4-84
 CRB (channel request block), 3-17 to 3-19

D

Data transfer
 zero byte count, 2-105, 2-131, 2-152
 DDB (device data block), 3-20 to 3-21
 DDT (driver dispatch table), 3-22 to 3-25
 DDTAB macro, 4-85 to 4-88
 Device
 disk, 2-146, 2-265
 tape, 2-265
 Device affinity, 2-236
 Device characteristics
 retrieving, 2-143 to 2-144
 setting, 2-145 to 2-147
 Device controller data channel
 releasing, 2-262 to 2-263, 4-146
 Device controller data channel wait queue, 2-262
 Device drivers
 Step 1, xiii
 Step 2, xiii
 DEVICELOCK macro, 4-89 to 4-90

Device unit
 operations count, 2-265
 DEVICE_LOCK macro, 5-2
 Diagnostic buffer, 2-236
 filling, 2-227 to 2-228
 Direct I/O
 checking accessibility of process buffer for,
 2-133 to 2-136, 2-154 to 2-157
 locking a process buffer for, 2-103 to 2-106,
 2-107 to 2-112, 2-129 to 2-132, 2-137 to
 2-142, 2-150 to 2-153, 2-158 to 2-163
 unlocking process buffer, 2-280 to 2-281
 Disk driver, 2-98 to 2-100
 DMA transfer
 for read operation, 2-129 to 2-132, 2-137 to
 2-142
 for read/write operation, 2-103 to 2-106
 for write operation, 2-107 to 2-112, 2-150 to
 2-153, 2-158 to 2-163
 Documentation comments, sending to Digital, iii
 DPTSV_SVP, 2-244, 2-247
 DPT (driver prologue table), 3-25 to 3-30
 DPTAB macro, 4-91 to 4-95
 DPT_STORE macro, 4-96 to 4-98
 DPT_STORE_ISR macro, 4-99
 Driver entry points, 1-1 to 1-48
 Driver macros, 4-1 to Index-1
 Driver unloading routines, 1-21
 \$DRIVER_ALTSTART_ENTRY macro, 1-2
 \$DRIVER_CANCEL_ENTRY macro, 1-4
 \$DRIVER_CANCEL_SELECTIVE_ENTRY macro,
 1-7
 \$DRIVER_CHANNEL_ASSIGN_ENTRY macro,
 1-9
 \$DRIVER_CLONEDUCB_ENTRY macro, 1-12
 DRIVER_CODE macro, 4-105
 \$DRIVER_CTRLINIT_ENTRY macro, 1-14
 DRIVER_DATA macro, 4-114
 \$DRIVER_DELIVER_ENTRY macro, 1-44
 \$DRIVER_ERRRTN_ENTRY macro, 1-25
 \$DRIVER_FDT_ENTRY macro, 1-22
 \$DRIVER_MNTVER_ENTRY macro, 1-31
 \$DRIVER_REGDUMP_ENTRY macro, 1-33
 \$DRIVER_START_ENTRY macro, 1-35
 \$DRIVER_UNITINIT_ENTRY macro, 1-46

E

ERL_STDS\$ALLOCEMB routine, 2-42 to 2-43
 ERL_STDS\$DEVICEATTN routine, 2-44 to 2-46
 ERL_STDS\$DEVICERR routine, 2-44 to 2-46
 ERL_STDS\$DEVICTMO routine, 2-44 to 2-46
 ERL_STDS\$RELEASEMB routine, 2-47
 Error logging, 2-44 to 2-46
 Error message buffer
 releasing, 2-265

