

```

*****
3424 Sat Jul 25 23:07:29 2015
new/usr/src/man/man1m/k5srvutil.1m
4514 k5srvutil(1m): The default keytab file is /etc/krb5/krb5.keytab
*****
1  \" te
2  \" Copyright (c) 2006, Sun Microsystems, Inc. All Rights Reserved
3  \" The contents of this file are subject to the terms of the Common Development
4  \" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
5  \" When distributing Covered Code, include this CDDL HEADER in each file and in
6  .TH K5SRVUTIL 1M \"Jul 25, 2015\"
6  .TH K5SRVUTIL 1M \"Aug 29, 2006\"
7  .SH NAME
8  k5srvutil \- host key table (keytab) manipulation utility
9  .SH SYNOPSIS
10 .LP
11 .nf
12 \fB/usr/sbin/k5srvutil\fR \fIoperation\fR [\fB-ik\fR] [\fB-f\fR \fIfilename\fR]
13 .fi

15 .SH DESCRIPTION
16 .sp
17 .LP
18 The \fBk5srvutil\fR command allows a system manager to list or change keys
19 currently in his keytab or to add new keys to the keytab.
20 .sp
21 .LP
22 The operand \fIoperation\fR must be one of the following:
23 .sp
24 .ne 2
25 .na
26 \fB\fBlist\fR\fR
27 .ad
28 .RS 10n
29 Lists the keys in a keytab, showing version number and principal name.
30 .RE

32 .sp
33 .ne 2
34 .na
35 \fB\fBchange\fR\fR
36 .ad
37 .RS 10n
38 Changes all the keys in the keytab to new randomly-generated keys, updating the
39 keys in the Kerberos server's database to match those by using the \fBkadmin\fR
40 protocol. If a key's version number does not match the version number stored in
41 the Kerberos server's database, the operation fails. The old keys are retained
42 so that existing tickets continue to work. If the \fB-i\fR flag is specified,
43 \fBk5srvutil\fR prompts for \fBByes\fR or \fBno\fR before changing each key. If
44 the \fB-k\fR option is used, the old and new keys are displayed.
45 .RE

47 .sp
48 .ne 2
49 .na
50 \fB\fBdelold\fR\fR
51 .ad
52 .RS 10n
53 Deletes keys that are not the most recent version from the keytab. This
54 operation should be used at some point after a change operation to remove old
55 keys. If the \fB-i\fR flag is specified, \fBk5srvutil\fR asks the user whether
56 the old keys associated with each principal should be removed.
57 .RE

59 .sp
60 .ne 2

```

```

61 .na
62 \fB\fBdelete\fR\fR
63 .ad
64 .RS 10n
65 Deletes particular keys in the keytab, interactively prompting for each key.
66 .RE

68 .sp
69 .LP
70 In all cases, the default keytab file is \fB/etc/krb5/krb5.keytab\fR file unless
71 in all cases, the default keytab file is \fB/etc/krb5.keytab\fR file unless
72 this is overridden by the \fB-f\fR option.
73 .LP
74 \fBk5srvutil\fR uses the \fBkadmin\fR(1M) program to edit the keytab in place.
75 However, old keys are retained, so they are available in case of failure.
76 .SH OPTIONS
77 .sp
78 .LP
79 The following options are supported:
80 .sp
81 .ne 2
82 .na
83 \fB\fB-f\fR \fIfilename\fR\fR
84 .ad
85 .RS 15n
86 Specify a keytab file other than the default file, \fB/etc/krb5/krb5.keytab\fR.
87 Specify a keytab file other than the default file, \fB/etc/krb5.keytab\fR.
87 .RE

89 .sp
90 .ne 2
91 .na
92 \fB\fB-i\fR\fR
93 .ad
94 .RS 15n
95 Prompts user before changing keys when using the \fBchange\fR or \fBdelold\fR
96 operands.
97 .RE

99 .sp
100 .ne 2
101 .na
102 \fB\fB-k\fR\fR
103 .ad
104 .RS 15n
105 Displays old and new keys when using the \fBchange\fR operand.
106 .RE

108 .SH ATTRIBUTES
109 .sp
110 .LP
111 See \fBattributes\fR(5) for descriptions of the following attributes:
112 .sp

114 .sp
115 .TS
116 box;
117 c | c
118 l | l .
119 ATTRIBUTE TYPE ATTRIBUTE VALUE
120 _
121 Interface Stability Committed
122 .TE

124 .SH SEE ALSO

```

**new/usr/src/man/man1m/k5srvutil.1m**

**3**

125 .sp  
126 .LP  
127 \fBktutil\fR(1), \fBkadmin\fR(1M), \fBattributes\fR(5)

```

*****
20641 Sat Jul 25 23:07:30 2015
new/usr/src/man/man3nsl/rpc_clnt_create.3nsl
5730 Typos in rpc_clnt_create(3nsl) man page
*****

```

```

1  \" te
2  .\" Copyright 1989 AT&T
3  .\" Copyright (C) 2009, Sun Microsystems, Inc. All Rights Reserved
4  .\" The contents of this file are subject to the terms of the Common Development
5  .\" See the License for the specific language governing permissions and limitat
6  .\" the fields enclosed by brackets "[]" replaced with your own identifying info
7  .TH RPC_CLNT_CREATE 3NSL "Jul 25, 2015"
8  .TH RPC_CLNT_CREATE 3NSL "Dec 16, 2013"
9  .SH NAME
10 rpc_clnt_create, clnt_control, clnt_create, clnt_create_timed,
11 clnt_create_vers, clnt_create_vers_timed, clnt_destroy, clnt_dg_create,
12 clnt_pcreateerror, clnt_raw_create, clnt_spcreateerror, clnt_tli_create,
13 clnt_tp_create, clnt_tp_create_timed, clnt_vc_create, rpc_createerr,
14 clnt_door_create \- library routines for dealing with creation and manipulation
15 of CLIENT handles
16 .SH SYNOPSIS
17 .LP
18 #include <rpc/rpc.h>
19
20 \fBbool_t\fR \fBclnt_control\fR(\fBCLIENT *fR\fIclnt\fR, \fBconst uint_t\fR \fI
21 .fi
22
23 .LP
24 .nf
25 \fBCLIENT *fR\fBclnt_create\fR(\fBconst char *fR\fIhost\fR, \fBconst rpcprog_t
26 \fBconst rpcvers_t\fR \fIversnum\fR, \fBconst char *fR\fInettype\fR);
27 .fi
28
29 .LP
30 .nf
31 \fBCLIENT *fR\fBclnt_create_timed\fR(\fBconst char *fR\fIhost\fR, \fBconst rpc
32 \fBconst rpcvers_t\fR \fIversnum\fR, \fBconst char *fR\fInettype\fR,
33 \fBconst struct timeval *fR\fItimeout\fR);
34 \fBconst rpcvers_t\fR \fIversnum\fR, \fBconst \fR \fInettype\fR,
35 \fBconst struct timeval *fR\fItimeout\fR);
36 .fi
37
38 .LP
39 .nf
40 \fBCLIENT *fR\fBclnt_create_vers\fR(\fBconst char *fR\fIhost\fR,
41 \fBconst rpcprog_t\fR \fIprognum\fR, \fBconst rpcvers_t *fR\fIvers_outp\fR,
42 \fBconst rpcvers_t\fR \fIvers_low\fR, \fBconst rpcvers_t\fR \fIvers_high\fR,
43 \fBconst char *fR\fInettype\fR);
44 .fi
45
46 .LP
47 .nf
48 \fBCLIENT *fR\fBclnt_create_vers_timed\fR(\fBconst char *fR\fIhost\fR,
49 \fBconst rpcprog_t\fR \fIprognum\fR, \fBconst rpcvers_t *fR\fIvers_outp\fR,
50 \fBconst rpcvers_t\fR \fIvers_low\fR, \fBconst rpcvers_t\fR \fIvers_high\fR,
51 \fBconst char *fR\fInettype\fR, \fBconst struct timeval *fR\fItimeout\fR);
52 .fi
53
54 .LP
55 .nf
56 \fBvoid\fR \fBclnt_destroy\fR(\fBCLIENT *fR\fIclnt\fR);
57 .fi
58
59 .LP
60 .nf

