

new/usr/src/man/man1m/share_nfs.1m

1

```
*****
19771 Sat Mar 22 14:25:25 2014
new/usr/src/man/man1m/share_nfs.1m
4398 Extra spaces in man pages
Reviewed by: Marcel Telka <marcel@telka.sk>
*****
1 \" te
2.\" Copyright (C) 2008, 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 SHARE_NFS 1M \"Mar 17, 2014\"
6.TH SHARE_NFS 1M \"May 6, 2009\"
7.SH NAME
8 share_nfs \- make local NFS file systems available for mounting by remote
9 systems
10.SH SYNOPSIS
11.LP
12.nf
13 \fBshare\fR [\fB-d\fR \fIdescription\fR] [\fB-F\fR nfs] [\fB-o\fR \fIspecific_op
14.fi

16.SH DESCRIPTION
17.sp
18.LP
19 The \fBshare\fR utility makes local file systems available for mounting by
20 remote systems. It starts the \fBnfsd\fR(1M) and \fBmountd\fR(1M) daemons if
21 they are not already running.
22.sp
23.LP
24 If no argument is specified, then \fBshare\fR displays all file systems
25 currently shared, including \fBNFS\fR file systems and file systems shared
26 through other distributed file system packages.
27.SH OPTIONS
28.sp
29.LP
30 The following options are supported:
31.sp
32.ne 2
33.na
34 \fB-d\fR \fIdescription\fR
35.ad
36.sp .6
37.RS 4n
38 Provide a comment that describes the file system to be shared.
39.RE

41.sp
42.ne 2
43.na
44 \fB-f\fR \fB-F\fR \fB-nfs\fR
45.ad
46.sp .6
47.RS 4n
48 Share \fBNFS\fR file system type.
49.RE

51.sp
52.ne 2
53.na
54 \fB-o\fR \fIspecific_options\fR
55.ad
56.sp .6
57.RS 4n
58 Specify \fIspecific_options\fR in a comma-separated list of keywords and
59 attribute-value-assertions for interpretation by the file-system-type-specific
```