Event flag
 handling for aborted I/O request, 2-77
 EXESBUS_DELAY, 2-48 to 2-49
 EXES\$DELAY, 2-50
 EXESILLIOFUNC, 2-90 to 2-91
 EXESKP_ALLOCATE_KPB, 2-51 to 2-53
 EXESKP_DEALLOCATE_KPB, 2-54 to 2-55
 EXESKP_END, 2-56 to 2-57
 EXESKP_FORK, 2-58 to 2-59
 EXESKP_FORK_WAIT, 2-60 to 2-61
 EXESKP_RESTART, 2-62 to 2-63
 EXESKP_STALL_GENERAL, 2-64 to 2-66
 EXESKP_START, 2-67 to 2-69
 EXESKP_STARTIO, 2-70 to 2-71
 EXESTIMEDWAIT_COMPLETE, 2-72 to 2-73
 EXESTIMEDWAIT_SETUP, 2-74 to 2-75
 EXESTIMEDWAIT_SETUP_10US, 2-74 to 2-75
 EXESWRTMAILBOX routine, 2-164 to 2-165
 EXE_STDSABORTIO, 2-146
 EXE_STDSABORTIO system routine, 2-76 to
 2-78
 EXE_STDS\$ALLOCBUF routine, 2-79 to 2-81
 EXE_STDS\$ALLOCIRP routine, 2-79 to 2-81
 EXE_STDS\$ALTQUEPKT routine, 2-82 to 2-83
 EXE_STDS\$CARRIAGE routine, 2-84
 EXE_STDS\$CHKCREACCES routine, 2-85 to
 2-86
 EXE_STDS\$CHKDELACCES routine, 2-85 to
 2-86
 EXE_STDS\$CHKEXEACCES routine, 2-85 to 2-86
 EXE_STDS\$CHKLOGACCES routine, 2-85 to
 2-86
 EXE_STDS\$CHKPHYACCES routine, 2-85 to 2-86
 EXE_STDS\$CHKRDACCES routine, 2-85 to 2-86
 EXE_STDS\$CHKWRTACCES routine, 2-85 to
 2-86
 EXE_STDS\$FINISHIO, 2-99, 2-100
 EXE_STDS\$FINISHIO routine, 2-87 to 2-89,
 2-147
 EXE_STDS\$INSERT_IRP routine, 2-92 to 2-93
 EXE_STDS\$INSIOQ, 2-124
 EXE_STDS\$INSIOQC routine, 2-94 to 2-95
 EXE_STDS\$INSIOQ routine, 2-94 to 2-95
 EXE_STDS\$IORSNWAIT routine, 2-96 to 2-97
 EXE_STDS\$LCLDSKVALID routine, 2-98 to 2-100
 EXE_STDS\$MNTVERSIO routine, 2-101 to 2-102
 EXE_STDS\$MODIFYLOCK, 2-280
 EXE_STDS\$MODIFYLOCK routine, 2-107 to
 2-112
 EXE_STDS\$MODIFY routine, 2-103 to 2-106
 EXE_STDS\$MOUNT_VER routine, 2-113 to 2-114
 EXE_STDS\$ONEPARM routine, 2-115 to 2-116
 EXE_STDS\$PRIMITIVE_FORK routine, 2-117 to
 2-118
 EXE_STDS\$PRIMITIVE_FORK_WAIT routine,
 2-119 to 2-120

EXE_STDSQIOACPPKT routine, 2-121 to 2-122
EXE_STDSQIODRVPKT, 2-100, 2-105, 2-116,
2-131, 2-167
EXE_STDSQIODRVPKT routine, 2-123 to 2-125,
2-147, 2-152
EXE_STDSQXQPPKT routine, 2-127 to 2-128
EXE_STDSREADCHK routine, 2-133 to 2-136
EXE_STDSREADLOCK, 2-280
EXE_STDSREADLOCK routine, 2-137 to 2-142
EXE_STDSREAD routine, 2-129 to 2-132
EXE_STDSSENSEMODE routine, 2-143 to 2-144
EXE_STDSSETCHAR routine, 2-145 to 2-147
EXE_STDSSETMODE routine, 2-145 to 2-147
EXE_STDSNDEVMSG routine, 2-148 to 2-149
EXE_STDSWRITECHK routine, 2-154 to 2-157
EXE_STDSWRITELOCK, 2-280
EXE_STDSWRITELOCK routine, 2-158 to 2-163
EXE_STDSWRITE routine, 2-150 to 2-153
EXE_STDSWRMAILBOX routine, 2-149
EXE_STDSZEROPARM, 2-166 to 2-167