```

```

61 \fBCLIENT *fR\fBclnt_dg_create\fR(\fBconst int\fR \fIfiles\fR,
62 \fBconst struct netbuf *fR\fIsvcaddr\fR, \fBconst rpcprog_t\fR \fIprognum\fR,
63 \fBconst rpcvers_t\fR \fIversnum\fR, \fBconst uint_t\fR \fIendsz\fR,
64 \fBconst uint_t\fR \fIrecsz\fR);
65 .fi
66
67 .LP
68 .nf
69 \fBvoid\fR \fBclnt_pcreateerror\fR(\fBconst char *fR\fIs\fR);
70 .fi
71
72 .LP
73 .nf
74 \fBCLIENT *fR\fBclnt_raw_create\fR(\fBconst rpcprog_t\fR \fIprognum\fR,
75 \fBconst rpcvers_t\fR \fIversnum\fR);
76 .fi
77
78 .LP
79 .nf
80 \fBchar *fR\fBclnt_spcreateerror\fR(\fBconst char *fR\fIs\fR);
81 .fi
82
83 .LP
84 .nf
85 \fBCLIENT *fR\fBclnt_tli_create\fR(\fBconst int\fR \fIfiles\fR,
86 \fBconst struct netconfig *fR\fInetconf\fR, \fBconst struct netbuf *fR\fI
87 \fBconst struct netconfig *fR\fInetconf\fR, \fBconst rpcvers_t\fR \fIversnum\fR,
88 \fBconst uint_t\fR \fIendsz\fR, \fBconst uint_t\fR \fIrecsz\fR);
89 .fi
90
91 .LP
92 .nf
93 \fBCLIENT *fR\fBclnt_tp_create\fR(\fBconst char *fR\fIhost\fR,
94 \fBconst rpcprog_t\fR \fIprognum\fR, \fBconst rpcvers_t\fR \fIversnum\fR,
95 \fBconst struct netconfig *fR\fInetconf\fR);
96 .fi
97
98 .LP
99 .nf
100 \fBCLIENT *fR\fBclnt_tp_create_timed\fR(\fBconst char *fR\fIhost\fR,
101 \fBconst rpcprog_t\fR \fIprognum\fR, \fBconst rpcvers_t\fR \fIversnum\fR,
102 \fBconst struct netconfig *fR\fInetconf\fR, \fBconst struct timeval *fR\fI
103 .LP
104 .nf
105 \fBCLIENT *fR\fBclnt_vc_create\fR(\fBconst int\fR \fIfiles\fR,
106 \fBconst struct netbuf *fR\fIsvcaddr\fR, \fBconst rpcprog_t\fR \fIprognum\fR,
107 \fBconst rpcvers_t\fR \fIversnum\fR, \fBconst uint_t\fR \fIendsz\fR,
108 \fBconst uint_t\fR \fIrecsz\fR);
109 .fi
110
111 .LP
112 .nf
113 \fBstruct rpc_createerr\fR \fBrpc_createerr\fR;
114 .fi
115
116 .LP
117 .nf
118 \fBCLIENT *fR\fBclnt_door_create\fR(\fBconst rpcprog_t\fR \fIprognum\fR,
119 \fBconst rpcvers_t\fR \fIversnum\fR, \fBconst uint_t\fR \fIendsz\fR);
120 .fi
121
122 .SH DESCRIPTION
123 .sp
124 .LP

```

```

125 \fBRPC\fR library routines allow \fBC\fR language programs to make procedure
126 calls on other machines across the network. First a \fBCLIENT\fR handle is
127 created and then the client calls a procedure to send a request to the server.
128 On receipt of the request, the server calls a dispatch routine to perform the
129 requested service, and then sends a reply.
130 .sp
131 .LP
132 These routines are MT-Safe. In the case of multithreaded applications, the
133 \fB-mt\fR option must be specified on the command line at compilation time.
134 When the \fB-mt\fR option is specified, \fBBrpc_createerr\fR becomes a macro
135 that enables each thread to have its own \fBBrpc_createerr\fR. See
136 \fBthreads\fR(5).
137 .SS "Routines"
138 .sp
139 .LP
140 See \fBBrpc\fR(3NSL) for the definition of the \fBCLIENT\fR data structure.
141 .sp
142 .ne 2
143 .na
144 \fB\fB\fBclnt_control()\fR\fR
145 .ad
146 .sp .6
147 .RS 4n
148 A function macro to change or retrieve various information about a client
149 object. \fBfireq\fR indicates the type of operation, and \fBfinfo\fR is a pointer
150 to the information. For both connectionless and connection-oriented transports,
151 the supported values of \fBfireq\fR and their argument types and what they do
152 are:
153 .sp
154 .in +2
155 .nf
156 CLSET_TIMEOUT struct timeval * set total timeout
157 CLGET_TIMEOUT struct timeval * get total timeout
158 .fi
159 .in -2

161 If the timeout is set using \fBclnt_control()\fR, the timeout argument passed
162 by \fBclnt_call()\fR is ignored in all subsequent calls. If the timeout value
163 is set to \fB0\fR, \fBclnt_control()\fR immediately returns
164 \fBBRPC_TIMEDOUT\fR. Set the timeout parameter to \fB0\fR for batching calls.
165 .sp
166 .in +2
167 .nf
168 CLGET_SERVER_ADDR struct netbuf * get server's address
169 CLGET_SVC_ADDR struct netbuf * get server's address
170 CLGET_FD int * get associated file descriptor
171 CLSET_FD_CLOSE void close the file descriptor when
172 destroying the client handle
173 (see \fBclnt_destroy()\fR)
174 CLSET_FD_NCLOSE void do not close the file
175 descriptor when destroying the client handle
176 CLGET_VERS rpcvers_t get the RPC program's version
177 number associated with the
178 client handle
179 CLSET_VERS rpcvers_t set the RPC program's version
180 number associated with the
181 client handle. This assumes
182 that the RPC server for this
183 new version is still listening
184 at the address of the previous
185 version.
186 CLGET_XID uint32_t get the XID of the previous
187 remote procedure call
188 CLSET_XID uint32_t set the XID of the next
189 remote procedure call
190 CLGET_PROG rpcprog_t get program number