new/usr/src/man/man1m/share_nfs.1m

2

```
60 command. If \fIspecific_options\fR is not specified, then by default sharing is
61 read-write to all clients. \fIspecific_options\fR can be any combination of the
62 following:
63.sp
64.ne 2
65.na
66 \fB-f\fR \fBaclok\fR
67.ad
68.sp .6
69.RS 4n
70 Allows the \fBNFS\fR server to do access control for \fBNFS\fR Version 2
71 clients (running SunOS 2.4 or earlier). When \fBaclok\fR is set on the server,
72 maximal access is given to all clients. For example, with \fBaclok\fR set, if
73 anyone has read permissions, then everyone does. If \fBaclok\fR is not set,
74 minimal access is given to all clients.
75.RE

77.sp
78.ne 2
79.na
80 \fB-f\fR \fBanon=\fR \fBuid\fR
81.ad
82.sp .6
83.RS 4n
84 Set \fBuid\fR to be the effective user \fBID\fR of unknown users. By default,
85 unknown users are given the effective user \fBID\fR \fBUID_NOBODY\fR. If
86 \fBuid\fR is set to \fB(mi1\fR, access is denied.
87.RE

89.sp
90.ne 2
91.na
92 \fB-f\fR \fBcharset\fR=\fBaccess_list\fR
93.ad
94.sp .6
95.RS 4n
96 Where \fBcharset\fR is one of: \fBBeuc-cn\fR, \fBBeuc-jp\fR, \fBBeuc-jpms\fR,
97 \fBBeuc-kr\fR, \fBBeuc-tw\fR, \fBiso8859-1\fR, \fBiso8859-2\fR, \fBiso8859-5\fR,
98 \fBiso8859-6\fR, \fBiso8859-7\fR, \fBiso8859-8\fR, \fBiso8859-9\fR,
99 \fBiso8859-13\fR, \fBiso8859-15\fR, \fBkoi8-r\fR.
100.sp
101 Clients that match the \fBaccess_list\fR for one of these properties will be
102 assumed to be using that character set and file and path names will be
103 converted to UTF-8 for the server.
104.RE

106.sp
107.ne 2
108.na
109 \fB-f\fR \fBindex=\fR \fBfile\fR
110.ad
111.sp .6
112.RS 4n
113 Load \fBfile\fR rather than a listing of the directory containing this file
114 when the directory is referenced by an \fBNFS URL\fR.
115.RE

117.sp
118.ne 2
119.na
120 \fB-f\fR \fBtag=\fR
121.ad
122.sp .6
123.RS 4n
124 Enables \fBNFS\fR server logging for the specified file system. The optional
125 tag determines the location of the related log files. The \fBtag\fR is defined
```

```

126 in \fBetc/nfs/nfslog.conf\fR. If no \fBtag\fR is specified, the default values
127 associated with the \fBglobal\fR \fBtag\fR in \fBetc/nfs/nfslog.conf\fR is
128 used. Support of NFS server logging is only available for NFS Version 2 and
129 Version 3 requests.
130 .RE

132 .sp
133 .ne 2
134 .na
135 \fB\fBnone=\fR\fIaccess_list\fR\fR
136 .ad
137 .sp .6
138 .RS 4n
139 Access is not allowed to any client that matches the access list. The exception
140 is when the access list is an asterisk (\fB*\fR), in which case \fBro\fR or
141 \fBrw\fR can override \fBnone\fR.
142 .RE

144 .sp
145 .ne 2
146 .na
147 \fB\fBnosub\fR\fR
148 .ad
149 .sp .6
150 .RS 4n
151 Prevents clients from mounting subdirectories of shared directories. For
152 example, if \fB/export\fR is shared with the \fBnosub\fR option on server
153 \fIfooey\fR then a \fBNFS\fR client cannot do:
154 .sp
155 .in +2
156 .nf
157 mount -F nfs fooey:/export/home/mnt
158 .fi
159 .in -2
160 .sp

162 NFS Version 4 does not use the \fBMOUNT\fR protocol. The \fBnosub\fR option
163 only applies to NFS Version 2 and Version 3 requests.
164 .RE

166 .sp
167 .ne 2
168 .na
169 \fB\fBnosuid\fR\fR
170 .ad
171 .sp .6
172 .RS 4n
173 By default, clients are allowed to create files on the shared file system with
174 the \fBsetuid\fR or \fBsetgid\fR mode enabled. Specifying \fBnosuid\fR causes the server
175 file system to silently ignore any attempt to enable the \fBsetuid\fR or \fBsetgid\fR mode
176 bits.
177 .RE

179 .sp
180 .ne 2
181 .na
182 \fB\fBpublic\fR\fR
183 .ad
184 .sp .6
185 .RS 4n
186 Moves the location of the public file handle from \fBroot\fR (\fB/\fR) to the
187 exported directory for \fBWeb\fBNFS\fR-enabled browsers and clients. This option
188 does not enable \fBWeb\fBNFS\fR service; \fBWeb\fBNFS\fR is always on. Only one file
189 system per server may use this option. Any other option, including the
190 \fB-ro=list\fR and \fB-rw=list\fR options can be included with the \fBpublic\fR
191 option.

```

```

192 .RE

194 .sp
195 .ne 2
196 .na
197 \fB\fBro\fR\fR
198 .ad
199 .sp .6
200 .RS 4n
201 Sharing is read-only to all clients.
202 .RE

204 .sp
205 .ne 2
206 .na
207 \fB\fBro=\fR\fIaccess_list\fR\fR
208 .ad
209 .sp .6
210 .RS 4n
211 Sharing is read-only to the clients listed in \fIaccess_list\fR; overrides the
212 \fBrw\fR suboption for the clients specified. See \fIaccess_list\fR below.
213 .RE

215 .sp
216 .ne 2
217 .na
218 \fB\fBroot=\fR\fIaccess_list\fR\fR
219 .ad
220 .sp .6
221 .RS 4n
222 Only root users from the hosts specified in \fIaccess_list\fR have root access.
223 See \fIaccess_list\fR below. By default, no host has root access, so root users
224 are mapped to an anonymous user \fBID\fR (see the \fBanon=\fR\fIuid\fR option
225 described above). Netgroups can be used if the file system shared is using UNIX
226 authentication (\fBAUTH_SYS\fR).
226 authentication ( \fBAUTH_SYS\fR).
227 .RE

229 .sp
230 .ne 2
231 .na
232 \fB\fBroot_mapping=\fR\fIuid\fR\fR
233 .ad
234 .sp .6
235 .RS 4n
236 For a client that is allowed root access, map the root UID to the specified
237 user id.
238 .RE

240 .sp
241 .ne 2
242 .na
243 \fB\fBrw\fR\fR
244 .ad
245 .sp .6
246 .RS 4n
247 Sharing is read-write to all clients.
248 .RE

250 .sp
251 .ne 2
252 .na
253 \fB\fBrw=\fR\fIaccess_list\fR\fR
254 .ad
255 .sp .6
256 .RS 4n