F

FDT error handling callback routines, 1-25 to
1-27
FDT routine
adjusting process quotas in, 2-80
completing an I/O operation in, 2-87 to 2-89
for direct I/O, 2-103 to 2-106, 2-129 to 2-132,
2-150 to 2-153
for disk I/O, 2-98 to 2-100
setting attention ASTs in, 2-36 to 2-38
setting control ASTs in, 2-39 to 2-41
unlocking process buffers in, 2-280
FDT routines, 1-22 to 1-24
FDT_ACT macro, 4-116 to 4-117
FDT_BUF macro, 4-118
FDT_INI macro, 4-29, 4-119
Feedback on documentation, sending to Digital, iii
FORKLOCK macro, 4-125 to 4-126
FORK macro, 4-120 to 4-121
Fork process
creation by IOC_STDS\$INITIATE, 2-235 to
2-237
FORK routine, 4-122
FORK_ROUTINE macro, 4-122
FORK_WAIT macro, 4-123 to 4-124
Full duplex device driver
I/O completion for, 2-34 to 2-35

I

I/O database, 3-1 to 3-3
I/O postprocessing
for aborted I/O request, 2-77
for full duplex device driver, 2-34 to 2-35
for I/O request involving no device activity,
2-87 to 2-89

I/O postprocessing queue, 2-34 to 2-35, 2-265
I/O request
aborting, 2-76 to 2-78
canceling, 2-215 to 2-216
completing, 2-264 to 2-266, 2-274 to 2-275
with no parameters, 2-166 to 2-167
with one parameter, 2-115 to 2-116
IDB (interrupt dispatch block), 3-30 to 3-32
IFNORD macro, 4-129 to 4-131
IFNOWRT macro, 4-129 to 4-131
IFRD macro, 4-129 to 4-131
IFWRT macro, 4-129 to 4-131
Interrupt service routines, 1-28 to 1-30
IOS_AVAILABLE function
servicing, 2-100
IOS_PACKACK function
servicing, 2-99
IOS_SENSECHAR function
servicing, 2-143 to 2-144
IOS_SENSEMODE function
servicing, 2-143 to 2-144
IOS_SETCHAR function
servicing, 2-145 to 2-147
IOS_SETMODE function
servicing, 2-145 to 2-147
IOS_UNLOAD function
servicing, 2-100
IOCSALLOCATE_CRAM, 2-178 to 2-179
IOCSALLOC_CNT_RES, 2-170 to 2-173
IOCSALLOC_CRAB, 2-174 to 2-175
IOCSALLOC_CRCTX, 2-176 to 2-177
IOCSALOALTMAP, 2-168
IOCSALOALTMAPN, 2-168
IOCSALOALTMAPSP, 2-168
IOCSALOUBAMAP, 2-169
IOCSALOUBAMAPN, 2-169
IOCSCANCEL_CNT_RES, 2-172, 2-180 to 2-181
IOCSGRAM_CMD, 2-182 to 2-184
IOCSGRAM_IO, 2-185 to 2-186
IOCSGRAM_QUEUE, 2-187 to 2-188
IOCSGRAM_WAIT, 2-189 to 2-190
IOCSDEALLOCATE_CRAM, 2-195
IOCSDEALLOC_CNT_RES, 2-191 to 2-192
IOCSDEALLOC_CRAB, 2-193
IOCSDEALLOC_CRCTX, 2-194
IOCSKP_REQCHAN, 2-196 to 2-197
IOCSKP_WFIKPCH, 2-198 to 2-199
IOCSKP_WFIRLCH, 2-198 to 2-199
IOCSLOAD_MAP, 2-200 to 2-201
IOCSMAP_IO, 2-202
IOCSNODE_FUNCTION, 2-204 to 2-206
IOCSRELCHAN routine, 4-146
IOCSREQCOM routine, 4-149
IOC_STDS\$ALTREQCOM routine, 2-211 to 2-212
IOC_STDS\$BROADCAST routine, 2-213 to 2-214
IOC_STDS\$CANCELIO routine, 2-215 to 2-216