```

```

191 CLSET_PROG rpcprog_t set program number
192 .fi
193 .in -2

195 The following operations are valid for connection-oriented transports only:
196 .sp
197 .in +2
198 .nf
199 CLSET_IO_MODE rpciomode_t* set the IO mode used
200 to send one-way requests. The argument for this operation
201 can be either:
202 - RPC_CL_BLOCKING all sending operations block
203 until the underlying transport protocol has
204 accepted requests. If you specify this argument
205 you cannot use flush and getting and setting buffer
206 size is meaningless.
207 - RPC_CL_NONBLOCKING sending operations do not
208 block and return as soon as requests enter the buffer.
209 You can now use non-blocking I/O. The requests in the
210 buffer are pending. The requests are sent to
211 the server as soon as a two-way request is sent
212 or a flush is done. You are responsible for flushing
213 the buffer. When you choose RPC_CL_NONBLOCKING argument
214 you have a choice of flush modes as specified by
215 CLSET_FLUSH_MODE.
216 CLGET_IO_MODE rpciomode_t* get the current IO mode
217 CLSET_FLUSH_MODE rpcflushmode_t* set the flush mode.
218 The flush mode can only be used in non-blocking I/O mode.
219 The argument can be either of the following:
220 - RPC_CL_BESTEFFORT_FLUSH: All flushes send requests
221 in the buffer until the transport end-point blocks.
222 If the transport connection is congested, the call
223 returns directly.
224 - RPC_CL_BLOCKING_FLUSH: Flush blocks until the
225 underlying transport protocol accepts all pending
226 requests into the queue.
227 CLGET_FLUSH_MODE rpcflushmode_t* get the current flush mode.
228 CLFLUSH rpcflushmode_t flush the pending requests.
229 This command can only be used in non-blocking I/O mode.
230 The flush policy depends on which of the following
231 parameters is specified:
232 - RPC_CL_DEFAULT_FLUSH, or NULL: The flush is done
233 according to the current flush mode policy
234 (see CLSET_FLUSH_MODE option).
235 - RPC_CL_BESTEFFORT_FLUSH: The flush tries
236 to send pending requests without blocking; the call
237 returns directly. If the transport connection is
238 congested, this call could return without the request
239 being sent.
240 - RPC_CL_BLOCKING_FLUSH: The flush sends all pending
241 requests. This call will block until all the requests
242 have been accepted by the transport layer.
243 CLSET_CONNMAXREC_SIZE int* set the buffer size.
244 It is not possible to dynamically
245 resize the buffer if it contains data.
246 The default size of the buffer is 16 kilobytes.
247 CLGET_CONNMAXREC_SIZE int* get the current size of the
248 buffer
249 CLGET_CURRENT_REC_SIZE int* get the size of
250 the pending requests stored in the buffer. Use of this
251 command is only recommended when you are in non-blocking
252 I/O mode. The current size of the buffer is always zero
253 when the handle is in blocking mode as the buffer is not
254 used in this mode.
255 .fi
256 .in -2

```

```

258 The following operations are valid for connectionless transports only:
259 .sp
260 .in +2
261 .nf
262 CLSET_RETRY_TIMEOUT struct timeval * set the retry timeout
263 CLGET_RETRY_TIMEOUT struct timeval * get the retry timeout
264 .fi
265 .in -2

267 The retry timeout is the time that \fBRPC\fR waits for the server to reply
268 before retransmitting the request.
269 .sp
270 \fBclnt_control()\fR returns \fBTRUE\fR on success and \fBFALSE\fR on failure.
271 .RE

273 .sp
274 .ne 2
275 .na
276 \fB\fBclnt_create()\fR\fR
277 .ad
278 .sp .6
279 .RS 4n
280 Generic client creation routine for program \fIprognum\fR and version
281 \fIversnum\fR. \fIhost\fR identifies the name of the remote host where the
282 server is located. \fInettype\fR indicates the class of transport protocol to
283 use. The transports are tried in left to right order in \fBNETPATH\fR variable
284 or in top to bottom order in the netconfig database.
285 .sp
286 \fBclnt_create()\fR tries all the transports of the \fInettype\fR class
287 available from the \fBNETPATH\fR environment variable and the netconfig
288 database, and chooses the first successful one. A default timeout is set and
289 can be modified using \fBclnt_control()\fR. This routine returns \fBNULL\fR if
290 it fails. The \fBclnt_pcreateerror()\fR routine can be used to print the reason
291 for failure.
292 .sp
293 Note that \fBclnt_create()\fR returns a valid client handle even if the
294 particular version number supplied to \fBclnt_create()\fR is not registered
295 with the \fBbrpcbind\fR service. This mismatch will be discovered by a
296 \fBclnt_call\fR later (see \fBbrpc_clnt_calls\fR(3NSL)).
297 .RE

299 .sp
300 .ne 2
301 .na
302 \fB\fBclnt_create_timed()\fR\fR
303 .ad
304 .sp .6
305 .RS 4n
306 Generic client creation routine which is similar to \fBclnt_create()\fR but
307 which also has the additional parameter \fItimeout\fR that specifies the
308 maximum amount of time allowed for each transport class tried. In all other
309 respects, the \fBclnt_create_timed()\fR call behaves exactly like the
310 \fBclnt_create()\fR call.
311 .RE

313 .sp
314 .ne 2
315 .na
316 \fB\fBclnt_create_vers()\fR\fR
317 .ad
318 .sp .6
319 .RS 4n
320 Generic client creation routine which is similar to \fBclnt_create()\fR but
321 which also checks for the version availability. \fIhost\fR identifies the name
322 of the remote host where the server is located. \fInettype\fR indicates the

