

```

*****
13502 Sat Feb 15 09:54:06 2020
new/usr/src/man/man3c/cond_init.3c
12309 errors in section 9e of the manual
*****
1 .\"
2 .\" Sun Microsystems, Inc. gratefully acknowledges The Open Group for
3 .\" permission to reproduce portions of its copyrighted documentation.
4 .\" Original documentation from The Open Group can be obtained online at
5 .\" http://www.opengroup.org/bookstore/.
6 .\"
7 .\" The Institute of Electrical and Electronics Engineers and The Open
8 .\" Group, have given us permission to reprint portions of their
9 .\" documentation.
10 .\"
11 .\" In the following statement, the phrase ``this text'' refers to portions
12 .\" of the system documentation.
13 .\"
14 .\" Portions of this text are reprinted and reproduced in electronic form
15 .\" in the SunOS Reference Manual, from IEEE Std 1003.1, 2004 Edition,
16 .\" Standard for Information Technology -- Portable Operating System
17 .\" Interface (POSIX), The Open Group Base Specifications Issue 6,
18 .\" Copyright (C) 2001-2004 by the Institute of Electrical and Electronics
19 .\" Engineers, Inc and The Open Group. In the event of any discrepancy
20 .\" between these versions and the original IEEE and The Open Group
21 .\" Standard, the original IEEE and The Open Group Standard is the referee
22 .\" document. The original Standard can be obtained online at
23 .\" http://www.opengroup.org/unix/online.html.
24 .\"
25 .\" This notice shall appear on any product containing this material.
26 .\"
27 .\" The contents of this file are subject to the terms of the
28 .\" Common Development and Distribution License (the "License").
29 .\" You may not use this file except in compliance with the License.
30 .\"
31 .\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
32 .\" or http://www.opensolaris.org/os/licensing.
33 .\" See the License for the specific language governing permissions
34 .\" and limitations under the License.
35 .\"
36 .\" When distributing Covered Code, include this CDDL HEADER in each
37 .\" file and include the License file at usr/src/OPENSOLARIS.LICENSE.
38 .\" If applicable, add the following below this CDDL HEADER, with the
39 .\" fields enclosed by brackets "[]" replaced with your own identifying
40 .\" information: Portions Copyright [yyyy] [name of copyright owner]
41 .\"
42 .\"
43 .\" Portions Copyright (c) 1995 IEEE. All Rights Reserved.
44 .\" Copyright (c) 2001, The IEEE and The Open Group. All Rights Reserved.
45 .\" Copyright (c) 2007, Sun Microsystems, Inc. All Rights Reserved.
46 .\"
47 .TH COND_INIT 3C "February 15, 2020"
47 .TH COND_INIT 3C "Jun 5, 2007"
48 .SH NAME
49 cond_init, cond_wait, cond_timedwait, cond_reltimedwait, cond_signal,
50 cond_broadcast, cond_destroy \- condition variables
51 .SH SYNOPSIS
52 .LP
52 .nf
53 cc -mt [ \fIflag\fR... ] \fIfile\fR... [ \fIlibrary\fR... ]
54 #include <thread.h>
55 #include <synch.h>

57 \fBint\fR \fBcond_init\fR(\fBcond_t * \fR\fIcvp\fR, \fBint\fR \fItype\fR, \fBvoid
58 .fi

```

```

60 .LP
61 .nf
62 \fBint\fR \fBcond_wait\fR(\fBcond_t * \fR\fIcvp\fR, \fBmutex_t * \fR\fImp\fR);
63 .fi

65 .LP
66 .nf
67 \fBint\fR \fBcond_timedwait\fR(\fBcond_t * \fR\fIcvp\fR, \fBmutex_t * \fR\fImp\fR,
68 \fBtimestruc_t * \fR\fIabstime\fR);
69 .fi

71 .LP
72 .nf
73 \fBint\fR \fBcond_reltimedwait\fR(\fBcond_t * \fR\fIcvp\fR, \fBmutex_t * \fR\fImp\fR,
74 \fBtimestruc_t * \fR\fIreltime\fR);
75 .fi

77 .LP
78 .nf
79 \fBint\fR \fBcond_signal\fR(\fBcond_t * \fR\fIcvp\fR);
80 .fi

82 .LP
83 .nf
84 \fBint\fR \fBcond_broadcast\fR(\fBcond_t * \fR\fIcvp\fR);
85 .fi

87 .LP
88 .nf
89 \fBint\fR \fBcond_destroy\fR(\fBcond_t * \fR\fIcvp\fR);
90 .fi

92 .SH DESCRIPTION
93 .SS "Initialize"
95 .sp
96 .LP
94 Condition variables and mutexes should be global. Condition variables that are
95 allocated in writable memory can synchronize threads among processes if they
96 are shared by the cooperating processes (see \fBmmap\fR(2)) and are initialized
97 for this purpose.
98 .sp
99 .LP
100 The scope of a condition variable is either intra-process or inter-process.
101 This is dependent upon whether the argument is passed implicitly or explicitly
102 to the initialization of that condition variable. A condition variable does not
103 need to be explicitly initialized. A condition variable is initialized with all
104 zeros, by default, and its scope is set to within the calling process. For
105 inter-process synchronization, a condition variable must be initialized once,
106 and only once, before use.
107 .sp
108 .LP
109 A condition variable must not be simultaneously initialized by multiple threads
110 or re-initialized while in use by other threads.
111 .sp
112 .LP
113 Attributes of condition variables can be set to the default or customized at
114 initialization.
115 .sp
116 .LP
117 The \fBcond_init()\fR function initializes the condition variable pointed to by
118 \fIcvp\fR. A condition variable can have several different types of behavior,
119 specified by \fItype\fR. No current type uses \fIarg\fR although a future type
120 may specify additional behavior parameters with \fIarg\fR. The \fItype\fR
121 argument c take one of the following values:
122 .sp
123 .ne 2

```

```

124 .na
125 \fB\FBUSYNC_THREAD\fR\fR
126 .ad
127 .RS 17n
128 The condition variable can synchronize threads only in this process. This is
129 the default.
130 .RE

132 .sp
133 .ne 2
134 .na
135 \fB\FBUSYNC_PROCESS\fR\fR
136 .ad
137 .RS 17n
138 The condition variable can synchronize threads in this process and other
139 processes. Only one process should initialize the condition variable. The
140 object initialized with this attribute must be allocated in memory shared
141 between processes, either in System V shared memory (see \fBshmop\fR(2)) or in
142 memory mapped to a file (see \fBmmap\fR(2)). It is illegal to initialize the
143 object this way and to not allocate it in such shared memory.
144 .RE

146 .sp
147 .LP
148 Initializing condition variables can also be accomplished by allocating in
149 zeroed memory, in which case, a \fItype\fR of \fB\FBUSYNC_THREAD\fR is assumed.
150 .sp
151 .LP
152 If default condition variable attributes are used, statically allocated
153 condition variables can be initialized by the macro \fBDEFAULTCV\fR.
154 .sp
155 .LP
156 Default condition variable initialization (intra-process):
157 .sp
158 .in +2
159 .nf
160 cond_t cvp;

162 cond_init(&cvp, NULL, NULL); /*initialize condition variable
163                               with default*/
164 .fi
165 .in -2

167 .sp
168 .LP
169 or
170 .sp
171 .in +2
172 .nf
173 cond_init(&cvp, USYNC_THREAD, NULL);
174 .fi
175 .in -2

177 .sp
178 .LP
179 or
180 .sp
181 .in +2
182 .nf
183 cond_t cond = DEFAULTCV;
184 .fi
185 .in -2

187 .sp
188 .LP
189 Customized condition variable initialization (inter-process):

```

```

190 .sp
191 .in +2
192 .nf
193 cond_init(&cvp, USYNC_PROCESS, NULL); /* initialize cv with
194                                           inter-process scope */
195 .fi
196 .in -2

198 .SS "Condition Wait"
202 .sp
203 .LP
199 The condition wait interface allows a thread to wait for a condition and
200 atomically release the associated mutex that it needs to hold to check the
201 condition. The thread waits for another thread to make the condition true and
202 that thread's resulting call to signal and wakeup the waiting thread.
203 .sp
204 .LP
205 The \fBcond_wait()\fR function atomically releases the mutex pointed to by
206 \fImp\fR and causes the calling thread to block on the condition variable
207 pointed to by \fIcvp\fR. The blocked thread may be awakened by
208 \fBcond_signal()\fR, \fBcond_broadcast()\fR, or when interrupted by delivery of
209 a \fBUNIX\fR signal or a \fBfork()\fR.
210 .sp
211 .LP
212 The \fBcond_wait()\fR, \fBcond_timedwait()\fR, and \fBcond_reltimedwait()\fR
213 functions always return with the mutex locked and owned by the calling thread
214 even when returning an error, except when the mutex has the \fBBLOCK_ROBUST\fR
215 attribute and has been left irrecoverable by the mutex's last owner. The
216 \fBcond_wait()\fR, \fBcond_timedwait()\fR, and \fBcond_reltimedwait()\fR
217 functions return the appropriate error value if they fail to internally
218 reacquire the mutex.
219 .SS "Condition Signaling"
225 .sp
226 .LP
220 A condition signal allows a thread to unblock a single thread waiting on the
221 condition variable, whereas a condition broadcast allows a thread to unblock
222 all threads waiting on the condition variable.
223 .sp
224 .LP
225 The \fBcond_signal()\fR function unblocks one thread that is blocked on the
226 condition variable pointed to by \fIcvp\fR.
227 .sp
228 .LP
229 The \fBcond_broadcast()\fR function unblocks all threads that are blocked on
230 the condition variable pointed to by \fIcvp\fR.
231 .sp
232 .LP
233 If no threads are blocked on the condition variable, then \fBcond_signal()\fR
234 and \fBcond_broadcast()\fR have no effect.
235 .sp
236 .LP
237 The \fBcond_signal()\fR or \fBcond_broadcast()\fR functions can be called by a
238 thread whether or not it currently owns the mutex that threads calling
239 \fBcond_wait()\fR, \fBcond_timedwait()\fR, or \fBcond_reltimedwait()\fR have
240 associated with the condition variable during their waits. If, however,
241 predictable scheduling behavior is required, then that mutex should be locked
242 by the thread prior to calling \fBcond_signal()\fR or \fBcond_broadcast()\fR.
243 .SS "Destroy"
251 .sp
252 .LP
244 The condition destroy functions destroy any state, but not the space,
245 associated with the condition variable.
246 .sp
247 .LP
248 The \fBcond_destroy()\fR function destroys any state associated with the
249 condition variable pointed to by \fIcvp\fR. The space for storing the condition

```

```

250 variable is not freed.
251 .SH RETURN VALUES
261 .sp
262 .LP
252 Upon successful completion, these functions return \fB0\fR. Otherwise, a
253 non-zero value is returned to indicate the error.
254 .SH ERRORS
266 .sp
267 .LP
255 The \fBcond_timedwait()\fR and \fBcond_reltimedwait()\fR functions will fail
256 if:
257 .sp
258 .ne 2
259 .na
260 \fB\fbetime\fR\fR
261 .ad
262 .RS 9n
263 The time specified by \fIabstime\fR or \fIreltime\fR has passed.
264 .RE

266 .sp
267 .LP
268 The \fBcond_wait()\fR, \fBcond_timedwait()\fR, and \fBcond_reltimedwait()\fR
269 functions will fail if:
270 .sp
271 .ne 2
272 .na
273 \fB\fbaintr\fR\fR
274 .ad
275 .RS 9n
276 Interrupted. The calling thread was awakened by the delivery of a UNIX signal.
277 .RE

279 .sp
280 .LP
281 If the mutex pointed to by \fImp\fR is a robust mutex (initialized with the
282 \fBLOCK_ROBUST\fR attribute), the \fBcond_wait()\fR, \fBcond_timedwait()\fR and
283 \fBcond_reltimedwait()\fR functions will, under the specified conditions,
284 return the following error values. For complete information, see the
285 description of the \fBmutex_lock()\fR function on the \fBmutex_init\fR(3C)
286 manual page.
287 .sp
288 .ne 2
289 .na
290 \fB\fbenotrecoverable\fR\fR
291 .ad
292 .RS 19n
293 The mutex was protecting the state that has now been left irrecoverable. The
294 mutex has not been acquired.
295 .RE

297 .sp
298 .ne 2
299 .na
300 \fB\fbownerdead\fR\fR
301 .ad
302 .RS 19n
303 The last owner of the mutex died while holding the mutex, possibly leaving the
304 state it was protecting inconsistent. The mutex is now owned by the caller.
305 .RE

307 .sp
308 .LP
309 These functions may fail if:
310 .sp
311 .ne 2