IOC_STDS\$CLONE_UCB routine, 2-217 to 2-218
 IOC_STDS\$COPY_UCB routine, 2-219 to 2-220
 IOC_STDS\$CREDIT_UCB routine, 2-221
 IOC_STDS\$CVTLOGPHY routine, 2-224 to 2-225
 IOC_STDS\$CVT_DEVNAM routine, 2-222 to 2-223
 IOC_STDS\$DELETE_UCB routine, 2-226
 IOC_STDS\$DIAGBUFILL routine, 2-227 to 2-228
 IOC_STDS\$FILSPT routine, 2-229 to 2-230
 IOC_STDS\$GETBYTE routine, 2-231 to 2-232
 IOC_STDS\$INITBUFWIND routine, 2-233 to 2-234
 IOC_STDS\$INITIATE routine, 2-235 to 2-237
 IOC_STDS\$IOPPOST
 unlocking process buffers, 2-280
 IOC_STDS\$LINK_UCB routine, 2-238 to 2-239
 IOC_STDS\$MAPVBLK routine, 2-240 to 2-241
 IOC_STDS\$MNTVER routine, 2-242
 IOC_STDS\$MOVFRUSER2 routine, 2-243 to 2-245
 IOC_STDS\$MOVFRUSER routine, 2-243 to 2-245
 IOC_STDS\$MOVTOUSER2 routine, 2-246 to 2-248
 IOC_STDS\$MOVTOUSER routine, 2-246 to 2-248
 IOC_STDS\$PARSDEVNAM routine, 2-249 to 2-250
 IOC_STDS\$POST_IRP routine, 2-251
 IOC_STDS\$PRIMITIVE_REQCHANH routine, 2-256 to 2-258
 IOC_STDS\$PRIMITIVE_REQCHANL routine, 2-256 to 2-258
 IOC_STDS\$PRIMITIVE_WFIKPCH routine, 2-259 to 2-261
 IOC_STDS\$PRIMITIVE_WFIRLCH routine, 2-259 to 2-261
 IOC_STDS\$PTETOPFN routine, 2-252 to 2-253
 IOC_STDS\$QNXTSEG1 routine, 2-254 to 2-255
 IOC_STDS\$RELCHAN routine, 2-262 to 2-263
 IOC_STDS\$REQCOM, 2-80
 IOC_STDS\$REQCOM routine, 2-264 to 2-266
 IOC_STDS\$SEARCHDEV routine, 2-267 to 2-268
 IOC_STDS\$SEARCHINT routine, 2-269 to 2-270
 IOC_STDS\$SENSEDISK routine, 2-271 to 2-272
 IOC_STDS\$SEVER_UCB routine, 2-273
 IOC_STDS\$SIMREQCOM routine, 2-274 to 2-275
 IOC_STDS\$THREADCRB routine, 2-276 to 2-277
 IOFORK macro, 4-127 to 4-128
 IOSB (I/O status block), 2-265
 IPL\$ASTDEL, 2-97, 2-121, 2-124, 2-127
 IPL\$IOPOST, 2-35, 2-77, 2-88, 2-265
 IPL\$MAILBOX, 2-149, 2-164
 IRP\$B_CARCON, 2-104, 2-130, 2-151
 IRP\$L_BCNT, 2-235, 2-236
 IRP\$L_BOFF, 2-235, 2-236
 IRP\$L_CHAN, 2-216
 IRP\$L_DIAGBUF, 2-235, 2-236

IRP\$L_MEDIA, 2-116, 2-147, 2-167
 IRP\$L_PID, 2-216
 IRP\$L_SVAPTE, 2-236
 IRP\$V_DIAGBUF, 2-235, 2-236
 IRP\$V_FUNC, 2-130
 IRP (I/O request packet), 3-32 to 3-38
 insertion in pending-I/O queue, 2-92 to 2-93, 2-94 to 2-95
 unlocking buffers specified in, 2-280
 IRPE (I/O request packet extension), 3-38 to 3-39
 deallocation, 2-280
 unlocking buffers specified in, 2-280

J

JIB\$L_BYTCNT, 2-80
 JIB\$L_BYTLM, 2-80
 Job controller
 sending a message to, 2-149, 2-164 to 2-165
 Job quota
 byte count, 2-80
 byte limit, 2-80

K

KPB (kernel process block), 3-39 to 3-46
 KP_ALLOCATE_KPB macro, 4-132
 KP_DEALLOCATE_KPB macro, 4-133
 KP_END macro, 4-134
 KP_REQCOM macro, 4-136
 KP_RESTART macro, 4-135
 KP_STALL_FORK, KP_STALL_IOFORK macro, 4-137
 KP_STALL_FORK_WAIT macro, 4-138
 KP_STALL_GENERAL macro, 4-139
 KP_STALL_REQCHAN macro, 4-140
 KP_STALL_WFIKPCH macro, 4-141
 KP_STALL_WFIRLCH macro, 4-141
 KP_START macro, 4-142
 KP_SWITCH_TO_KP_STACK macro, 4-143