```

```

323 class transport protocols to be used. If the routine is successful it returns a
324 client handle created for the highest version between \fIvers_low\fR and
325 \fIvers_high\fR that is supported by the server. \fIvers_outp\fR is set to this
326 value. That is, after a successful return \fIvers_low\fR <= \fI*vers_outp\fR <=
327 \fIvers_high\fR. If no version between \fIvers_low\fR and \fIvers_high\fR is
328 supported by the server then the routine fails and returns \fBNULL\fR.
329 default timeout is set and can be modified using \fBclnt_control()\fR. This
330 routine returns \fBNULL\fR if it fails. The \fBclnt_pcreateerror()\fR routine
331 can be used to print the reason for failure.
332 .sp
333 Note: \fBclnt_create()\fR returns a valid client handle even if the particular
334 version number supplied to \fBclnt_create()\fR is not registered with the
335 \fBbrpcbind\fR service. This mismatch will be discovered by a \fBclnt_call\fR
336 later (see \fBbrpc_clnt_calls\fR(3NSL)). However, \fBclnt_create_vers()\fR does
337 this for you and returns a valid handle only if a version within the range
338 supplied is supported by the server.
339 .RE

341 .sp
342 .ne 2
343 .na
344 \fB\fBclnt_create_vers_timed()\fR\fR
345 .ad
346 .sp .6
347 .RS 4n
348 Generic client creation routine similar to \fBclnt_create_vers()\fR but with
349 the additional parameter \fItimeout\fR, which specifies the maximum amount of
350 time allowed for each transport class tried. In all other respects, the
351 \fBclnt_create_vers_timed()\fR call behaves exactly like the
352 \fBclnt_create_vers()\fR call.
353 .RE

355 .sp
356 .ne 2
357 .na
358 \fB\fBclnt_destroy()\fR\fR
359 .ad
360 .sp .6
361 .RS 4n
362 A function macro that destroys the client's \fBRPC\fR handle. Destruction
363 usually involves deallocation of private data structures, including \fBclnt\fR
364 itself. Use of \fBclnt\fR is undefined after calling \fBclnt_destroy()\fR. If
365 the \fBRPC\fR library opened the associated file descriptor, or
366 \fBCLSET_FD_CLOSE\fR was set using \fBclnt_control()\fR, the file descriptor
367 will be closed.
368 .sp
369 The caller should call \fBauth_destroy(\fR\fBclnt\fR->\fBauth)\fR (before
370 calling \fBclnt_destroy()\fR) to destroy the associated \fBBAUTH\fR structure
371 (see \fBbrpc_clnt_auth\fR(3NSL)).
372 .RE

374 .sp
375 .ne 2
376 .na
377 \fB\fBclnt_dg_create()\fR\fR
378 .ad
379 .sp .6
380 .RS 4n
381 This routine creates an \fBRPC\fR client for the remote program \fIprognum\fR
382 and version \fIversnum\fR; the client uses a connectionless transport. The
383 remote program is located at address \fIsvcadddr\fR. The parameter \fIfildes\fR
384 is an open and bound file descriptor. This routine will resend the call message
385 in intervals of 15 seconds until a response is received or until the call times
386 out. The total time for the call to time out is specified by \fBclnt_call()\fR
387 (see \fBclnt_call()\fR in \fBbrpc_clnt_calls\fR(3NSL)). The retry time out and
388 the total time out periods can be changed using \fBclnt_control()\fR. The user

```

```

389 may set the size of the send and receive buffers with the parameters
390 \fIisendsz\fR and \fIrecvsz\fR; values of \fB0\fR choose suitable defaults. This
391 routine returns \fINULL\fR if it fails.
392 .RE

394 .sp
395 .ne 2
396 .na
397 \fB\fBclnt_pcreateerror()\fR\fR
398 .ad
399 .sp .6
400 .RS 4n
401 Print a message to standard error indicating why a client \fBRPC\fR handle
402 could not be created. The message is prepended with the string \fIs\fR and a
403 colon, and appended with a newline.
404 .RE

406 .sp
407 .ne 2
408 .na
409 \fB\fBclnt_raw_create()\fR\fR
410 .ad
411 .sp .6
412 .RS 4n
413 This routine creates an \fBRPC\fR client handle for the remote program
414 \fIprognum\fR and version \fIversnum\fR. The transport used to pass messages to
415 the service is a buffer within the process's address space, so the
416 corresponding \fBRPC\fR server should live in the same address space; (see
417 \fBsvc_raw_create()\fR in \fBrpc_svc_create(3NLS)\fR). This allows simulation
418 of \fBRPC\fR and measurement of \fBRPC\fR overheads, such as round trip times,
419 without any kernel or networking interference. This routine returns \fINULL\fR
420 if it fails. \fBclnt_raw_create()\fR should be called after
421 \fBsvc_raw_create()\fR.
422 .RE

424 .sp
425 .ne 2
426 .na
427 \fB\fBclnt_spcreateerror()\fR\fR
428 .ad
429 .sp .6
430 .RS 4n
431 Like \fBclnt_pcreateerror()\fR, except that it returns a string instead of
432 printing to the standard error. A newline is not appended to the message in
433 this case.
434 .sp
435 Warning: returns a pointer to a buffer that is overwritten on each call. In
436 multithread applications, this buffer is implemented as thread-specific data.
437 .RE

439 .sp
440 .ne 2
441 .na
442 \fB\fBclnt_tli_create()\fR\fR
443 .ad
444 .sp .6
445 .RS 4n
446 This routine creates an \fBRPC\fR client handle for the remote program
447 \fIprognum\fR and version \fIversnum\fR. The remote program is located at
448 address \fIsvccaddr\fR. If \fIsvccaddr\fR is \fINULL\fR and it is
449 connection-oriented, it is assumed that the file descriptor is connected. For
450 connectionless transports, if \fIsvccaddr\fR is \fINULL\fR,
451 \fBRPC_UNKNOWADDR\fR error is set. \fIfildes\fR is a file descriptor which may
452 be open, bound and connected. If it is \fBRPC_ANYFD\fR, it opens a file
453 descriptor on the transport specified by \fInetconf\fR. If \fIfildes\fR is
454 \fBRPC_ANYFD\fR and \fInetconf\fR is \fINULL\fR, a \fBRPC_UNKNOWPROTO\fR error

```

```

455 is set. If \fIfildes\fR is unbound, then it will attempt to bind the
456 descriptor. The user may specify the size of the buffers with the parameters
457 \fIisendsz\fR and \fIrecvsz\fR; values of \fB0\fR choose suitable defaults.
458 Depending upon the type of the transport (connection-oriented or
459 connectionless), \fBclnt_tli_create()\fR calls appropriate client creation
460 routines. This routine returns \fINULL\fR if it fails. The
461 \fBclnt_pcreateerror()\fR routine can be used to print the reason for failure.
462 The remote \fBbrpcbind\fR service (see \fBbrpcbind(1M)\fR) is not consulted for
463 the address of the remote service.
464 .RE

466 .sp
467 .ne 2
468 .na
469 \fB\fBclnt_tp_create()\fR\fR
470 .ad
471 .sp .6
472 .RS 4n
473 Like \fBclnt_create()\fR except \fBclnt_tp_create()\fR tries only one transport
474 specified through \fInetconf\fR.
475 .sp
476 \fBclnt_tp_create()\fR creates a client handle for the program \fIprognum\fR,
477 the version \fIversnum\fR, and for the transport specified by \fInetconf\fR.
478 Default options are set, which can be changed using \fBclnt_control()\fR calls.
479 The remote \fBbrpcbind\fR service on the host \fIhost\fR is consulted for the
480 address of the remote service. This routine returns \fINULL\fR if it fails. The
481 \fBclnt_pcreateerror()\fR routine can be used to print the reason for failure.
482 .RE

484 .sp
485 .ne 2
486 .na
487 \fB\fBclnt_tp_create_timed()\fR\fR
488 .ad
489 .sp .6
490 .RS 4n
491 Like \fBclnt_tp_create()\fR except \fBclnt_tp_create_timed()\fR has the extra
492 parameter \fItimeout\fR which specifies the maximum time allowed for the
493 creation attempt to succeed. In all other respects, the
494 \fBclnt_tp_create_timed()\fR call behaves exactly like the
495 \fBclnt_tp_create()\fR call.
496 .RE

498 .sp
499 .ne 2
500 .na
501 \fB\fBclnt_vc_create()\fR\fR
502 .ad
503 .sp .6
504 .RS 4n
505 This routine creates an \fBRPC\fR client for the remote program \fIprognum\fR
506 and version \fIversnum\fR; the client uses a connection-oriented transport. The
507 remote program is located at address \fIsvccaddr\fR. The parameter \fIfildes\fR
508 is an open and bound file descriptor. The user may specify the size of the send
509 and receive buffers with the parameters \fIisendsz\fR and \fIrecvsz\fR; values
510 of \fB0\fR choose suitable defaults. This routine returns \fINULL\fR if it
511 fails.
512 .sp
513 The address \fIsvccaddr\fR should not be \fINULL\fR and should point to the
514 actual address of the remote program. \fBclnt_vc_create()\fR does not consult
515 the remote \fBbrpcbind\fR service for this information.
516 .RE

518 .sp
519 .ne 2
520 .na