```

```

312 .na
313 \fB\fbefault\fR\fR
314 .ad
315 .RS 10n
316 The \fIcond\fR, \fIattr\fR, \fIcvp\fR, \fIarg\fR, \fIabstime\fR, or \fImutex\fR
317 argument points to an illegal address.
318 .RE

320 .sp
321 .ne 2
322 .na
323 \fB\fbEINVAL\fR\fR
324 .ad
325 .RS 10n
326 Invalid argument. For \fBcond_init()\fR, \fItype\fR is not a recognized type.
327 For \fBcond_timedwait()\fR, the number of nanoseconds is greater than or equal
328 to 1,000,000,000.
329 .RE

331 .SH EXAMPLES
345 .LP
332 \fBExample 1\fR Use \fBcond_wait()\fR in a loop to test some condition.
333 .sp
334 .LP
335 The \fBcond_wait()\fR function is normally used in a loop testing some
349 The \fBcond_wait()\fR function is normally used in a loop testing some
336 condition, as follows:

338 .sp
339 .in +2
340 .nf
341 (void) mutex_lock(mp);
342 while (cond == FALSE) {
343     (void) cond_wait(cvp, mp);
344 }
345 unchanged portion omitted
397 (void) mutex_unlock(mp);
398 .fi
399 .in -2

401 .SH ATTRIBUTES
416 .sp
417 .LP
402 See \fBattributes\fR(5) for descriptions of the following attributes:
403 .sp

405 .sp
406 .TS
407 box;
408 c | c
409 l | l .
410 ATTRIBUTE TYPE ATTRIBUTE VALUE
411 _
412 MT-Level MT-Safe
413 .TE

415 .SH SEE ALSO
432 .sp
433 .LP
416 \fBfork\fR(2), \fBmmap\fR(2), \fBsetitimer\fR(2), \fBshmop\fR(2),
417 \fBmutex_init\fR(3C), \fBsignal\fR(3C), \fBattributes\fR(5),
418 \fBcondition\fR(5), \fBmutex\fR(5), \fBstandards\fR(5)
419 .SH NOTES
438 .sp
439 .LP
420 If more than one thread is blocked on a condition variable, the order in which

```

421 threads are unblocked is determined by the scheduling policy. When each thread,
422 unblocked as a result of a `\fBcond_signal()\fR` or `\fBcond_broadcast()\fR`,
423 returns from its call to `\fBcond_wait()\fR` or `\fBcond_timedwait()\fR`, the
424 thread owns the mutex with which it called `\fBcond_wait()\fR`,
425 `\fBcond_timedwait()\fR`, or `\fBcond_reltimedwait()\fR`. The thread(s) that are
426 unblocked compete for the mutex according to the scheduling policy and as if
427 each had called `\fBmutex_lock(3C)`.

428 .sp
429 .LP
430 When `\fBcond_wait()\fR` returns the value of the condition is indeterminate and
431 must be reevaluated.

432 .sp
433 .LP
434 The `\fBcond_timedwait()\fR` and `\fBcond_reltimedwait()\fR` functions are similar
435 to `\fBcond_wait()\fR`, except that the calling thread will not wait for the
436 condition to become true past the absolute time specified by `\fIabstime\fR` or
437 the relative time specified by `\fIreltime\fR`. Note that `\fBcond_timedwait()\fR`
438 or `\fBcond_reltimedwait()\fR` might continue to block as it tries to reacquire
439 the mutex pointed to by `\fImp\fR`, which may be locked by another thread. If
440 either `\fBcond_timedwait()\fR` or `\fBcond_reltimedwait()\fR` returns because of a
441 timeout, it returns the error value `\fBETIME\fR`.

```

*****
3621 Sat Feb 15 09:54:06 2020
new/usr/src/man/man9e/awrite.9e
12309 errors in section 9e of the manual
*****
1 \" te
2.\" Copyright (c) 2008, Sun Microsystems, Inc. All Rights Reserved.
3.\" Copyright 1989 AT&T
4.\" The contents of this file are subject to the terms of the Common Development
5.\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
6.\" When distributing Covered Code, include this CDDL HEADER in each file and in
7.TH AWRITE 9E "February 15, 2020"
7.TH AWRITE 9E "Mar 28, 1997"
8.SH NAME
9 awrite \- asynchronous write to a device
10 .SH SYNOPSIS
11 .LP
11 .nf
12 #include <sys/uio.h>
13 #include <sys/aio_req.h>
14 #include <sys/cred.h>
15 #include <sys/ddi.h>
16 #include <sys/sunddi.h>
17
18 \fBint prefix\fR\fBawrite\fR(\fBdev_t\fR \fBidev\fR, \fBstruct aio_req *\fR\fBiaio_
19 \fBintprefix\fR\fBawrite\fR(\fBdev_t\fR \fBidev\fR, \fBstruct aio_req *\fR\fBiaio_
19 \fBcred_t *\fR\fBicred_p\fR);
20 .fi
21
22 .SH INTERFACE LEVEL
24 .sp
25 .LP
23 Solaris \fBDDI\fR-specific (Solaris DDI). This entry point is optional. Drivers
24 that do not support an \fBawrite()\fR entry point should use \fBnodev()\fR(9F)
25 .SH PARAMETERS
29 .sp
26 .ne 2
27 .na
28 \fBidev\fR
29 .ad
30 .RS 12n
31 Device number.
32 .RE
33
34 .sp
35 .ne 2
36 .na
37 \fBiaio_reqp\fR
38 .ad
39 .RS 12n
40 Pointer to the \fBaio_req\fR(9S) structure that describes where the data is
41 stored.
42 .RE
43
44 .sp
45 .ne 2
46 .na
47 \fBicred_p\fR
48 .ad
49 .RS 12n
50 Pointer to the credential structure.
51 .RE
52
53 .SH DESCRIPTION
58 .sp
59 .LP

```

```

54 The driver's \fBawrite()\fR routine is called to perform an asynchronous write.
55 \fBgetminor()\fR(9F) can be used to access the minor number component of the
56 \fBidev\fR argument. \fBawrite()\fR may use the credential structure pointed to
57 by \fBicred_p\fR to check for superuser access by calling \fBdrv_priv()\fR(9F).
58 The \fBawrite()\fR routine may also examine the \fBaio\fR(9S) structure
59 through the \fBaio_req\fR structure pointer, \fBaio_reqp\fR. \fBawrite()\fR
60 must call \fBaphysio()\fR(9F) with the \fBaio_req\fR pointer and a pointer to the
61 driver's \fBstrategy()\fR(9E) routine.
62 .sp
63 .LP
64 No fields of the \fBaio\fR(9S) structure pointed to by \fBaio_req\fR, other
65 than \fBoffset\fR or \fBloffset\fR, may be modified for non-seekable
66 devices.
67 .SH RETURN VALUES
74 .sp
75 .LP
68 The \fBawrite()\fR routine should return \fB0\fR for success, or the
69 appropriate error number.
70 .SH CONTEXT
79 .sp
80 .LP
71 This function is called from user context only.
72 .SH EXAMPLES
83 .LP
73 \fBExample 1\fR Using the \fBawrite()\fR routine:
74 .sp
75 .LP
76 The following is an example of an \fBawrite()\fR routine:
77
78 .sp
79 .in +2
80 .nf
81 static
82 xxawrite(dev_t dev, struct aio_req *aio, cred_t *cred_p)
83 {
84     int instance;
85     struct xxstate *xsp;
86
87     instance = getminor(dev);
88     xsp = ddi_get_soft_state(statep, instance);
89     /*Verify soft state structure has been allocated */
90     if (xsp == NULL)
91         return (ENXIO);
92     return (aphysio(xxstrategy, anocancel, dev, B_WRITE, \e
93     xxminphys, aio));
94 }
95 .fi
96 .in -2
97
98 .SH SEE ALSO
110 .sp
111 .LP
99 \fBwrite()\fR(2), \fBaiowrite()\fR(3C), \fBaread()\fR(9E), \fBread()\fR(9E),
100 \fBstrategy()\fR(9E), \fBwrite()\fR(9E), \fBanocancel()\fR(9F), \fBaphysio()\fR(9F),
101 \fBddi_get_soft_state()\fR(9F), \fBdrv_priv()\fR(9F), \fBgetminor()\fR(9F),
102 \fBminphys()\fR(9F), \fBnodev()\fR(9F), \fBaio_req()\fR(9S), \fBcb_ops()\fR(9S),
103 \fBaio()\fR(9S)
104 .sp
105 .LP
106 \fIWriting Device Drivers\fR
107 .SH BUGS
121 .sp
122 .LP
108 There is no way other than calling \fBaphysio()\fR(9F) to accomplish an
109 asynchronous write.

```

14121 Sat Feb 15 09:54:06 2020

new/usr/src/man/man9e/ddi_ufm.9e

12309 errors in section 9e of the manual

```

1 .\"
2 .\" This file and its contents are supplied under the terms of the
3 .\" Common Development and Distribution License ("CDDL"), version 1.0.
4 .\" You may only use this file in accordance with the terms of version
5 .\" 1.0 of the CDDL.
6 .\"
7 .\" A full copy of the text of the CDDL should have accompanied this
8 .\" source. A copy of the CDDL is also available via the Internet at
9 .\" http://www.illumos.org/license/CDDL.
10 .\"
11 .\"
12 .\" Copyright 2019 Joyent, Inc.
13 .\"
14 .Dd February 15, 2020
14 .Dd Apr 30, 2019
15 .Dt DDI_UFM 9E
16 .Os
17 .Sh NAME
18 .Nm ddi_ufm ,
19 .Nm ddi_ufm_op_nimages ,
20 .Nm ddi_ufm_op_fill_image ,
21 .Nm ddi_ufm_op_fill_slot ,
22 .Nm ddi_ufm_op_getcaps
23 .Nd DDI upgradable firmware module entry points
24 .Sh SYNOPSIS
25 .Vt typedef struct ddi_ufm_handle ddi_ufm_handle_t
26 .Vt typedef struct ddi_ufm_ops ddi_ufm_ops_t
27 .In sys/ddi_ufm.h
28 .Ft int
29 .Fo ddi_ufm_op_getcaps
30 .Fa "ddi_ufm_handle_t *uhp"
31 .Fa "void *drv_arg"
32 .Fa "ddi_ufm_cap_t *caps"
33 .Fc
34 .Ft int
35 .Fo ddi_ufm_op_nimages
36 .Fa "ddi_ufm_handle_t *uhp"
37 .Fa "void *drv_arg"
38 .Fa "uint_t *nimgp"
39 .Fc
40 .Ft int
41 .Fo ddi_ufm_op_fill_image
42 .Fa "ddi_ufm_handle_t *uhp"
43 .Fa "void *drv_arg"
44 .Fa "uint_t imgid"
45 .Fa "ddi_ufm_image_t *uip"
46 .Fc
47 .Ft int
48 .Fo ddi_ufm_op_fill_slot
49 .Fa "ddi_ufm_handle_t *uhp"
50 .Fa "void *drv_arg"
51 .Fa "uint_t imgid"
52 .Fa "uint_t slotid"
53 .Fa "ddi_ufm_slot_t *usp"
54 .Fc
55 .Sh INTERFACE LEVEL
56 .Sy Evolving - This interface is evolving still in illumos. API and ABI stabilit
57 .Sh PARAMETERS
58 .Bl -tag -width Fa
59 .It Fa uhp
60 A handle corresponding to the device's UFM handle.