L

Local disk
 online count, 2-98
 valid bit, 2-98
 Local disk UCB extension
 required for EXE_STDS\$LCLDSKVALID routine, 2-100
 LOCK macro, 4-144 to 4-145
 Logical I/O function
 translation to physical function, 2-103, 2-129, 2-150

M

Mailbox

- sending a message to, 2-148 to 2-149, 2-164 to 2-165
- MAILBOX spin lock, 2-149, 2-164
- MMG spin lock, 2-280
- MMG_STDSIOLOCK routine, 2-278 to 2-279
- MMG_STDSUNLOCK routine, 2-280 to 2-281
- Mount verification routines, 1-31 to 1-32
- MT_STDSCHECK_ACCESS routine, 2-282 to 2-283

Mutex

- locking, 2-284 to 2-285, 2-286 to 2-287, 4-38, 4-39
- unlocking, 2-288, 4-41

N

Nonpaged pool

- allocating, 2-79 to 2-81
- deallocating, 2-28 to 2-29
- lookaside list, 2-80

O

OPCOM process

- sending a message to, 2-149, 2-164 to 2-165
- ORB (object rights block), 3-46 to 3-47

P

PCBSL_PID, 2-216

- Pending-I/O queue, 2-92 to 2-93, 2-94 to 2-95, 2-123 to 2-125, 2-265
 - bypassing, 2-82 to 2-83
- PMI (processor-memory interconnect), 3-3

Q

- QUEUEAST spin lock, 2-37

R

- Register dumping routines, 1-33 to 1-34
- RELCHAN macro, 4-146
- REQCHAN macro, 4-147 to 4-148
- REQCOM macro, 4-149
- REQPCHAN macro, 4-150
- Resource wait mode, 2-80

S

- SCH_STDSIOLOCKR routine, 2-284 to 2-285
- SCH_STDSIOLOCKW routine, 2-286 to 2-287

- SCH_STDSIOUNLOCK routine, 2-288

- SS\$_ACCVIO, 2-146

- SS\$_ILLIOFUNC, 2-146

Start I/O routine

- activating, 2-94 to 2-95
- checking for zero length buffer, 2-105, 2-131, 2-152
- transferring control to, 2-123 to 2-125, 2-235 to 2-237

Start-I/O routines

- kernel process, 1-38 to 1-39
- Simple Fork, CALL Environment, 1-35
- Step 1 device driver, xiii
- Step 2 device driver, xiii
- SYSDISP macro, 4-151
- System page-table entry
 - allocating permanent, 2-244, 2-247

T

- TBI_ALL macro, 4-152

- TBI_DATA_64 macro, 4-153

- TBI_SINGLE macro, 4-154

- TBI_SINGLE_64 macro, 4-155

- TIMEDWAIT macro, 4-156 to 4-158

Timeout handling code

- kernel process, 1-42 to 1-43
- traditional, 1-40 to 1-41

U

- UCBSB_DEVCLASS, 2-147

- UCBSB_DEVTYPE, 2-147

- UCBSB_ONLCNT, 2-98

- UCBSL_AFFINITY, 2-236

- UCBSL_BCNT, 2-236

- UCBSL_BOFF, 2-236

- UCBSL_DEVDEPEND, 2-144

- UCBSL_IRP, 2-236

- UCBSL_SVAPTE, 2-236, 2-244, 2-247

- UCBSL_SVPN, 2-243, 2-246

- UCBSV_BSY, 2-216

- UCBSV_CANCEL, 2-216, 2-236

- UCBSV_ERLOGIP, 2-265

- UCBSV_LCL_VALID, 2-98

- UCBSV_TIMEOUT, 2-236

- UCBSW_BUFQUO

- in mailbox UCB, 2-165

- UCBSW_DEVBUFSIZ, 2-147

- in mailbox UCB, 2-165

- UCB (unit control block), 3-47 to 3-68

- error log extension, 3-60

- local disk extension, 2-100, 3-61

- local tape extension, 3-60 to 3-61

- terminal extension, 3-61 to 3-68

- UCBLQ_DEVDEPEND, 2-147

Unit delivery routines, 1-44 to 1-45
Unit initialization routines, 1-46 to 1-48

V

VEC (interrupt transfer vector block), 3-19 to
3-20
VLE (vector list extension), 3-68 to 3-69

W

WFIKPCH macro, 4-159 to 4-161
WFIRLCH macro, 4-159 to 4-161