```

```
521 \fB\fBrpc_createerr\fR\fR
522 .ad
523 .sp .6
524 .RS 4n
525 A global variable whose value is set by any \fBRPC\fR client handle creation
526 routine that fails. It is used by the routine \fBclnt_pcreateerror()\fR to
527 print the reason for the failure.
528 .sp
529 In multithreaded applications, \fBrpc_createerr\fR becomes a macro which
530 enables each thread to have its own \fBrpc_createerr\fR.
531 .RE

533 .sp
534 .ne 2
535 .na
536 \fB\fBclnt_door_create()\fR\fR
537 .ad
538 .sp .6
539 .RS 4n
540 This routine creates an RPC client handle over doors for the given program
541 \fIprognum\fR and version \fIversnum\fR. Doors is a transport mechanism that
542 facilitates fast data transfer between processes on the same machine. The user
543 may set the size of the send buffer with the parameter \fIisendsz\fR. If
544 \fIisendsz\fR is 0, the corresponding default buffer size is 16 Kbyte. The
545 \fBclnt_door_create()\fR routine returns \fINULL\fR if it fails and sets a
546 value for \fBrpc_createerr\fR.
547 .RE

549 .SH ATTRIBUTES
550 .sp
551 .LP
552 See \fBattributes\fR(5) for descriptions of the following attributes:
553 .sp

555 .sp
556 .TS
557 box;
558 c | c
559 l | l .
560 ATTRIBUTE TYPE ATTRIBUTE VALUE
561 -
562 Architecture All
563 -
564 Interface Stability Committed
565 -
566 MT-Level MT-Safe
567 .TE

569 .SH SEE ALSO
570 .sp
571 .LP
572 \fBbrpcbind\fR(1M), \fBbrpc\fR(3NSL), \fBbrpc_clnt_auth\fR(3NSL),
573 \fBbrpc_clnt_calls\fR(3NSL), \fBbrpc_svc_create\fR(3NSL),
574 \fBbsvc_raw_create\fR(3NSL), \fBbthreads\fR(5), \fBattributes\fR(5)
```

```

*****
6662 Sat Jul 25 23:07:30 2015
new/usr/src/man/man3nsl/rpcbnd.3nsl
5657 Typo in rpcbind(3nsl): ssvcaddr
*****
1  \" te
2  .\" Copyright 2014 Nexenta Systems, Inc. All Rights Reserved.
3  .\" Copyright 1989 AT&T Copyright (c) 1997, Sun Microsystems, Inc. All Rights
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 RPCBIND 3NSL \"Jul 25, 2015\"
8  .TH RPCBIND 3NSL \"Nov 24, 2014\"
9  .SH NAME
10 rpcbind, rpcb_getmaps, rpcb_getaddr, rpcb_gettime, rpcb_rmtcall, rpcb_set,
11 rpcb_unset \- library routines for RPC bind service
12 .SH SYNOPSIS
13 .LP
14 .nf
15
16
17
18 \fBstruct rpcblist *\fR\fBpcb_getmaps\fR(\fBconst struct netconfig *\fR\fInetco
19 \fBconst char *\fR\fIhost\fR);
20 .fi
21
22 .LP
23 .nf
24 \fBbool_t\fR \fBpcb_getaddr\fR(\fBconst rpcprog_t\fR \fIprognum\fR, \fBconst
25 \fBconst struct netconfig *\fR\fInetconf\fR, \fBstruct netbuf *\fR\fIsvcad
26 \fBconst struct netconfig *\fR\fInetconf\fR, \fBstruct netbuf *\fR\fIsvcad
27 \fBconst char *\fR\fIhost\fR);
28 .fi
29 .LP
30 .nf
31 \fBbool_t\fR \fBpcb_gettime\fR(\fBconst char *\fR\fIhost\fR, \fBtime_t *\fR\fIt
32 .fi
33
34 .LP
35 .nf
36 \fBenum clnt_stat\fR \fBpcb_rmtcall\fR(\fBconst struct netconfig *\fR\fInetconf
37 \fBconst char *\fR\fIhost\fR, \fBconst rpcprog_t\fR \fIprognum\fR,
38 \fBconst rpcvers_t\fR \fIversnum\fR, \fBconst rpcproc_t\fR \fIprocnum\fR,
39 \fBconst xdrproc_t\fR \fIinproc\fR, \fBconst caddr_t\fR \fIin\fR,
40 \fBconst xdrproc_t\fR \fIoutproc\fR, \fBcaddr_t\fR \fIout\fR,
41 \fBconst struct timeval\fR \fIitout\fR, \fBstruct netbuf *\fR\fIsvcaddr\fR)
42 .fi
43
44 .LP
45 .nf
46 \fBbool_t\fR \fBpcb_set\fR(\fBconst rpcprog_t\fR \fIprognum\fR, \fBconst rpcver
47 \fBconst struct netconfig *\fR\fInetconf\fR, \fBconst struct netbuf *\fR\fI
48 .fi
49
50 .LP
51 .nf
52 \fBbool_t\fR \fBpcb_unset\fR(\fBconst rpcprog_t\fR \fIprognum\fR, \fBconst rpcv
53 \fBconst struct netconfig *\fR\fInetconf\fR);
54 .fi
55
56 .SH DESCRIPTION
57 .LP
58 These routines allow client C programs to make procedure calls to the RPC
59 binder service. \fBpcb\fR maintains a list of mappings between programs and