```

```

61 This is the same value as returned in
62 .Xr ddi_ufm_init 9F .
63 .It Fa drv_arg
64 This is a private value that the driver passed in when calling
65 .Xr ddi_ufm_init 9F .
66 .It Fa nimgp
67 A pointer that the driver should set with a number of images.
68 .It Fa nslotp
69 A pointer that the driver should set with a number of slots.
70 .It Fa imgid
71 An integer indicating which image information is being requested for.
72 .It Fa uip
73 An opaque pointer that represents a UFM image.
74 .It Fa slotid
75 An integer indicating which slot information is being requested for.
76 .It Fa usp
77 An opaque pointer that represents a UFM slot.
78 .El
79 .Sh DESCRIPTION
80 Upgradable firmware modules (UFM) are a potential component of many
81 devices.
82 These interfaces aim to provide a simple series of callbacks
83 for a device driver to implement such that it is easy to report
84 information and in the future, manipulate firmware modules.
85 .Ss UFM Background
86 UFM's may come in different flavors and styles ranging from a
87 firmware blob, to an EEPROM image, to microcode, and more.
88 Take for example a hard drive.
89 While it is a field replaceable unit (FRU), it also contains some amount
90 of firmware that manages the drive which can be updated independently of
91 replacing the drive.
92 .Pp
93 The motherboard often has a UFM in the form of the BIOS or UEFI.
94 The Lights out management controller on a system has a UFM, which is usually
95 the entire system image.
96 CPUs also have a UFM in the form of microcode.
97 .Pp
98 An important property of a UFM is that it is a property of the device
99 itself.
100 For example, many WiFi device drivers are required to send a binary blob of
101 firmware to the device after every reset.
102 Because these images are not properties of the device and must be upgraded by
103 either changing the device driver or related system files, we do not consider
104 these UFM's.
105 .Pp
106 There are also devices that have firmware which is a property of the
107 device, but may not be upgradable from the running OS.
108 This may be because the vendor doesn't have tooling to upgrade the image or
109 because the firmware image itself cannot be upgraded in the field at all.
110 For example, a YubiKey has a firmware image that's burned into it in the
111 factory, but there is no way to change the firmware on it short of
112 replacing the device in its entirety.
113 However, because these images are a permanent part of the device, we also
114 consider them a UFM.
115 .Ss Images and Slots
116 A device that supports UFM's is made up of one or more distinct firmware
117 images.
118 Each image has its own unique purpose.
119 For example, a motherboard may have both a BIOS and a CPLD image, each of which
120 has independent firmware revisions.
121 .Pp
122 A given image may have a number of slots.
123 A slot represents a particular version of the image.
124 Only one slot can be active at a given time.
125 Devices support slots such that a firmware image can be downloaded
126 to the device without impacting the current device if it fails half-way

```

127 through.
 128 The slot that's currently in use is referred to as the
 129 .Em active
 130 slot.
 131 .Pp
 132 The various entry points are designed such that all a driver has to do
 133 is provide information about the image and its slots to the kernel, it
 134 does not have to wrangle with how that is marshalled to users and the
 135 appearance of those structures.
 136 .Ss Registering with the UFM Subsystem
 137 During a device driver's
 138 .Xr attach 9E
 139 entry point, a device driver should register with the UFM subsystem by
 140 filling out a UFM operations vector and then calling
 141 .Xr ddi_ufm_init 9F .
 142 The driver may pass in a value, usually a pointer to its soft state
 143 pointer, which it will then receive when its subsequent entry points are
 144 called.
 145 .Pp
 146 Once the driver has finished initializing, it must call
 147 .Xr ddi_ufm_update 9F
 148 to indicate that the driver is in a state where it's ready to receive
 149 calls to the entry points.
 150 .Pp
 151 The various UFM entry points may be called from an arbitrary kernel
 152 context.
 153 However, they will only ever be called from a single thread at
 154 a given time.
 155 .Ss UFM operations vector
 156 The UFM operations vector is a structure that has the following members:
 157 .Bd -literal -offset indent
 158 typedef struct ddi_ufm_ops {
 159 int (*ddi_ufm_op_nimages)(ddi_ufm_handle_t *uhp, void *arg,
 160 uint_t *nimgp);
 161 int (*ddi_ufm_op_fill_image)(ddi_ufm_handle_t *uhp, void *arg,
 162 uint_t imgid, ddi_ufm_image_t *img);
 163 int (*ddi_ufm_op_fill_slot)(ddi_ufm_handle_t *uhp, void *arg,
 164 int imgid, ddi_ufm_image_t *img, uint_t slotid,
 165 ddi_ufm_slot_t *slotp);
 166 int (*ddi_ufm_op_getcaps)(ddi_ufm_handle_t *uhp, void *arg,
 167 ddi_ufm_cap_t *caps);
 168 } ddi_ufm_ops_t;
 169 .Ed
 170 .Pp
 171 The
 172 .Fn ddi_ufm_op_nimages
 173 entry point is optional.
 174 If a device only has a single image, then there is no reason to implement the
 175 .Fn ddi_ufm_op_nimages entry point.
 176 The system will assume that there is only a single image.
 177 .Pp
 178 Slots and images are numbered starting at zero.
 179 If a driver indicates support for multiple images or slots then the images
 180 or slots will be numbered sequentially going from 0 to the number of images or
 181 slots minus one.
 182 These values will be passed to the various entry points to indicate which image
 183 and slot the system is interested in.
 184 It is up to the driver to maintain a consistent view of the images and slots
 185 for a given UFM.
 186 .Pp
 187 The members of this structure should be filled in the following ways:
 188 .Bl -tag -width Fn
 189 .It Fn ddi_ufm_op_nimages
 190 The
 191 .Fn ddi_ufm_op_nimages
 192 entry point is an optional entry point that answers the question of how

193 many different, distinct firmware images are present on the device.
 194 Once the driver determines how many are present, it should set the value in
 195 .Fa nimgp to the determined value.
 196 .Pp
 197 It is legal for a device to pass in zero for this value, which indicates
 198 that there are none present.
 199 .Pp
 200 Upon successful completion, the driver should return
 201 .Sy 0 .
 202 Otherwise, the driver should return the appropriate error number.
 203 For a full list of error numbers, see
 204 .Xr Intro 2 .
 205 Common values are:
 206 .Bl -tag -width Er -offset width
 207 .It Er EIO
 208 An error occurred while communicating with the device to determine the
 209 number of firmware images.
 210 .El
 211 .It Fn ddi_ufm_op_fill_image
 212 The
 213 .Fn ddi_ufm_op_fill_image
 214 entry point is used to fill in information about a given image.
 215 The value in
 216 .Fa imgid
 217 is used to indicate which image the system is asking to fill
 218 information about.
 219 If the driver does not recognize the image ID in
 220 .Fa imgid
 221 then it should return an error.
 222 .Pp
 223 The
 224 .Ft ddi_ufm_image_t
 225 structure passed in
 226 .Fa uip
 227 is opaque.
 228 To fill in information about the image, the driver should call the functions
 229 described in
 230 .Xr ddi_ufm_image 9F .
 231 .Pp
 232 The driver should call the
 233 .Xr ddi_ufm_image_set_desc 9F
 234 function to set a description of the image which indicates its purpose.
 235 This should be a human-readable string.
 236 The driver may also set any ancillary data that it deems may be useful with the
 237 .Xr ddi_ufm_image_set_misc 9F function.
 238 This function takes an nvlist, allowing the driver to set arbitrary keys and val
 239 .Pp
 240 Once the driver has finished setting all of the information about the
 241 image then the driver should return
 242 .Sy 0 .
 243 Otherwise, the driver should return the appropriate error number.
 244 For a full list of error numbers, see
 245 .Xr Intro 2 .
 246 Common values are:
 247 .Bl -tag -width Er -offset width
 248 .It Er EIO
 249 The image indicated by
 250 .Fa imgid
 251 is unknown.
 252 .It Er EIO
 253 An error occurred talking to the device while trying to fill out
 254 firmware image information.
 255 .It Er ENOMEM
 256 The driver was unable to allocate memory while filling out image
 257 information.
 258 .El

```

259 .It Fn ddi_ufm_op_fill_slot
260 The
261 .Fn ddi_ufm_op_fill_slot
262 function is used to fill in information about a specific slot for a
263 specific image.
264 The value in
265 .Fa imgid
266 indicates the image the system wants slot information for and the value
267 in
268 .Fa slotid
269 indicates which slot of that image the system is interested in.
270 If the device driver does not recognize the value in either or
271 .Fa imgid
272 or
273 .Fa slotid ,
274 then it should return an error.
275 .Pp
276 The
277 .Ft ddi_ufm_slot_t
278 structure passed in
279 .Fa usp
280 is opaque.
281 To fill in information about the image the driver should call the functions
282 described in
283 .Xr ddi_ufm_slot 9F .
284 .Pp
285 The driver should call the
286 .Xr ddi_ufm_slot_set_version 9F
287 function to indicate the version of the UFM.
288 The version is a device-specific character string.
289 It should contain the current version of the UFM as a human can understand it
290 and it should try to match the format used by device vendor.
291 .Pp
292 The
293 .Xr ddi_ufm_slot_set_attrs 9F
294 function should be used to set the attributes of the UFM slot.
295 These attributes include the following enumeration values:
296 .Bl -tag -width Dv
297 .It Dv DDI_UFM_ATTR_READABLE
298 This attribute indicates that the firmware image in the specified slot
299 may be read, even if the device driver does not currently support such
300 functionality.
301 .It Dv DDI_UFM_ATTR_WRITEABLE
302 This attributes indicates that the firmware image in the specified slot
303 may be updated, even if the driver does not currently support such
304 functionality.
305 .It Dv DDI_UFM_ATTR_ACTIVE
306 This attributes indicates that the firmware image in the specified slot
307 is the active
308 .Pq i.e. currently running
309 firmware.
310 Only one slot should be marked active.
311 .It Dv DDI_UFM_ATTR_EMPTY
312 This attributes indicates that the specified slot does not currently contain
313 any firmware image.
314 .El
315 .Pp
316 Finally, if there are any device-specific key-value pairs that form
317 useful, ancillary data, then the driver should assemble an nvlist and
318 pass it to the
319 .Xr ddi_ufm_set_misc 9F
320 function.
321 .Pp
322 Once the driver has finished setting all of the information about the
323 slot then the driver should return
324 .Sy 0 .

```

```

325 Otherwise, the driver should return the appropriate error number.
326 For a full list of error numbers, see
327 .Xr Intro 2 .
328 Common values are:
329 .Bl -tag -width Er -offset width
330 .It Er EINVAL
331 The image or slot indicated by
332 .Fa imgid
333 and
334 .Fa slotid
335 is unknown.
336 .It Er EIO
337 An error occurred talking to the device while trying to fill out
338 firmware slot information.
339 .It Er ENOMEM
340 The driver was unable to allocate memory while filling out slot
341 information.
342 .El
343 .It Fn ddi_ufm_op_getcaps
344 The
345 .Fn ddi_ufm_op_getcaps
346 function is used to indicate which DDI UFM capabilities are supported by this
347 driver instance.
348 Currently there is only a single capability
349 .Pq DDI_UFM_CAP_REPORT
350 which indicates that the driver is capable of reporting UFM information for this
351 instance.
352 Future UFM versions may add additional capabilities such as the ability to
353 obtain a raw dump of the firmware image or to upgrade the firmware.
354 .Pp
355 The driver should indicate the supported capabilities by setting the value in
356 the
357 .Ft caps
358 parameter.
359 Once the driver has populated
360 .Ft caps
361 with an appropriate value, then the driver should return
362 .Sy 0 .
363 Otherwise, the driver should return the appropriate error number.
364 For a full list of error numbers, see
365 .Xr Intro 2 .
366 Common values are:
367 .Bl -tag -width Er -offset width
368 .It Er EIO
369 An error occurred talking to the device while trying to discover firmware
370 capabilities.
371 .It Er ENOMEM
372 The driver was unable to allocate memory.
373 .El
374 .El
375 .Ss Caching and Updates
376 The system will fetch firmware and slot information on an as-needed
377 basis.
378 Once it obtains some information, it may end up caching this information on
379 behalf of the driver.
380 Whenever the driver believes that something could have changed -- it need know
381 that it has -- then the driver must call
382 .Xr ddi_ufm_update 9F .
383 .Ss Locking
384 All UFM operations on a single UFM handle will always be run serially.
385 However, the device driver may still need to apply adequate locking to
386 its structure members as other may be accessing the same data structure
387 or trying to communicate with the device.
388 .Ss Unregistering from the UFM subsystem
389 When a device driver is detached, it should unregister from the UFM

```



```
390 subsystem.
391 To do so, the driver should call
392 .Xr ddi_ufm_fini 9F .
393 By the time this function returns, the driver is guaranteed that no UFM
394 entry points will be called.
395 However, if there are outstanding UFM related activity, the function will
396 block until it is terminated.
397 .Ss ioctl Interface
398 Userland consumers can access UFM information via a set of ioctls that are
399 implemented by the
400 .Xr ufm 7D
401 driver.
402 .Sh CONTEXT
403 The various UFM entry points that a device driver must implement will
404 always be called from
405 .Sy kernel
406 context.
407 .Sh SEE ALSO
408 .Xr Intro 2 ,
409 .Xr ufd 7D ,
410 .Xr attach 9E ,
411 .Xr ddi_ufm_fini 9F ,
412 .Xr ddi_ufm_image 9F ,
413 .Xr ddi_ufm_image_set_desc 9F ,
414 .Xr ddi_ufm_image_set_misc 9F ,
415 .Xr ddi_ufm_image_set_nslots 9F ,
416 .Xr ddi_ufm_init 9F ,
417 .Xr ddi_ufm_slot 9F ,
418 .Xr ddi_ufm_slot_set_attrs 9F ,
419 .Xr ddi_ufm_slot_set_misc 9F ,
420 .Xr ddi_ufm_slot_set_version 9F ,
421 .Xr ddi_ufm_update 9F
```

```

*****
12469 Sat Feb 15 09:54:06 2020
new/usr/src/man/man9e/mac_capab_transceiver.9e
12309 errors in section 9e of the manual
*****

```

```

1 .\"
2 .\" This file and its contents are supplied under the terms of the
3 .\" Common Development and Distribution License ("CDDL"), version 1.0.
4 .\" You may only use this file in accordance with the terms of version
5 .\" 1.0 of the CDDL.
6 .\"
7 .\" A full copy of the text of the CDDL should have accompanied this
8 .\" source. A copy of the CDDL is also available via the Internet at
9 .\" http://www.illumos.org/license/CDDL.
10 .\"
11 .\"
12 .\" Copyright (c) 2017, Joyent, Inc.
13 .\"
14 .Dd February 15, 2020
15 .Dt MAC_CAPAB_TRANSCEIVER 9E
16 .Os
17 .Sh NAME
18 .Nm mac_capab_transceiver ,
19 .Nm mct_info ,
20 .Nm mct_read
21 .Nd MAC capability for networking transceivers
22 .Sh SYNOPSIS
23 .In sys/mac_provider.h
24 .Vt typedef struct mac_capab_transceiver mac_capab_transceiver_t;
25 .Ft int
26 .Fo "mct_info"
27 .Fa "void *driver"
28 .Fa "uint_t id"
29 .Fa "mac_transceiver_info_t *infop"
30 .Fc
31 .Ft int
32 .Fo mct_read
33 .Fa "void *driver"
34 .Fa "uint_t id"
35 .Fa "uint_t page"
36 .Fa "void *buf"
37 .Fa "size_t nbytes"
38 .Fa "off_t offset"
39 .Fa "size_t *nread"
40 .Fc
41 .Sh INTERFACE LEVEL
42 .Sy Volatile -
43 This interface is still evolving in illumos.
44 API and ABI stability is
45 not guaranteed.
46 .Sh PARAMETERS
47 .Bl -tag -width Fa
48 .It Fa driver
49 A pointer to the driver's private data that was passed in via the
50 .Sy m_pdata
51 member of the
52 .Xr mac_register 9S
53 structure to the
54 .Xr mac_register 9F
55 function.
56 .It Fa id
57 An integer value indicating which transceiver is being inquired about.
58 .It Fa infop
59 An opaque structure which is used to set information about the
60 transceiver.

```

```

61 .It Fa page
62 A value that indicates which page from the i2c bus is being requested.
63 .It Fa buf
64 A pointer to which data should be written to when reading from the
65 device.
66 .It Fa nbytes
67 A value indicating the number of bytes being asked to read into
68 .Fa buf .
69 .It Fa offset
70 A value indicating the offset into the page to start reading data.
71 .It Fa nread
72 A value to be updated by the driver with the number of successfully read
73 bytes.
74 .El
75 .Sh DESCRIPTION
76 The
77 .Sy MAC_CAPAB_TRANSCEIVER
78 capability allows for GLDv3 networking device drivers to provide
79 information to the system about their transceiver.
80 Implementing this capability is optional.
81 For more information on how to handle capabilities and how to indicate
82 that a capability is not supported, see
83 .Xr mc_getcapab 9E .
84 .Pp
85 This capability should be implemented if the device in question supports
86 a Small Form Factor (SFF) transceiver.
87 These are more commonly known by names such as SFP, SFP+, SFP28, QSFP+,
88 and QSFP28.
89 This interface does not apply to traditional copper Ethernet phys.
90 These transceivers provide standardized information over the i2c bus at
91 specific pages.
92 .Ss Supported Standards
93 .Bl -tag -width Sy
94 .It Sy INF-8074
95 The
96 .Sy INF-8084
97 standard was the original multiple source agreement (MSA) for SFP
98 devices.
99 It proposed the original series of management pages at i2c page 0xa0.
100 This page contained up to 512 bytes, however, only the first
101 96 bytes are standardized.
102 Bytes 97 to 127 are reserved for the vendor.
103 The remaining bytes are reserved by the specification.
104 The management page was subsequently adopted by SFP+ devices.
105 .It Sy SFF-8472
106 The
107 .Sy SFF-8472
108 standard extended the original SFP MSA.
109 This standard added a second i2c page at 0xa2, while maintaining the
110 original page at 0xa0.
111 The page at 0xa0 is now explicitly 256 bytes.
112 The page at 0xa2 is also 256 bytes.
113 This standard was also adopted for all SFP28 parts, which are commonly
114 used in transceivers for 25 Gb/s Ethernet.
115 .It Sy SFF-8436
116 The
117 .Sy SFF-8436
118 standard was developed for QSFP+ transceivers, which involve the
119 bonding of 4 SFP+ links.
120 QSFP+ is commonly used in the transceivers for 40 Gb/s Ethernet.
121 This standard uses i2c page 0xa0 for read-only identification purposes.
122 The lower half of the page is used for control, while the upper 128
123 bytes is similar to the
124 .Sy INF-8084
125 and
126 .Sy SFF-8472

```

```

127 standards.
128 .It Sy SFF-8636
129 The
130 .Sy SFF-8636
131 standard is a common management standard which is shared between both
132 SAS and QSFP+ 28 Gb/s transceivers.
133 The latter transceiver is commonly found in 100 Gb/s Ethernet.
134 The transceiver's memory map is similar to that found in the
135 .Sy SFF-8436
136 specification.
137 The identification information is found in the upper 128
138 bytes of page 0xa0, while the lower part of the page is used for
139 control, among other purposes.
140 .El
141 .Pp
142 The following table summarizes the above information.
143 .Bl -column "Sy SFF-8636" "1 Gb/s, 10 Gb/s, 25 Gb/s" "256 bytes" "0xa0, 0xa2" -o
144 .Em "Standard" Ta Em Speeds Ta Em Size Ta Em i2c pages
145 .It INF-8074 Ta 1 Gb/s, 10 Gb/s Ta 128 bytes Ta 0xa0
146 .It SFF-8472 Ta 1 Gb/s, 10 Gb/s, 25 GB/s Ta 512 bytes Ta 0xa0, 0xa2
147 .It SFF-8436 Ta 40 Gb/s Ta 256 bytes Ta 0xa0
148 .It SFF-8636 Ta 100 Gb/s Ta 256 bytes Ta 0xa0
149 .El
150 .Ss MAC Capability Structure
151 When the device driver's
152 .Xr mc_getcapab 9E
153 function entry point is called with the capability requested set to
154 .Sy MAC_CAPAB_TRANSCEIVER ,
155 then the value of the capability structure is the following structure:
156 .Bd -literal -offset indent
157 typedef struct mac_capab_transceiver {
158     uint_t mct_flags;
159     uint_t mct_ntransceivers;
160     uint_t mct_ntransceiverses;
161     int (*mct_info)(void *driver, uint_t id,
162                 mac_transceiver_info_t *infp),
163     int (*mct_read)(void *driver, uint_t id, uint_t page,
164                 void *buf, size_t nbytes, off_t offset,
165                 size_t *nread)
166 } mac_capab_transceiver_t;
167 .Ed
168 .Pp
169 If the device driver supports the
170 .Sy MAC_CAPAB_TRANSCEIVER
171 capability, it should fill in this structure, based on the following
172 rules:
173 .Bl -tag -width Sy
174 .It Sy mct_flags
175 .Vt mct_flags
176 member is used to negotiate extensions with the driver.
177 MAC will set the value of
178 .Vt mct_flags
179 to include all of the currently known extensions.
180 The driver should intersect this list with the set that they actually
181 support.
182 At this time, no such features are defined and the driver should set the
183 member to
184 .Sy 0 .
185 .It Sy mct_ntransceivers
186 The value of
187 .Sy mct_ntransceivers
188 indicates the number of transceivers present in the device.
189 For most devices, it is expected that this value will be set to one.
190 However, some devices do support multiple transceivers and PHYs that

```

```

191 show up behind a single logical MAC.
192 .Pp
193 It is expected that this value will not change across the lifetime of
194 the device being attached.
195 It is important to remember that this represents the total possible
196 number of transceivers in the device, not how many are currently present
197 and powered on.
198 .Pp
199 The number of transceivers will influence the
200 .Fa id
201 argument used in the
202 .Fn mct_info
203 and
204 .Fn mct_read
205 entry points.
206 The transceiver IDs will start at zero and go to the value of
207 .Fa mct_ntransceivers - 1
208 It is up to the driver to keep the mapping between actual transceivers
209 and the transceiver identifiers consistent.
210 .It Sy mct_info
211 The
212 .Fn mct_info
213 entry point is used to set basic information about the transceiver.
214 This entry point is
215 .Em required .
216 If the device driver cannot implement this entry point, then it should
217 not indicate that it supports the capability.
218 .Pp
219 The
220 .Fn mct_info
221 entry point should fill in information about the transceiver with an
222 identifier of
223 .Fa id .
224 See the description above of
225 .Sy mct_ntransceivers
226 for more information on how the IDs are determined.
227 .Pp
228 The driver should then proceed to fill in basic information by calling
229 the functions described in the section
230 .Sx Information Functions .
231 After successfully calling all of the functions, the driver should
232 return
233 .Sy 0 .
234 Otherwise, it should return the appropriate error number.
234 Othewise, it should return the appropriate error number.
235 For a full list of error numbers, see
236 .Xr Intro 2 .
237 Common values are:
238 .Bl -tag -width Er -offset width
239 .It Er EINVAL
240 The transceiver identifier
241 .Fa id
242 was invalid.
243 .It Er ENOTSUP
244 This instance of the devices does not support a transceiver.
245 For example, a device which sometimes has copper PHYs and therefore this
246 instance does not have any PHYs.
247 .It Er EIO
248 An error occurred while trying to read device registers.
249 For example, an FM-aware device had an error.
250 .El
251 .It Sy mct_read
252 The
253 .Fn mct_read
254 function is used to read information from a transceiver's i2c bus.
255 The

```

```

256 .Fn mct_read
257 entry point is an
258 .Em optional
259 entry point.
260 .Pp
261 The transceiver should first check the value of
262 .Fa id ,
263 which indicates which transceiver information is being requested.
264 See the description above of
265 .Sy mct_ntransceivers
266 for more information on how the IDs are determined.
267 .Pp
268 The driver should try to read up to
269 .Fa nbytes
270 of data from the i2c bus at page
271 .Fa page .
272 The driver should start reading at offset
273 .Fa offset .
274 Finally, it should update the value in
275 .Fa nread
276 with the number of bytes written to the buffer
277 .Fa buf .
278 .Pp
279 If for some reason the driver cannot read all of the requested bytes,
280 that is acceptable.
281 Instead it should perform a short read.
282 This may occur because the transceiver does not allow reads at a
283 requested region or the region is shorter than is common for most
284 devices.
285 .Pp
286 Upon successful completion, the driver should ensure that
287 .Fa nread
288 has been updated and then return
289 .Sy 0 .
290 Otherwise, the driver should return the appropriate error number.
291 For
292 a full list of error numbers, see
293 .Xr Intro 2 .
294 Common values are:
295 .Bl -tag -width Er -offset width
296 .It Er EINVAL
297 The value of
298 .Fa id
299 represented an invalid transceiver identifier.
300 The transceiver i2c page
301 .Fa page
302 is not valid for this type of device.
303 The value of
304 .Fa offset
305 is beyond the range supported for this
306 .Fa page .
307 .It Er EIO
308 An error occurred while trying to read the device i2c pages.
309 .El
310 .El
311 .Ss Transceiver Information Functions
312 The
313 .Fn mct_info
314 entry point is the primary required entry point for a device driver
315 which supports this capability.
316 The information structure is opaque to the device driver.
317 Instead, a series of informational functions is
318 available to the device driver to call on the transceiver.
319 The device drivers should try to call and fill in as many of these as
320 possible.
321 There are two different properties that a driver can set:

```

```

322 .Bl -enum -offset indent
323 .It
324 Whether the transceiver is present.
325 .It
326 Whether the transceiver is usable.
327 .El
328 .Pp
329 To set whether or not the transceiver is present, the driver should call
330 .Xr mac_transceiver_info_set_present 9F .
331 This is used to indicate whether the transceiver is plugged in or not.
332 If the transceiver is a part of the NIC, then this function should
333 always be called with the value set to
334 .Dv B_TRUE .
335 .Pp
336 Finally, the driver has the ability to provide information about whether
337 or not the transceiver is usable or not.
338 A transceiver may be present, but not usable, if the hardware and
339 firmware support a limited number of transceivers.
340 To set this information, the driver should call
341 .Xr mac_transceiver_info_set_usable 9F .
342 If the transceiver is not present, then the driver should not call this
343 function.
344 .Ss Opaque Transceivers
345 Some devices abstract the nature of the transceiver and do not allow
346 direct access to the transceiver.
347 In this case, if the device driver still has access to enough
348 information to know if the transceiver is at least present, then it
349 should still implement the
350 .Fn mct_info
351 entry point.
352 .Ss Locking and Data Access
353 Calls to get information about the transceivers may come at the same
354 time as general I/O requests to the device to send or receive data.
355 The driver should make sure that reading data from the i2c bus of the
356 transceiver does not interfere with the device's functionality in this
357 regard.
358 Different locks should be used.
359 .Pp
360 On some devices, reading from the transceiver's i2c bus might cause a
361 disruption of service to the device.
362 For example, on some devices a phy reset may be required or come about
363 as a side effect of trying to read the device.
364 If any kind of disruption would be caused, then the driver
365 must not implement the
366 .Fn mct_read
367 entry point.
368 .Sh CONTEXT
369 The various callback functions will be called from
370 .Sy kernel
371 context.
372 These functions will never be called from
373 .Sy interrupt
374 context.
375 .Sh SEE ALSO
376 .Xr Intro 2 ,
377 .Xr mac 9E ,
378 .Xr mc_getcapab 9E ,
379 .Xr mac_register 9F ,
380 .Xr mac_transceiver_info_set_present 9F ,
381 .Xr mac_transceiver_info_set_usable 9F ,
382 .Xr mac_register 9S
383 .Rs
384 .%N INF-8074i
385 .%T SFP (Small Formfactor Pluggable) Interface
386 .%Q SFF Committee
387 .%O Revision 1.0

```

388 .%D May 12, 2001
389 .Re
390 .Rs
391 .%N SFF-8472
392 .%T Diagnostic Monitoring Interface for Optical Transceivers
393 .%O Revision 12.2
394 .%D November 21, 2014
395 .Re
396 .Rs
397 .%N SFF-8436
398 .%T QSFP+ 10 Gbs 4X PLUGGABLE TRANSCEIVER
399 .%O Revision 4.8
400 .%D October 31, 2013
401 .Re
402 .Rs
403 .%N SFF-8636
404 .%T Management Interface for Cabled Environments
405 .%O Revision 2.7
406 .%D January 26, 2016
407 .Re

```

*****
6020 Sat Feb 15 09:54:06 2020
new/usr/src/man/man9e/mc_getprop.9e
12309 errors in section 9e of the manual
*****

```

```

1 .\"
2 .\" This file and its contents are supplied under the terms of the
3 .\" Common Development and Distribution License ("CDDL"), version 1.0.
4 .\" You may only use this file in accordance with the terms of version
5 .\" 1.0 of the CDDL.
6 .\"
7 .\" A full copy of the text of the CDDL should have accompanied this
8 .\" source. A copy of the CDDL is also available via the Internet at
9 .\" http://www.illumos.org/license/CDDL.
10 .\"
11 .\"
12 .\" Copyright 2016 Joyent, Inc.
13 .\"
14 .Dd February 15, 2020
14 .Dd November 15, 2016
15 .Dt MC_GETPROP 9E
16 .Os
17 .Sh NAME
18 .Nm mc_getprop
19 .Nd get device properties
20 .Sh SYNOPSIS
21 .In sys/mac_provider.h
22 .Ft int
23 .Fo prefix_m_getprop
24 .Fa "void *driver"
25 .Fa "const char *pr_name"
26 .Fa "mac_prop_id_t pr_num"
27 .Fa "uint_t pr_valsize"
28 .Fa "void *pr_val"
29 .Fc
30 .Sh INTERFACE LEVEL
31 illumos DDI specific
32 .Sh PARAMETERS
33 .Bl -tag -width Fa
34 .It Fa driver
35 A pointer to the driver's private data that was passed in via the
36 .Sy m_pdata
37 member of the
38 .Xr mac_register 9S
39 structure to the
40 .Xr mac_register 9F
41 function.
42 .It Fa pr_name
43 A null-terminated string that contains the name of the property.
44 .It Fa pr_num
45 A constant that is used to identify the property.
46 .It Fa pr_valsize
47 A value that indicates the size in bytes of
48 .Fa pr_val .
49 .It Fa pr_val
50 A pointer to a
51 .Fa pr_valsize
52 byte buffer that can store the property.
53 .El
54 .Sh DESCRIPTION
55 The
56 .Fn mc_getprop
57 entry point is used to obtain the value of a given device's property and
58 place it into
59 .Fa pr_val .
60 .Pp

```

```

61 When the
62 .Fn mc_getprop
63 entry point is called, the driver needs to first identify the property.
64 The set of possible properties and their meaning is listed in the
65 .Sx PROPERTIES
66 section of
67 .Xr mac 9E .
68 It should identify the property based on the value of
69 .Fa pr_num .
70 Most drivers will use a
71 .Sy switch
72 statement and for any property that it supports it should then check if
73 the value in
74 .Fa pr_valsize
75 is sufficient for the property, comparing it to the minimum size
76 listed for the property in
77 .Xr mac 9E .
78 If it is not, then it should return an error.
79 Otherwise, it should copy the property's value into
80 .Fa pr_val .
81 When an unknown or unsupported
82 property is encountered, generally the
83 .Sy default
84 case of the switch statement, the device driver should return an error.
85 .Pp
86 The special property
87 .Sy MAC_PROP_PRIVATE
88 indicates that this is a device driver specific private property.
89 The device driver must then look at the value of the
90 .Fa pr_name
91 argument and use
92 .Xr strcmp 9F
93 on it, comparing it to each of its private (bounded-size) properties to
94 identify which one it is.
95 .Pp
96 At this time, private properties are limited to being string based properties.
97 If other types of property values are used, they will not be rendered
98 correctly by
99 .Xr dladm 1M .
100 .Pp
101 The device
102 driver can access its device soft state by casting the
103 .Fa device
104 pointer to the appropriate structure.
105 As this may be called while other operations are ongoing, the device driver
106 should employ the appropriate locking while reading the properties.
107 .Sh CONTEXT
108 The
109 .Fn mc_getprop
110 function is generally called from
111 .Sy kernel
112 context.
113 .Sh RETURN VALUES
114 Upon successful completion, the device driver should have copied the
115 value of the property into
116 .Fa pr_val
117 and return
118 .Sy 0 .
119 Otherwise, a positive error should be returned to indicate failure.
120 .Sh EXAMPLES
121 The following example shows how a device driver might structure its
122 .Fn mc_getprop
123 entry point.
124 .Bd -literal
125 #include <sys/mac_provider.h>

```

```
127 /*
128 * Note, this example merely shows the structure of this function.
129 * Different devices will manage their state in different ways. Like other
130 * examples, this assumes that the device has state in a structure called
131 * example_t and that there is a lock which keeps track of that state.
132 */
133 static char *example_priv_props[] = {
134     "_rx_intr_throttle",
135     "_tx_intr_throttle",
136     NULL
137 };
138
139 static int
140 example_m_getprop_private(example_t *ep, const char *pr_name, uint_t pr_valsize,
141     void *pr_val)
142 {
143     uint32_t val;
144
145     ASSERT(MUTEX_HELD(&ep->ep_lock));
146     if (strcmp(pr_name, example_priv_props[0] == 0) {
147         val = ep->ep_rx_itr;
148     } else if (strcmp(pr_name, example_priv_props[1] == 0) {
149     } else if (strcmp(pr_name, exampe_priv_props[1] == 0) {
150         val = ep->ep_tx_itr;
151     } else {
152         return (ENOTSUP);
153     }
154
155     /*
156     * Due to issues in the GLDv3, these must be returned as string
157     * properties.
158     */
159     if (snprintf(pr_val, pr_valsize, "%d", val) >= pr_valsize)
160         return (EOVERFLOW);
161
162     return (0);
163 }
164
165 _____unchanged_portion_omitted_____
```

```

*****
6951 Sat Feb 15 09:54:06 2020
new/usr/src/man/man9e/mc_setprop.9e
12309 errors in section 9e of the manual
*****

```

```

1 .\"
2 .\" This file and its contents are supplied under the terms of the
3 .\" Common Development and Distribution License ("CDDL"), version 1.0.
4 .\" You may only use this file in accordance with the terms of version
5 .\" 1.0 of the CDDL.
6 .\"
7 .\" A full copy of the text of the CDDL should have accompanied this
8 .\" source. A copy of the CDDL is also available via the Internet at
9 .\" http://www.illumos.org/license/CDDL.
10 .\"
11 .\"
12 .\" Copyright 2016 Joyent, Inc.
13 .\"
14 .Dd February 15, 2020
14 .Dd June 02, 2016
15 .Dt MC_SETPROP 9E
16 .Os
17 .Sh NAME
18 .Nm mc_setprop
19 .Nd set device properties
20 .Sh SYNOPSIS
21 .In sys/mac_provider.h
22 .Ft int
23 .Fo prefix_m_setprop
24 .Fa "void *driver"
25 .Fa "const char *pr_name"
26 .Fa "mac_prop_id_t pr_num"
27 .Fa "uint_t pr_valsize"
28 .Fa "const void *pr_val"
29 .Fc
30 .Sh INTERFACE LEVEL
31 illumos DDI specific
32 .Sh PARAMETERS
33 .Bl -tag -width Fa
34 .It Fa driver
35 A pointer to the driver's private data that was passed in via the
36 .Sy m_pdata
37 member of the
38 .Xr mac_register 9S
39 structure to the
40 .Xr mac_register 9F
41 function.
42 .It Fa pr_name
43 A null-terminated string that contains the name of the property.
44 .It Fa pr_num
45 A constant that is used to identify the property.
46 .It Fa pr_valsize
47 A value that indicates the size in bytes of
48 .Fa pr_val .
49 .It Fa pr_val
50 A pointer to a
51 .Fa pr_valsize
52 byte buffer that contains the new value of the property.
53 .El
54 .Sh DESCRIPTION
55 The
56 .Fn mc_setprop
57 entry point is used to set the value of a given device's property from
58 the copy stored in
59 .Fa pr_val .
60 .Pp

```

```

61 When the
62 .Fn mc_setprop
63 entry point is called, the driver needs to first identify the property.
64 The set of possible properties and their meaning is listed in the
65 .Sx PROPERTIES
66 section of
67 .Xr mac 9E .
68 It should identify the property based on the value of
69 .Fa pr_num .
70 Most drivers will use a
71 .Sy switch
72 statement and for any property that it supports it should then check if
73 the value in
74 .Fa pr_valsize
75 is sufficient for the property, comparing it to the minimum size
76 listed for the property in
77 .Xr mac 9E .
78 If it is not, then it should return an error.
79 Otherwise, it should update the property based on the value in
80 .Fa pr_val .
81 When an unknown or unsupported property is encountered, generally the
82 .Sy default
83 case of the switch statement, the device driver should return an error.
84 .Pp
85 The special property
86 .Sy MAC_PROP_PRIVATE
87 indicates that this is a device driver specific private property.
88 The device driver must then look at the value of the
89 .Fa pr_name
90 argument and use
91 .Xr strcmp 9F
92 on it, comparing it to each of its private properties to identify which
93 one it is.
94 .Pp
95 Not all properties are supposed to be writable.
96 Some devices may opt to not allow a property that is designated as read/write to
97 be set.
98 When such a property is encountered, the driver should return the appropriate
99 error.
100 .Pp
101 The device
102 driver can access its device soft state by casting the
103 .Fa device
104 pointer to the appropriate structure.
105 As this may be called while other operations are ongoing, the device driver
106 should employ the appropriate locking while writing the properties.
107 .Sh RETURN VALUES
108 Upon successful completion, the device driver should have copied the
109 value of the property into
110 .Fa pr_val
111 and return
112 .Sy 0 .
113 Otherwise, a positive error should be returned to indicate failure.
114 .Sh EXAMPLES
115 The following examples shows how a device driver might structure its
116 .Fn mc_setprop
116 .Fn mc_setprop
117 entry point.
118 .Bd -literal
119 #include <sys/mac_provider.h>
121 /*
122 * Note, this example merely shows the structure of this function.
123 * Different devices will manage their state in different ways. Like other
124 * examples, this assumes that the device has state in a structure called
125 * example_t and that there is a lock which keeps track of that state.