```

```

60 their universal addresses. See \fBpcb\fR(1M).
61 .SS "Routines"
62 .ne 2
63 .na
64 \fBpcb_getmaps()\fR
65 .ad
66 .RS 18n
67 An interface to the \fBpcb\fR service, which returns a list of the current
68 \fBPC\fR program-to-address mappings on \fIhost\fR. It uses the transport
69 specified through \fInetconf\fR to contact the remote \fBpcb\fR service on
70 \fIhost\fR. This routine will return \fBNULL\fR if the remote \fBpcb\fR
71 could not be contacted.
72 .RE
73
74 .sp
75 .ne 2
76 .na
77 \fBpcb_getaddr()\fR
78 .ad
79 .RS 18n
80 An interface to the \fBpcb\fR service, which finds the address of the
81 service on \fIhost\fR that is registered with program number \fIprognum\fR,
82 version \fIversnum\fR, and speaks the transport protocol associated with
83 \fInetconf\fR. The address found is returned in \fIsvcaddr\fR. \fIsvcaddr\fR
84 should be preallocated. This routine returns \fBTRUE\fR if it succeeds. A
85 return value of \fBFALSE\fR means that the mapping does not exist or that the
86 \fBPC\fR system failed to contact the remote \fBpcb\fR service. In the
87 latter case, the global variable \fBpcb_createerr\fR contains the \fBPC
88 status. See \fBpcb_create\fR(3NSL).
89 .RE
90
91 .sp
92 .ne 2
93 .na
94 \fBpcb_gettime()\fR
95 .ad
96 .RS 18n
97 This routine returns the time on \fIhost\fR in \fIitimep\fR. If \fIhost\fR is
98 \fBNULL\fR, \fBpcb_gettime()\fR returns the time on its own machine. This
99 routine returns \fBTRUE\fR if it succeeds, \fBFALSE\fR if it fails.
100 \fBpcb_gettime()\fR can be used to synchronize the time between the client and
101 the remote server. This routine is particularly useful for secure RPC.
102 .RE
103
104 .sp
105 .ne 2
106 .na
107 \fBpcb_rmtcall()\fR
108 .ad
109 .RS 18n
110 An interface to the \fBpcb\fR service, which instructs \fBpcb\fR on
111 \fIhost\fR to make an \fBPC\fR call on your behalf to a procedure on that
112 host. The \fBnetconfig\fR structure should correspond to a connectionless
113 transport. The parameter \fB*\fR\fIsvcaddr\fR will be modified to the server's
114 address if the procedure succeeds. See \fBpcb_call()\fR and \fBclnt_call()\fR
115 in \fBpcb_calls\fR(3NSL) for the definitions of other parameters.
116 .sp
117 This procedure should normally be used for a "ping" and nothing else. This
118 routine allows programs to do lookup and call, all in one step.
119 .sp
120 Note: Even if the server is not running \fBpcb\fR does not return any error
121 messages to the caller. In such a case, the caller times out.
122 .sp
123 Note: \fBpcb_rmtcall()\fR is only available for connectionless transports.
124 .RE

```



```

126 .sp
127 .ne 2
128 .na
129 \fB\fBrpcb_set()\fR\fR
130 .ad
131 .RS 18n
132 An interface to the \fBrpcbnd\fR service, which establishes a mapping between
133 the triple [\fIprognum\fR, \fIversnum\fR, \fInetconf\fR->\fInc_netid\fR] and
134 \fIsvccaddr\fR on the machine's \fBrpcbnd\fR service. The value of
135 \fInc_netid\fR must correspond to a network identifier that is defined by the
136 netconfig database. This routine returns \fBTRUE\fR if it succeeds, \fBFALSE\fR
137 otherwise. See also \fBsvc_reg()\fR in \fBrpc_svc_reg\fR(3NSL). If there
138 already exists such an entry with \fBrpcbnd\fR, \fBrpcb_set()\fR will fail.
139 .RE

```

```

141 .sp
142 .ne 2
143 .na
144 \fB\fBrpcb_unset()\fR\fR
145 .ad
146 .RS 18n
147 An interface to the \fBrpcbnd\fR service, which destroys the mapping between
148 the triple [\fIprognum\fR, \fIversnum\fR, \fInetconf\fR->\fInc_netid\fR] and
149 the address on the machine's \fBrpcbnd\fR service. If \fInetconf\fR is
150 \fINULL\fR, \fBrpcb_unset()\fR destroys all mapping between the triple
151 [\fIprognum\fR, \fIversnum\fR, \fIall-transport\fR] and the addresses on the
152 machine's \fBrpcbnd\fR service. This routine returns \fBTRUE\fR if it
153 succeeds, \fBFALSE\fR otherwise. Only the owner of the service or the
154 super-user can destroy the mapping. See also \fBsvc_unreg()\fR in
155 \fBrpc_svc_reg\fR(3NSL).
156 .RE

```

```

158 .SH ATTRIBUTES
159 .LP
160 See \fBattributes\fR(5) for descriptions of the following attributes:
161 .sp

```

```

163 .sp
164 .TS
165 box;
166 c | c
167 l | l .
168 ATTRIBUTE TYPE ATTRIBUTE VALUE
169 -
170 MT-Level MT-Safe
171 .TE

```

```

173 .SH SEE ALSO
174 .LP
175 \fBrpcbnd\fR(1M), \fBrpcinfo\fR(1M), \fBrpc_clnt_calls\fR(3NSL),
176 \fBrpc_clnt_create\fR(3NSL), \fBrpc_svc_calls\fR(3NSL), \fBattributes\fR(5)