```



```

126 *
127 * For the purpose of this example, we assume that this device supports 100 Mb,
128 * 1 GB, and 10 Gb full duplex speeds.
129 */

131 static int
132 example_m_setprop(void *arg, const char *pr_name, mac_prop_id_t pr_num,
133 example_m_setprop(void *arg, const char *pr_name, mac_prop_id_t pr_num,
134 uint_t pr_valsize, const void *pr_val)
135 {
136     uint32_t new_mtu;
137     int ret = 0;
138     example_t *ep = arg;

139     mutex_enter(&ep->ep_lock);
140     switch (pr_num) {
141     /*
142      * These represent properties that can never be changed, regardless of
143      * the type of PHY on the device (copper, fiber, etc.)
144      */
145     case MAC_PROP_DUPLEX:
146     case MAC_PROP_SPEED:
147     case MAC_PROP_STATUS:
148     case MAC_PROP_ADV_100FDX_CAP:
149     case MAC_PROP_ADV_1000FDX_CAP:
150     case MAC_PROP_ADV_10GFDX_CAP:
151         ret = ENOTSUP;
152         break;

154     /*
155      * These EN properties are used to control the advertised speeds of the
156      * device. For this example, we assume that this device does not have a
157      * copper phy, at which point auto-negotiation and the speeds in
158      * question cannot be changed. These are called out separately as they
159      * should be controllable for copper based devices or it may need to be
160      * conditional depending on the type of phy present.
161      */
162     case MAC_PROP_EN_100FDX_CAP:
163     case MAC_PROP_EN_1000FDX_CAP:
164     case MAC_PROP_EN_10GFDX_CAP:
165     case MAC_PROP_AUTONEG:
166         ret = ENOTSUP;
167         break;

169     case MAC_PROP_MTU:
170         if (pr_valsize < sizeof (uint32_t)) {
171             ret = EOVERFLOW;
172             break;
173         }
174         bcopy(&new_mtu, pr_val, sizeof (uint32_t));

176         if (new_mtu < ep->ep_min_mtu ||
177             new_mtu > ep->ep_max_mtu) {
178             ret = EINVAL;
179             break;
180         }

182     /*
183      * We first ask MAC to update the MTU before we do anything.
184      * This may fail. It returns zero on success. The
185      * example_update_mtu function does device specific updates to
186      * ensure that the MTU on the device is updated and any internal
187      * data structures are up to date.
188      */
189     ret = mac_maxdsu_update(&ep->ep_mac_hdl, new_mtu);
190     if (ret == 0) {

```

```

191         example_update_mtu(ep, new_mtu);
192     }
193     break;

195     /*
196      * Devices may have their own private properties. If they do, they
197      * should not return ENOTSUP, but instead see if it's a property they
198      * recognize and handle it similar to those above. If it doesn't
199      * recognize the name, then it should return ENOTSUP.
200      */
201     case MAC_PROP_PRIVATE:
202         ret = ENOTSUP;
203         break;

205     default:
206         ret = ENOTSUP;
207         break;
208     }
209     mutex_exit(&ep->ep_lock);

211     return (ret);
212 }

```

unchanged portion omitted

3716 Sat Feb 15 09:54:07 2020

new/usr/src/man/man9e/probe.9e

12309 errors in section 9e of the manual

```

1  \" te
2  .\" Copyright (c) 2000, Sun Microsystems, Inc.
3  .\" All Rights Reserved
4  .\" The contents of this file are subject to the terms of the Common Development
5  .\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
6  .\" When distributing Covered Code, include this CDDL HEADER in each file and in
7  .TH PROBE 9E \"February 15, 2020\"
8  .TH PROBE 9E \"Nov 18, 1992\"
9  .SH NAME
10 probe \- determine if a non-self-identifying device is present
11 .SH SYNOPSIS
12 .LP
13 .nf
14 #include <sys/conf.h>
15 #include <sys/ddi.h>
16 #include <sys/sunddi.h>
17
18 \fBstatic int prefix\fR(\fBprobe\fR(\fBdev_info_t * \fR\fR\fR);
19 \fBstatic int prefix\fR(\fBprobe\fR(\fBdev_info_t * \fR\fR\fR);
20 .fi
21 .SH INTERFACE LEVEL
22 .sp
23 .LP
24 Solaris DDI specific (Solaris DDI). This entry point is required for
25 non-self-identifying devices. You must write it for such devices. For
26 self-identifying devices, \fBnulldev\fR(9F) should be specified in the
27 \fBdev_ops\fR(9S) structure if a probe routine is not necessary.
28 .SH ARGUMENTS
29 .sp
30 .ne 2
31 .na
32 \fB\fR\fR\fR
33 .ad
34 .RS 8n
35 Pointer to the device's \fBdev_info\fR structure.
36 .RE
37
38 .SH DESCRIPTION
39 .sp
40 .LP
41 \fBprobe()\fR determines whether the device corresponding to \fR\fR\fR actually
42 exists and is a valid device for this driver. \fBprobe()\fR is called after
43 \fBidentify\fR(9E) and before \fBattach\fR(9E) for a given \fR\fR\fR. For
44 example, the \fBprobe()\fR routine can map the device registers using
45 \fBddi_map_regs\fR(9F) then attempt to access the hardware using
46 \fBddi_peek\fR(9F) or \fBddi_poke\fR(9F) and determine if the device exists.
47 Then the device registers should be unmapped using \fBddi_unmap_regs\fR(9F).
48 .sp
49 .LP
50 To probe a device that was left powered off after the last \fBdetach()\fR, it
51 might be necessary to power it up. If so, the driver must power up the device
52 by accessing device registers directly. \fBpm_raise_power\fR(9F) will be not be
53 available until \fBattach\fR(9E). The framework ensures that the ancestors of
54 the node being probed and all relevant platform-specific power management
55 hardware is at full power at the time that \fBprobe()\fR is called.
56 .sp
57 .LP
58 \fBprobe()\fR should only probe the device. It should not change any software

```

```

54 state and should not create any software state. Device initialization should be
55 done in \fBattach\fR(9E).
56 .sp
57 .LP
58 For a self-identifying device, this entry point is not necessary. However, if a
59 device exists in both self-identifying and non-self-identifying forms, a
60 \fBprobe()\fR routine can be provided to simplify the driver.
61 \fBddi_dev_is_sid\fR(9F) can then be used to determine whether \fBprobe()\fR
62 needs to do any work. See \fBddi_dev_is_sid\fR(9F) for an example.
63 .SH RETURN VALUES
64 .sp
65 .ne 2
66 .na
67 \fB\fR\fR\fR
68 .ad
69 .RS 23n
70 If the probe was successful.
71 .RE
72
73 .sp
74 .ne 2
75 .na
76 \fB\fR\fR\fR
77 .ad
78 .RS 23n
79 If the probe failed.
80 .RE
81
82 .sp
83 .ne 2
84 .na
85 \fB\fR\fR\fR
86 .ad
87 .RS 23n
88 If the probe was unsuccessful, yet \fBattach\fR(9E) should still be called.
89 .RE
90
91 .sp
92 .ne 2
93 .na
94 \fB\fR\fR\fR
95 .ad
96 .RS 23n
97 If the instance is not present now, but may be present in the future.
98 .RE
99
100 .SH SEE ALSO
101 .sp
102 .LP
103 \fBattach\fR(9E), \fBidentify\fR(9E), \fBddi_dev_is_sid\fR(9F),
104 \fBddi_map_regs\fR(9F), \fBddi_peek\fR(9F), \fBddi_poke\fR(9F),
105 \fBnulldev\fR(9F), \fBdev_ops\fR(9S)

```

```

*****
2623 Sat Feb 15 09:54:07 2020
new/usr/src/man/man9e/usba_hcdi_hub_update.9e
12309 errors in section 9e of the manual
*****

```

```

1 .\"
2 .\" This file and its contents are supplied under the terms of the
3 .\" Common Development and Distribution License ("CDDL"), version 1.0.
4 .\" You may only use this file in accordance with the terms of version
5 .\" 1.0 of the CDDL.
6 .\"
7 .\" A full copy of the text of the CDDL should have accompanied this
8 .\" source. A copy of the CDDL is also available via the Internet at
9 .\" http://www.illumos.org/license/CDDL.
10 .\"
11 .\"
12 .\" Copyright 2016 Joyent, Inc.
13 .\"
14 .Dd February 15, 2020
14 .Dd Dec 20, 2016
15 .Dt USBA_HCDI_HUB_UPDATE 9E
16 .Os
17 .Sh NAME
18 .Nm usba_hcdi_hub_update
19 .Nd USB HCD hub update entry point
20 .Sh SYNOPSIS
21 .In sys/usb/usba/hcdi.h
22 .Ft int
23 .Fo prefix_hcdi_hub_update
24 .Fa "usba_device_t *ud"
25 .Fa "uint8_t nports"
26 .Fa "uint8_t tt"
27 .Fc
28 .Sh INTERFACE LEVEL
29 .Sy Volatile -
30 illumos USB HCD private function
31 .Pp
32 This is a private function that is not part of the stable DDI.
33 It may be removed or changed at any time.
34 .Sh PARAMETERS
35 .Bl -tag -width Fa
36 .It Fa ud
37 Pointer to a USB device structure being updated.
38 See
39 .Xr usba_device 9S
40 for more information.
41 .It Fa nports
42 The number of ports present on the hub.
43 .It Fa tt
44 The value of the Think Time property as defined in the USB
45 specification's hub descriptor.
46 .El
47 .Sh DESCRIPTION
48 The
49 .Fn usba_hcdi_hub_update
50 entry point is an optional entry point for USB host controller drivers.
51 It is used by some controllers to allow them to update information about
52 a device in the controller after a device has been determined to be a
53 hub during enumeration.
54 If a host controller does not need to take any specific action after enumerating
55 a hub, then it should simply set this entry point in the
56 .Xr usba_hcdi_ops 9S
57 structure to
58 .Dv NULL .
59 .Pp
60 The

```

```

61 .Fa nports
62 and
63 .Fa tt
64 members provide relevant information from the device's hub class
65 descriptor which can be used to help program the host controller.
66 Any programming should be performed synchronously and be completed before
67 this function returns.
68 .Pp
69 This function will be called after
70 .Xr usba_hcdi_device_init 9E
71 has been called.
72 Any private data registered with that function will be available.
73 .Pp
74 If this function fails, the enumeration of this device will fail, the
75 hub driver will not attach to this USB device, and all devices plugged
76 into this hub will not be detected by the system.
77 .Sh CONTEXT
78 This function is called from kernel context only.
79 This function is called from kernel context only.
80 .Sh RETURN VALUES
81 Upon successful completion, the
82 .Fn usba_hcdi_hub_update
83 function should return
84 .Sy USB_SUCCESS .
85 Otherwise, it should return the appropriate USB error.
86 If uncertain, use
87 .Sy USB_FAILURE .
88 .Sh SEE ALSO
89 .Xr usba_hcdi_device_init 9E ,
90 .Xr usba_device 9S ,
91 .Xr usba_hcdi_ops 9S

```

5117 Sat Feb 15 09:54:07 2020

new/usr/src/man/man9e/usba_hcdi_pipe_open.9e

12309 errors in section 9e of the manual

```

1 .\"
2 .\" This file and its contents are supplied under the terms of the
3 .\" Common Development and Distribution License ("CDDL"), version 1.0.
4 .\" You may only use this file in accordance with the terms of version
5 .\" 1.0 of the CDDL.
6 .\"
7 .\" A full copy of the text of the CDDL should have accompanied this
8 .\" source. A copy of the CDDL is also available via the Internet at
9 .\" http://www.illumos.org/license/CDDL.
10 .\"
11 .\"
12 .\" Copyright 2016 Joyent, Inc.
13 .\"
14 .Dd February 15, 2020
15 .Dt USBA_HCDI_PIPE_OPEN 9E
16 .Os
17 .Sh NAME
18 .Nm usba_hcdi_pipe_open ,
19 .Nm usba_hcdi_pipe_close
20 .Nd open and close a USB pipe
21 .Sh SYNOPSIS
22 .In sys/usb/usba/hcdi.h
23 .Ft int
24 .Fo prefix_hcdi_pipe_open
25 .Fa "usba_pipe_handle_data_t *ph"
26 .Fa "usb_flags_t usb_flags"
27 .Fc
28 .Ft int
29 .Fo prefix_hcdi_pipe_close
30 .Fa "usba_pipe_handle_data_t *ph"
31 .Fa "usb_flags_t usb_flags"
32 .Fc
33 .Sh INTERFACE LEVEL
34 .Sy Volatile -
35 illumos USB HCD private function
36 .Pp
37 This is a private function that is not part of the stable DDI.
38 It may be removed or changed at any time.
39 .Sh PARAMETERS
40 .Bl -tag -width Fa
41 .It Fa ph
42 A pointer to a USB pipe handle as defined in
43 .Xr usba_pipe_handle_data 9S .
44 .It Fa usb_flags
45 Flags which describe how allocations should be performed.
46 Valid flags are:
47 .Bl -tag -width Sy
48 .It Sy USB_FLAGS_NOSLEEP
49 Do not block waiting for memory.
50 If memory is not available the allocation will fail.
51 .It Sy USB_FLAGS_SLEEP
52 Perform a blocking allocation.
53 If memory is not available, the function will wait until memory is made
54 available.
55 .Pp
56 Note, the request may still fail even if
57 .Sy USB_FLAGS_SLEEP
58 is specified.
59 .El
60 .El

```

```

61 .Sh DESCRIPTION
62 The
63 .Fn usba_hcdi_pipe_open
64 and
65 .Fn usba_hcdi_pipe_close
66 entry points are called by the USB framework whenever a client, or the
67 framework itself, need to open or close a specific pipe.
68 For additional background see
69 .Xr usba_hcdi 9E .
70 .Pp
71 When a pipe is opened, the host controller driver is responsible for
72 preparing the specified endpoint for performing transfers.
73 This may include allocating bandwidth, programming the controller, and more.
74 When the pipe is closed, the host controller driver is responsible for
75 cleaning up any resources that were allocated during the open call.
76 .Pp
77 The pipe handle,
78 .Fa ph ,
79 identifies the endpoint that it the USB is trying to open or close
80 through its endpoint descriptor in the
81 .Sy p_ep
82 member.
83 The endpoint descriptor is described in
84 .Xr usb_ep_descr 9S .
85 From the endpoint descriptor the driver can determine the type of
86 endpoint, what the address of the endpoint is, and what direction the
87 endpoint is in.
88 When combined, these uniquely describe the pipe.
89 .Pp
90 To open a pipe, the driver may need additional companion endpoint
91 descriptors.
92 If these are available, they will be in the
93 .Sy p_xep
94 member of the pipe handle.
95 See
96 .Xr usb_ep_xdescr 9S
97 for more information on how to determine which descriptors are present
98 and get the information encoded in them.
99 .Pp
100 Host controller drivers should check the USB address of the
101 USB device that
102 .Fa ph
103 belongs to.
104 The driver may be asked to open a pipe to the root hub.
105 As the root hub is often synthetic, the driver may need to take a different
106 path than normal.
107 .Ss Pipe open specifics
108 A given endpoint on a device can only be opened once.
109 If there's a request to open an already open endpoint, then the request to open
110 the pipe should be failed.
111 .Pp
112 By the time the call to open a pipe returns, the driver should expect
113 that any of the pipe transfer or reset entry points will be called on
114 the pipe.
115 .Pp
116 A driver can establish private data on an endpoint.
117 During pipe open it may set the
118 .Sy p_hcd_private
119 member to any value.
120 Generally this points to an allocated structure that contains data specific to
121 the host controller.
122 This value will remain on the pipe handle.
123 It is the responsibility of the driver to clear the data when the pipe is
124 closed.
125 .Ss Pipe close specifics
126 When a pipe is closed, the driver must clean up all of the resources

```

127 that it allocated when opening the pipe.
128 **For non-periodic transfers, the host controller driver may assume that there**
128 *For non-periodic transfers, the host controller driver may assume that there*
129 are no outstanding transfers that need to be cleaned up.
130 However, the same is not true for periodic pipes.
131 .Pp
132 For pipes that have outstanding periodic transfers, the host controller
133 driver needs to clean them up and quiesce them as though a call to
134 either
135 .Xr usba_hcdi_pipe_stop_intr_polling 9E
136 or
137 .Xr usba_hcdi_pipe_stop_isoc_polling 9E
138 had been called.
139 .Pp
140 Just as with opening the pipe, the driver should pay attention to the
141 address of the USB device, as it may be the root hub, which may be a
142 synthetic pipe.
143 .Pp
144 When a call to
145 .Fn usba_hcdi_pipe_close
146 completes, the device should be in a state that the pipe can be opened
147 again.
148 .Sh RETURN VALUES
149 Upon successful completion, the
150 .Fn usba_hcdi_pipe_open
151 and
152 .Fn usba_hcdi_pipe_close
153 functions should return
154 .Sy USB_SUCCESS .
155 Otherwise, it should return the appropriate USB error.
156 If uncertain, use
157 .Sy USB_FAILURE .
158 .Sh SEE ALSO
159 .Xr usba_hcdi 9E ,
160 .Xr usba_hcdi_pipe_stop_intr_polling 9E ,
161 .Xr usba_hcdi_pipe_stop_isoc_polling 9E ,
162 .Xr usb_ep_descr 9S ,
163 .Xr usb_ep_xdescr 9S ,
164 .Xr usba_pipe_handle_data 9S

8807 Sat Feb 15 09:54:07 2020

new/usr/src/man/man9f/gld.9f

12309 errors in section 9e of the manual

```

1  \" te
2  .\" Copyright (c) 2003, Sun Microsystems, Inc.
3  .\" All Rights Reserved
4  .\" The contents of this file are subject to the terms of the Common Development
5  .\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
6  .\" When distributing Covered Code, include this CDDL HEADER in each file and in
7  .TH GLD 9F \"February 15, 2020\"
8  .TH GLD 9F \"Aug 28, 2003\"
9  .SH NAME
10 gld, gld_mac_alloc, gld_mac_free, gld_register, gld_unregister, gld_recv,
11 gld_sched, gld_intr \- Generic LAN Driver service routines
12 .SH SYNOPSIS
13 .LP
14 .nf
15 #include <sys/gld.h>
16 .fi
17 .sp
18 .nf
19 .sp
20 \fBvoid\fR \fBgld_mac_alloc\fR(\fBdev_info_t * \fR \fIdip\fR);
21 .fi
22 .LP
23 .nf
24 .sp
25 \fBint\fR \fBgld_register\fR(\fBdev_info_t * \fR \fIdip\fR, \fBchar * \fR \fIname\fR
26 .fi
27 .LP
28 .nf
29 .sp
30 \fBint\fR \fBgld_unregister\fR(\fBgld_mac_info_t * \fR \fImacinfo\fR);
31 .fi
32 .LP
33 .nf
34 .sp
35 \fBvoid\fR \fBgld_recv\fR(\fBgld_mac_info_t * \fR \fImacinfo\fR, \fBmbblk_t * \fR \fI
36 .fi
37 .LP
38 .nf
39 .sp
40 \fBvoid\fR \fBgld_sched\fR(\fBgld_mac_info_t * \fR \fImacinfo\fR);
41 .fi
42 .LP
43 .nf
44 .sp
45 \fBuint_t\fR \fBgld_intr\fR(\fBcaddr_t); \fR
46 .fi
47 .LP
48 .nf
49 .sp
50 \fBvoid\fR \fBgld_linkstate\fR(\fBgld_mac_info_t * \fR \fImacinfo\fR, \fBint32_t\fR
51 .fi
52 .SH INTERFACE LEVEL
53 .sp
54 .LP
55 Solaris architecture specific (Solaris DDI).
56 .SH PARAMETERS
57 .sp
58 .ne 2

```

```

57 .na
58 \fB\fImacinfo\fR \fR
59 .ad
60 .RS 13n
61 Pointer to a \fBgld_mac_info\fR(9S) structure.
62 .RE
63 .sp
64 .ne 2
65 .na
66 .ad
67 \fB\fIdip\fR \fR
68 .ad
69 .RS 13n
70 Pointer to \fBdev_info\fR structure.
71 .RE
72 .sp
73 .ne 2
74 .na
75 .ad
76 \fB\fIname\fR \fR
77 .ad
78 .RS 13n
79 Device interface name.
80 .RE
81 .sp
82 .ne 2
83 .na
84 .ad
85 \fB\fImp\fR \fR
86 .ad
87 .RS 13n
88 Pointer to a message block containing a received packet.
89 .RE
90 .sp
91 .ne 2
92 .na
93 .ad
94 \fB\fInewstate\fR \fR
95 .ad
96 .RS 13n
97 Media link state.
98 .RE
99 .sp
100 .SH DESCRIPTION
101 .sp
102 .LP
103 \fBgld_mac_alloc\fR(\fR) allocates a new \fBgld_mac_info\fR(9S) structure and
104 returns a pointer to it. Some of the GLD-private elements of the structure may
105 be initialized before \fBgld_mac_alloc\fR(\fR) returns; all other elements are
106 initialized to zero. The device driver must initialize some structure members,
107 as described in \fBgld_mac_info\fR(9S), before passing the mac_info pointer to
108 \fBgld_register\fR(\fR).
109 .sp
110 .LP
111 \fBgld_mac_free\fR(\fR) frees a \fBgld_mac_info\fR(9S) structure previously
112 allocated by \fBgld_mac_alloc\fR(\fR).
113 .sp
114 .LP
115 \fBgld_register\fR(\fR) is called from the device driver's \fBattach\fR(9E)
116 routine, and is used to link the GLD-based device driver with the GLD
117 framework. Before calling \fBgld_register\fR(\fR) the device driver's
118 \fBattach\fR(9E) routine must first use \fBgld_mac_alloc\fR(\fR) to allocate a
119 \fBgld_mac_info\fR(9S) structure, and initialize several of its structure
120 elements. See \fBgld_mac_info\fR(9S) for more information. A successful call to
121 \fBgld_register\fR(\fR) performs the following actions:
122 .RS +4

```

```

121 .TP
122 .ie t \(\bu
123 .el o
124 links the device-specific driver with the GLD system;
125 .RE
126 .RS +4
127 .TP
128 .ie t \(\bu
129 .el o
130 sets the device-specific driver's private data pointer (using
131 \fBddi_set_driver_private\fR(9F)) to point to the \fBmacinfo\fR structure;
132 .RE
133 .RS +4
134 .TP
135 .ie t \(\bu
136 .el o
137 creates the minor device node.
138 .RE
139 .sp
140 .LP
141 The device interface name passed to \fBgld_register\fR(\|) must exactly match
142 the name of the driver module as it exists in the filesystem.
143 .sp
144 .LP
145 The driver's \fBattach\fR(9E) routine should return \fBDDI_SUCCESS\fR if
146 \fBgld_register\fR(\|) succeeds. If \fBgld_register\fR(\|) returns
147 \fBDDI_FAILURE\fR, the \fBattach\fR(9E) routine should deallocate any resources
148 it allocated before calling \fBgld_register\fR(\|) and then also return
149 \fBDDI_FAILURE\fR.
150 .sp
151 .LP
152 \fBgld_unregister\fR(\|) is called by the device driver's \fBdetach\fR(9E)
153 function, and if successful, performs the following tasks:
154 .RS +4
155 .TP
156 .ie t \(\bu
157 .el o
158 ensures the device's interrupts are stopped, calling the driver's
159 \fBgldm_stop\fR(\|) routine if necessary;
160 .RE
161 .RS +4
162 .TP
163 .ie t \(\bu
164 .el o
165 removes the minor device node;
166 .RE
167 .RS +4
168 .TP
169 .ie t \(\bu
170 .el o
171 unlinks the device-specific driver from the GLD system.
172 .RE
173 .sp
174 .LP
175 If \fBgld_unregister\fR(\|) returns \fBDDI_SUCCESS\fR, the \fBdetach\fR(9E)
176 routine should deallocate any data structures allocated in the \fBattach\fR(9E)
177 routine, using \fBgld_mac_free\fR(\|) to deallocate the \fBmacinfo\fR
178 structure, and return \fBDDI_SUCCESS\fR. If \fBgld_unregister\fR(\|) returns
179 \fBDDI_FAILURE\fR, the driver's \fBdetach\fR(9E) routine must leave the device
180 operational and return \fBDDI_FAILURE\fR.
181 .sp
182 .LP
183 \fBgld_rcv\fR(\|) is called by the driver's interrupt handler to pass a
184 received packet upstream. The driver must construct and pass a STREAMS
185 \fBDM_DATA\fR message containing the raw packet. \fBgld_rcv\fR(\|) determines
186 which STREAMS queues, if any, should receive a copy of the packet, duplicating