```

\*\*\*\*\*
6103 Sat Jul 25 23:07:30 2015
new/usr/src/man/man9e/detach.9e
4648 detach(9e): Extra space between 'prefix' and 'detach'
\*\*\*\*\*

```
1 \" te
2.\" Copyright (c) 2003, Sun Microsystems, Inc. All Rights Reserved
3.\" The contents of this file are subject to the terms of the Common Development
4.\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
5.\" When distributing Covered Code, include this CDDL HEADER in each file and in
6.TH DETACH 9E \"Jul 25, 2015\"
6.TH DETACH 9E \"Dec 7, 2003\"
7.SH NAME
8 detach \- detach or suspend a device
9.SH SYNOPSIS
10.LP
11.nf
12#include <sys/ddi.h>
13#include <sys/sunddi.h>
17 \fBint prefix\fR \fBdetach\fR(\fBdev_info_t * \fR \fIdip\fR, \fBddi_detach_cmd_t \fR
17 \fBint prefix\fR \fBdetach\fR(\fBdev_info_t * \fR \fIdip\fR, \fBddi_detach_cmd_t \fR
18 .fi
20.SH INTERFACE LEVEL
21.LP
22 Solaris DDI specific (Solaris DDI)
23.SH PARAMETERS
24.ne 2
25.na
26 \fB\fIdip\fR\fR
27.ad
28.RS 7n
29 A pointer to the device's \fBdev_info\fR structure.
30.RE
32.sp
33.ne 2
34.na
35 \fB\fIcmd\fR\fR
36.ad
37.RS 7n
38 Type of detach; the driver should return \fBDDI_FAILURE\fR if any value other
39 than \fBDDI_DETACH\fR or \fBDDI_SUSPEND\fR is passed to it.
40.RE
42.SH DESCRIPTION
43.LP
44 The \fBdetach()\fR function complements the \fBattach\fR(9E) routine.
45.SS \"DDI_DETACH\"
46.LP
47 If \fIcmd\fR is set to \fBDDI_DETACH\fR, \fBdetach()\fR is used to remove the
48 state associated with a given instance of a device node prior to the removal of
49 that instance from the system.
50.sp
51.LP
52 The \fBdetach()\fR function will be called once for each instance of the device
53 for which there has been a successful \fBattach()\fR, once there are no longer
54 any opens on the device. An attached instance of a driver can be successfully
55 detached only once. The \fBdetach()\fR function should clean up any per
56 instance data initialized in \fBattach\fR(9E) and call \fBkmem_free\fR(9F) to
57 free any heap allocations. For information on how to unregister interrupt
58 handlers, see \fBddi_add_intr\fR(9F). This should also include putting the
59 underlying device into a quiescent state so that it will not generate
```

```
60 interrupts.
61.sp
62.LP
63 Drivers that set up \fBtimeout\fR(9F) routines should ensure that they are
64 cancelled before returning \fBDDI_SUCCESS\fR from \fBdetach()\fR.
65.sp
66.LP
67 If \fBdetach()\fR determines a particular instance of the device cannot be
68 removed when requested because of some exceptional condition, \fBdetach()\fR
69 must return \fBDDI_FAILURE\fR, which prevents the particular device instance
70 from being detached. This also prevents the driver from being unloaded. A
71 driver instance failing the detach must ensure that no per instance data or
72 state is modified or freed that would compromise the system or subsequent
73 driver operation.
74.sp
75.LP
76 The system guarantees that the function will only be called for a particular
77 \fBdev_info\fR node after (and not concurrently with) a successful
78 \fBattach\fR(9E) of that device. The system also guarantees that \fBdetach()\fR
79 will only be called when there are no outstanding \fBopen\fR(9E) calls on the
80 device.
81.SS \"DDI_SUSPEND\"
82.LP
83 The \fBDDI_SUSPEND\fR \fIcmd\fR is issued when the entire system is being
84 suspended and power removed from it or when the system must be made quiescent.
85 It will be issued only to devices which have a \fBreg\fR property or which
86 export a \fBbpm-hardware-state\fR property with the value needs-suspend-resume.
87.sp
88.LP
89 If \fIcmd\fR is set to \fBDDI_SUSPEND\fR, \fBdetach()\fR is used to suspend all
90 activity of a device before power is (possibly) removed from the device. The
91 steps associated with suspension must include putting the underlying device
92 into a quiescent state so that it will not generate interrupts or modify or
93 access memory. Once quiescence has been obtained, \fBdetach()\fR can be called
94 with outstanding \fBopen\fR(9E) requests. It must save the hardware state of
95 the device to memory and block incoming or existing requests until
96 \fBattach()\fR is called with \fBDDI_RESUME\fR.
97.sp
98.LP
99 If the device is used to store file systems, then after \fBDDI_SUSPEND\fR is
100 issued, the device should still honor \fBdump\fR(9E) requests as this entry
101 point may be used by suspend-resume operation (see \fBbcp\r(7)) to save state
102 file. It must do this, however, without disturbing the saved hardware state of
103 the device.
104.sp
105.LP
106 If the device driver uses automatic device Power Management interfaces (driver
107 exports \fBbpm-components\fR(9P) property), it might need to call
108 \fBbpm_raise_power\fR(9F) if the current power level is lower than required to
109 complete the \fBdump\fR(9E) request.
110.sp
111.LP
112 Before returning successfully from a call to \fBdetach()\fR with a command of
113 \fBDDI_SUSPEND\fR, the driver must cancel any outstanding timeouts and make any
114 driver threads quiescent.
115.sp
116.LP
117 If \fBDDI_FAILURE\fR is returned for the \fBDDI_SUSPEND\fR \fIcmd\fR, either
118 the operation to suspend the system or to make it quiescent will be aborted.
119.SH RETURN VALUES
120.ne 2
121.na
122 \fB\bBDDI_SUCCESS\fR\fR
123.ad
124.RS 15n
125 For \fBDDI_DETACH\fR, the state associated with the given device was
```

```
126 successfully removed. For \fBDDI_SUSPEND\fR, the driver was successfully
127 suspended.
128 .RE

130 .sp
131 .ne 2
132 .na
133 \fB\fBDDI_FAILURE\fR\fR
134 .ad
135 .RS 15n
136 The operation failed or the request was not understood. The associated state is
137 unchanged.
138 .RE

140 .SH CONTEXT
141 .LP
142 This function is called from user context only.
143 .SH ATTRIBUTES
144 .LP
145 See \fBattributes\fR(5) for descriptions of the following attributes:
146 .sp

148 .sp
149 .TS
150 box;
151 c | c
152 l | l .
153 ATTRIBUTE TYPE ATTRIBUTE VALUE
154 _
155 Interface Stability Committed
156 .TE

158 .SH SEE ALSO
159 .LP
160 \fBbcpr\fR(7), \fBbpm\fR(7D), \fBbpm\fR(9P), \fBbpm-components\fR(9P),
161 \fBattach\fR(9E), \fBdump\fR(9E), \fBopen\fR(9E), \fBpower\fR(9E),
162 \fBddi_add_intr\fR(9F), \fBddi_dev_is_needed\fR(9F), \fBddi_map_regs\fR(9F),
163 \fBkmem_free\fR(9F), \fBpm_raise_power\fR(9F), \fBtimeout\fR(9F)
164 .sp
165 .LP
166 \fIWriting Device Drivers\fR
```

```

*****
6934 Sat Jul 25 23:07:30 2015
new/usr/src/man/man9f/taskq.9f
5036 taskq(9f): Typos in the man page
*****
1  \' te
2  .\" Copyright (c) 2005, 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 TASKQ 9F \"Jul 25, 2015\"
8  .TH TASKQ 9F \"Mar 1, 2005\"
9  .SH NAME
10 taskq, ddi_taskq_create, ddi_taskq_destroy, ddi_taskq_dispatch, ddi_taskq_wait,
11 ddi_taskq_suspend, taskq_suspended, ddi_taskq_resume \- Kernel task queue
12 operations
13 .SH SYNOPSIS
14 .LP
15 .nf
16 #include <sys/sunddi.h>
17 \fBddi_taskq_t *\fR\fBddi_taskq_create\fR(\fBdev_info_t *\fR\fBIdip\fR, \fBconst
18 \fBint\fR \fBInthreads\fR, \fBpri_t\fR \fBipri\fR, \fBbuint_t\fR \fBicflags\fR
19 .fi
20
21 .LP
22 .nf
23 \fBvoid\fR \fBddi_taskq_destroy\fR(\fBddi_taskq_t *\fR\fBItq\fR);
24 .fi
25
26 .LP
27 .nf
28 \fBint\fR \fBddi_taskq_dispatch\fR(\fBddi_taskq_t *\fR\fBItq\fR, \fBvoid (*\fR \fB
29 \fBvoid *\fR\fBarg\fR, \fBbuint_t\fR \fBiflags\fR);
30 .fi
31
32 .LP
33 .nf
34 \fBvoid\fR \fBddi_taskq_wait\fR(\fBddi_taskq_t *\fR\fBItq\fR);
35 .fi
36
37 .LP
38 .nf
39 \fBvoid\fR \fBddi_taskq_suspend\fR(\fBddi_taskq_t *\fR\fBItq\fR);
40 .fi
41
42 .LP
43 .nf
44 \fBboolean_t\fR \fBddi_taskq_suspended\fR(\fBddi_taskq_t *\fR\fBItq\fR);
45 .fi
46
47 .LP
48 .nf
49 \fBvoid\fR \fBddi_taskq_resume\fR(\fBddi_taskq_t *\fR\fBItq\fR);
50 .fi
51
52 .SH INTERFACE LEVEL
53 .sp
54 .LP
55 Solaris DDI specific (Solaris DDI)
56 .SH PARAMETERS
57 .sp
58 .ne 2
59 .na
60 \fB\fIdip\fR

```

```

61 .ad
62 .RS 12n
63 Pointer to the device's dev_info structure. May be NULL for kernel
64 modules that do not have an associated dev_info structure.
65 .RE
66
67 .sp
68 .ne 2
69 .na
70 \fB\fIname\fR
71 .ad
72 .RS 12n
73 Descriptive string. Only alphanumeric characters can be used in name
74 and spaces are not allowed. The name should be unique.
75 .RE
76
77 .sp
78 .ne 2
79 .na
80 \fB\fInthreads\fR
81 .ad
82 .RS 12n
83 Number of threads servicing the task queue. Note that the request ordering is
84 guaranteed (tasks are processed in the order scheduled) if the \fBtaskq\fR is
85 created with a single servicing thread.
86 .RE
87
88 .sp
89 .ne 2
90 .na
91 \fB\fIpri\fR
92 .ad
93 .RS 12n
94 Priority of threads servicing the task queue. Drivers and modules should
95 specify TASKQ_DEFAULTPRI.
96 .RE
97
98 .sp
99 .ne 2
100 .na
101 \fB\fIcflags\fR
102 .ad
103 .RS 12n
104 Should pass 0 as flags.
105 .RE
106
107 .sp
108 .ne 2
109 .na
110 \fB\fIfunc\fR
111 .ad
112 .RS 12n
113 Callback function to call.
114 .RE
115
116 .sp
117 .ne 2
118 .na
119 \fB\fIarg\fR
120 .ad
121 .RS 12n
122 Argument to the callback function.
123 .RE
124
125 .sp
126 .ne 2

```

```

127 .na
128 \fB\fIdflags\fR\fR
129 .ad
130 .RS 12n
131 Possible \fIdflags\fR are:
132 .sp
133 .ne 2
134 .na
135 \fBDDI_SLEEP\fR
136 .ad
137 .RS 15n
138 Allow sleeping (blocking) until memory is available.
139 .RE

141 .sp
142 .ne 2
143 .na
144 \fBDDI_NOSLEEP\fR
145 .ad
146 .RS 15n
147 Return DDI_FAILURE immediately if memory is not available.
148 .RE

150 .RE

152 .sp
153 .ne 2
154 .na
155 \fB\fItq\fR\fR
156 .ad
157 .RS 12n
158 Pointer to a task queue (ddi_taskq_t *).
159 .RE

161 .sp
162 .ne 2
163 .na
164 \fB\fItq\fR\fR
165 .ad
166 .RS 12n
167 Pointer to a thread structure.
168 .RE

170 .SH DESCRIPTION
171 .sp
172 .LP
173 A kernel task queue is a mechanism for general-purpose asynchronous task
174 scheduling that enables tasks to be performed at a later time by another
175 thread. There are several reasons why you may utilize asynchronous task
176 scheduling:
177 .RS +4
178 .TP
179 1.
180 You have a task that isn't time-critical, but a current code path that is.
181 .RE
182 .RS +4
183 .TP
184 2.
185 You have a task that may require grabbing locks that a thread already holds.
186 .RE
187 .RS +4
188 .TP
189 3.
190 You have a task that needs to block (for example, to wait for memory), but you
190 You have a task that needs to block (for example, to wait for memory), but a
191 have a thread that cannot block in its current context.

```

```

192 .RE
193 .RS +4
194 .TP
195 4.
196 You have a code path that can't complete because of a specific condition,
197 but also can't sleep or fail. In this case, the task is immediately queued and
198 then is executed after the condition disappears.
199 .RE
200 .RS +4
201 .TP
202 5.
203 A task queue is just a simple way to launch multiple tasks in parallel.
204 .RE
205 .sp
206 .LP
207 A task queue consists of a list of tasks, together with one or more threads to
208 service the list. If a task queue has a single service thread, all tasks are
209 guaranteed to execute in the order they were dispatched. Otherwise they can be
210 executed in any order. Note that since tasks are placed on a list, execution of
211 one task should not depend on the execution of another task or a deadlock
212 may occur.
212 one task and should not depend on the execution of another task or a deadlock
212 may occur. A \fBtaskq\fR created with a single servicing thread guarantees that
213 all the tasks are serviced in the order in which they are scheduled.
213 .sp
214 .LP
215 The \fBddi_taskq_create()\fR function creates a task queue instance.
216 .sp
217 .LP
218 The \fBddi_taskq_dispatch()\fR function places \fBtaskq\fR on the list for
219 later execution. The \fBflag\fR argument specifies whether it is allowed sleep
220 waiting for memory. DDI_SLEEP dispatches can sleep and are guaranteed to
221 succeed. DDI_NOSLEEP dispatches are guaranteed not to sleep but may fail
222 (return \fBDDI_FAILURE\fR) if resources are not available.
223 .sp
224 .LP
225 The \fBddi_taskq_destroy()\fR function waits for any scheduled tasks to
226 complete, then destroys the \fBtaskq\fR. The caller should guarantee that no
227 new tasks are scheduled for the closing \fBtaskq\fR.
228 .sp
229 .LP
230 The \fBddi_taskq_wait()\fR function waits for all previously scheduled tasks to
231 complete. Note that this function does not stop any new task dispatches.
232 .sp
233 .LP
234 The \fBddi_taskq_suspend()\fR function suspends all task execution until
235 \fBddi_taskq_resume()\fR is called. Although \fBddi_taskq_suspend()\fR attempts
236 to suspend pending tasks, there are no guarantees that they will be suspended.
237 The only guarantee is that all tasks dispatched after \fBddi_taskq_suspend()\fR
238 will not be executed. Because it will trigger a deadlock, the
239 \fBddi_taskq_suspend()\fR function should never be called by a task executing
240 on a \fBtaskq\fR.
241 .sp
242 .LP
243 The \fBddi_taskq_suspended()\fR function returns \fBTRUE\fR if \fBtaskq\fR is
244 suspended, and \fBFALSE\fR otherwise. It is intended to ASSERT that the task
245 queue is suspended.
246 .sp
247 .LP
248 The \fBddi_taskq_resume()\fR function resumes task queue execution.
249 .SH RETURN VALUES
250 .sp
251 .LP
252 The \fBddi_taskq_create()\fR function creates an opaque handle that is used for
253 all other \fBtaskq\fR operations. It returns a \fBtaskq\fR pointer on success
254 and NULL on failure.

```

255 .sp  
256 .LP  
257 The \fBddi\_taskq\_dispatch()\fR function returns \fBDDI\_FAILURE\fR if it can't  
258 dispatch a task and returns \fBDDI\_SUCCESS\fR if dispatch succeeded.  
259 .sp  
260 .LP  
261 The \fBddi\_taskq\_suspended()\fR function returns \fBTRUE\fR if \fBtaskq\fR is  
262 suspended. Otherwise \fBFALSE\fR is returned.  
263 .SH CONTEXT  
264 .sp  
265 .LP  
266 All functions may be called from the user or kernel contexts.  
267 .sp  
268 .LP  
269 Additionally, the \fBddi\_taskq\_dispatch\fR function may be called from the  
270 interrupt context only if the DDI\_NOSLEEP flag is set.