```

```

187 it if necessary. It then formats a \fBDL_UNITDATA_IND\fR message, if required,
188 and passes the data up all appropriate streams.
189 .sp
190 .LP
191 The driver should avoid holding mutex or other locks during the call to
192 \fBgld_rcv\fR(\|). In particular, locks that could be taken by a transmit
193 thread may not be held during a call to \fBgld_rcv\fR(\|): the interrupt
194 thread that calls \fBgld_rcv\fR(\|) may in some cases carry out processing
195 that includes sending an outgoing packet, resulting in a call to the driver's
196 \fBgldm_send\fR(\|) routine. If the \fBgldm_send\fR(\|) routine were to try to
197 acquire a mutex being held by the \fBgldm_intr\fR(\|) routine at the time it
198 calls \fBgld_rcv\fR(\|), this could result in a panic due to recursive mutex
199 entry.
200 .sp
201 .LP
202 \fBgld_sched\fR(\|) is called by the device driver to reschedule stalled
203 outbound packets. Whenever the driver's \fBgldm_send\fR(\|) routine has
204 returned \fBGLD_NORESOURCES\fR, the driver must later call \fBgld_sched\fR(\|)
205 to inform the GLD framework that it should retry the packets that previously
206 could not be sent. \fBgld_sched\fR(\|) should be called as soon as possible
207 after resources are again available, to ensure that GLD resumes passing
208 outbound packets to the driver's \fBgldm_send\fR(\|) routine in a timely way.
209 (If the driver's \fBgldm_stop\fR(\|) routine is called, the driver is absolved
210 from this obligation until it later again returns \fBGLD_NORESOURCES\fR from
211 its \fBgldm_send\fR(\|) routine; however, extra calls to \fBgld_sched\fR(\|)
212 will not cause incorrect operation.)
213 .sp
214 .LP
215 \fBgld_intr\fR(\|) is GLD's main interrupt handler. Normally it is specified as
216 the interrupt routine in the device driver's call to \fBddi_add_intr\fR(9F).
217 The argument to the interrupt handler (specified as \fBint_handler_arg\fR in
218 the call to \fBddi_add_intr\fR(9F)) must be a pointer to the
219 \fBgld_mac_info\fR(9S) structure. \fBgld_intr\fR(\|) will, when appropriate,
220 call the device driver's \fBgldm_intr\fR(\|) function, passing that pointer to
221 the \fBgld_mac_info\fR(9S) structure. However, if the driver uses a high-level
222 interrupt, it must provide its own high-level interrupt handler, and trigger a
223 soft interrupt from within that. In this case, \fBgld_intr\fR(\|) may be
224 specified as the soft interrupt handler in the call to
225 \fBddi_add_softintr\fR(\|).
226 .sp
227 .LP
228 \fBgld_linkstate()\fR is called by the device driver to notify GLD of changes
229 in the media link state. The newstate argument should be set to one of the
230 following:
231 .sp
232 .ne 2
233 .na
234 \fB\fBGLD_LINKSTATE_DOWN\fR \fR
235 .ad
236 .RS 26n
237 The media link is unavailable.
238 .RE
239 .sp
240 .sp
241 .ne 2
242 .na
243 \fB\fBGLD_LINKSTATE_UP\fR \fR
244 .ad
245 .RS 26n
246 The media link is unavailable.
247 .RE
248 .sp
249 .sp
250 .ne 2
251 .na
252 \fB\fBGLD_LINKSTATE_UNKNOWN\fR \fR

```

```
253 .ad
254 .RS 26n
255 The status of the media link is unknown.
256 .RE

258 .sp
259 .LP
260 If a driver calls \fBgld_linkstate()\fR, it must also set the GLD_CAP_LINKSTATE
261 bit in the gldm_capabilities field of the \fBgld_mac_info\fR(9S) structure.
262 bit in the gldm_capabilities field of the \fBgld_mac_info\fR(9S) structure.
263 .SH RETURN VALUES
264 .sp
265 .LP
266 \fBgld_mac_alloc\fR(\fR\|) returns a pointer to a new \fBgld_mac_info\fR(9S)
267 structure.
268 .sp
269 .LP
270 \fBgld_register\fR(\fR\|) and \fBgld_unregister\fR(\fR\|) return:
271 .sp
272 .na
273 \fB\fBDDI_SUCCESS\fR \fR
274 .ad
275 .RS 16n
276 on success.
277 .RE

279 .sp
280 .na
281 \fB\fBDDI_FAILURE\fR \fR
282 .ad
283 .RS 16n
284 on failure.
285 .RE

287 .sp
288 .LP
289 \fBgld_intr\fR(\fR\|) returns a value appropriate for an interrupt handler.
290 .SH SEE ALSO
291 .sp
292 .LP
293 \fBgld\fR(7D), \fBgld\fR(9E), \fBgld_mac_info\fR(9S), \fBgld_stats\fR(9S),
294 \fBdmpi\fR(7P), \fBattach\fR(9E), \fBbdi_add_intr\fR(9F)
```



```
*****
6797 Sat Feb 15 09:54:07 2020
new/usr/src/man/man9f/mac_register.9f
12309 errors in section 9e of the manual
*****
```

```
1 .\"
2 .\" This file and its contents are supplied under the terms of the
3 .\" Common Development and Distribution License ("CDDL"), version 1.0.
4 .\" You may only use this file in accordance with the terms of version
5 .\" 1.0 of the CDDL.
6 .\"
7 .\" A full copy of the text of the CDDL should have accompanied this
8 .\" source. A copy of the CDDL is also available via the Internet at
9 .\" http://www.illumos.org/license/CDDL.
10 .\"
11 .\"
12 .\" Copyright (c) 2017, Joyent, Inc.
13 .\"
14 .Dd February 15, 2020
14 .Dd September 22, 2017
15 .Dt MAC_REGISTER 9F
16 .Os
17 .Sh NAME
18 .Nm mac_register ,
19 .Nm mac_unregister
20 .Nd register and unregister a device driver from the MAC framework
21 .Sh SYNOPSIS
22 .In sys/mac_provider.h
23 .Ft int
24 .Fo mac_register
25 .Fa "mac_register_t *mregp"
26 .Fa "mac_handle_t *mhp"
27 .Fc
28 .Ft int
29 .Fo mac_unregister
30 .Fa "mac_handle_t mh"
31 .Fc
32 .Sh INTERFACE LEVEL
33 illumos DDI specific
34 .Sh PARAMETERS
35 .Bl -tag -width Fa
36 .It Fa mregp
37 A pointer to a
38 .Xr mac_register 9S
39 structure allocated by calling
40 .Xr mac_alloc 9F
41 and filled in by the device driver.
42 .It Fa mhp
43 A pointer to a driver-backed handle to the MAC framework.
44 .It Fa mh
45 The driver-backed handle to the MAC framework.
46 .El
47 .Sh DESCRIPTION
48 The
49 .Fn mac_register
50 function is used to register an instance of a device driver with the
51 .Xr mac 9E
52 framework.
53 Upon successfully calling the
54 .Fn mac_register
55 function, the device will start having its
56 .Xr mac_callbacks 9S
57 entry points called.
58 The device driver should call this function during it's
59 .Xr attach 9E
60 entry point after the device has been configured and is set up.
```

```
61 For a more detailed explanation of the exact steps that the device driver
62 should take and where in the sequence of a driver's
63 .Xr attach 9E
64 entry point this function should be called, see the
65 .Sx Registering with MAC
66 section of
67 .Xr mac 9E .
68 .Pp
69 The driver should provide a pointer to a
70 .Ft mac_handle_t
71 structure as the second argument to the
72 .Fn mac_register
73 function.
74 This handle will be used when the device driver needs to interact with the
75 framework in various ways throughout its life.
76 It is also where the driver gets the
77 .Fa mh
78 argument for calling the
79 .Fn mac_unregister
80 function.
81 It is recommended that the device driver keep the handle around in its soft
82 state structure for a given instance.
83 .Pp
84 If the call to the
85 .Fn mac_register
86 function fails, the device driver should unwind its
87 .Xr attach 9E
88 entry point, tear down everything that it initialized, and ultimately
89 return an error from its
90 .Xr attach 9E
91 entry point.
92 .Pp
93 If the
94 .Xr attach 9E
95 routine fails for some reason after the call to the
96 .Fn mac_register
97 function has succeeded, then the driver should call the
98 .Fn mac_unregister
99 function as part of unwinding all of its state.
100 .Pp
101 When a driver is in its
102 .Xr detach 9E
103 entry point, it should call the
104 .Fn mac_unregister
105 function immediately after draining any of its transmit and receive
106 resources that might have been given to the rest of the operating system
107 through DMA binding.
108 See the
109 .Sx MBLKS AND DMA
110 section of
111 .Xr mac 9E
112 for more information.
113 This should be done before the driver does any tearing down.
114 The call to the
115 .Fn mac_unregister
116 function may fail.
117 This may happen because the networking stack is still using the device.
118 In such a case, the driver should fail the call to
119 .Xr detach 9E
120 and return
121 .Sy DDI_FAILURE .
122 .Sh CONTEXT
123 The
124 .Fn mac_register
125 function is generally only called from a driver's
126 .Xr attach 9E
```

```

127 entry point.
128 The
129 .Fn mac_unregister
130 function is generally only called from a driver's
131 .Xr attach 9E
132 and
133 .Xr detach 9E
134 entry point.
135 However, both functions may be called from either
136 .Sy user
137 or
138 .Sy kernel
139 context.
140 .Sh RETURN VALUES
141 Upon successful completion, the
142 .Fn mac_register
143 and
144 .Fn mac_unregister
145 functions both return
146 .Sy 0 .
147 Otherwise, they return an error number.
148 .Sh EXAMPLES
149 The following example shows how a device driver might call the
150 .Fn mac_register
151 function.
152 .Bd -literal
153 #include <sys/mac_provider.h>
154 #include <sys/mac_ether.h>
155
156 /*
157  * The call to mac_register(9F) generally comes from the context of
158  * attach(9E). This function encapsulates setting up and initializing
159  * the mac_register_t structure and should be assumed to be called from
160  * attach.
161  *
162  * The exact set of callbacks and private properties will vary based
163  * upon the driver.
164  */
165
166 static char *example_priv_props[] = {
167     "_rx_intr_throttle",
168     "_tx_intr_throttle",
169     NULL
170 };
171
172 static mac_callbacks_t example_m_callbacks = {
173     .mc_callbacks = MC_GETCAPAB | MC_SETPROP | MC_GETPROP | MC_PROPINFO |
174     .mc_callbacksk = MC_GETCAPAB | MC_SETPROP | MC_GETPROP | MC_PROPINFO |
175     MC_IOCTL,
176     .mc_start = example_m_start,
177     .mc_stop = example_m_stop,
178     .mc_setpromisc = example_m_setpromisc,
179     .mc_multicst = example_m_multicst,
180     .mc_unicst = example_m_unicst,
181     .mc_tx = example_m_tx,
182     .mc_ioctl = example_m_ioctl,
183     .mc_getcapab = example_m_getcapab,
184     .mc_getprop = example_m_getprop,
185     .mc_setprop = example_m_setprop,
186     .mc_propinfo = example_m_propinfo
187 };
188
189 static boolean_t
190 example_register_mac(example_t *ep)
191 {
192     int status;

```

```

192     mac_register_t *mac;
193
194     mac = mac_alloc(MAC_VERSION);
195     if (mac == NULL)
196         return (B_FALSE);
197
198     mac->m_type_ident = MAC_PLUGIN_IDENT_ETHER;
199     mac->m_driver = ep;
200     mac->m_dip = ep->ep_dev_info;
201     mac->m_src_addr = ep->ep_mac_addr;
202     mac->m_callbacks = &example_m_callbacks;
203     mac->m_min_sdu = 0;
204     mac->m_max_sdu = ep->ep_sdu;
205     mac->m_margin = VLAN_TAGSZ;
206     mac->m_priv_props = example_priv_props;
207
208     status = mac_register(mac, &ep->ep_mac_hdl);
209     mac_free(mac);
210
211     return (status == 0);
212 }

```

unchanged portion omitted