```

257 Sharing is read-write to the clients listed in \fIaccess_list\fR; overrides the
 258 \fBro\fR suboption for the clients specified. See \fIaccess_list\fR below.
 259 .RE

261 .sp
 262 .ne 2
 263 .na
 264 \fB\fBsec=\fR\fImode\fR[\fB:\fR\fImode\fR].\|.\|.\fR
 265 .ad
 266 .sp .6
 267 .RS 4n
 268 Sharing uses one or more of the specified security modes. The \fImode\fR in the
 269 \fBsec=\fR\fImode\fR option must be a node name supported on the client. If the
 270 \fBsec=\fR option is not specified, the default security mode used is
 271 \fBAUTH_SYS.\fR Multiple \fBsec=\fR options can be specified on the command
 272 line, although each mode can appear only once. The security modes are defined
 273 in \fBnfssec\fR(5).
 274 .sp
 275 Each \fBsec=\fR option specifies modes that apply to any subsequent \fBwindow=,
 276 rw, ro, rw=, ro=\fR and \fBroot=\fR options that are provided before another
 277 \fBsec=\fR option. Each additional \fBsec=\fR resets the security mode context,
 278 so that more \fBwindow=,\fR \fBrw,\fR \fR \fBro,\fR \fBrw=,\fR \fBro=\fR and
 279 \fBroot=\fR options can be supplied for additional modes.
 280 .RE

282 .sp
 283 .ne 2
 284 .na
 285 \fB\fBsec=\fR\fInone\fR\fR
 286 .ad
 287 .sp .6
 288 .RS 4n
 289 If the option \fBsec=\fR\fInone\fR is specified when the client uses
 290 \fBAUTH_NONE,\fR or if the client uses a security mode that is not one that the
 291 file system is shared with, then the credential of each \fBNFS\fR request is
 292 treated as unauthenticated. See the \fBanon=\fR\fIuid\fR option for a
 293 description of how unauthenticated requests are handled.
 294 .RE

296 .sp
 297 .ne 2
 298 .na
 299 \fB\fBsecure\fR\fR
 300 .ad
 301 .sp .6
 302 .RS 4n
 303 This option has been deprecated in favor of the \fBsec=\fR\fRIdh\fR option.
 304 .RE

306 .sp
 307 .ne 2
 308 .na
 309 \fB\fBwindow=\fR\fRvalue\fR\fR
 310 .ad
 311 .sp .6
 312 .RS 4n
 313 When sharing with \fBsec=\fR\fRIdh\fR, set the maximum life time (in seconds) of
 314 the \fBRPC\fR request's credential (in the authentication header) that the
 315 \fBNFS\fR server allows. If a credential arrives with a life time larger than
 316 what is allowed, the \fBNFS\fR server rejects the request. The default value is
 317 30000 seconds (8.3 hours).
 318 .RE

320 .RE

322 .SS "\fIaccess_list\fR"

323 .sp
 324 .LP
 325 The \fIaccess_list\fR argument is a colon-separated list whose components may
 326 be any number of the following:
 327 .sp
 328 .ne 2
 329 .na
 330 \fBhostname\fR
 331 .ad
 332 .sp .6
 333 .RS 4n
 334 The name of a host. With a server configured for \fBDNS\fR or \fBLDAP\fR naming
 335 in the \fBnsswitch\fR "hosts" entry, any hostname must be represented as a
 336 fully qualified \fBDNS\fR or \fBLDAP\fR name.
 337 .RE

339 .sp
 340 .ne 2
 341 .na
 342 \fBnetgroup\fR
 343 .ad
 344 .sp .6
 345 .RS 4n
 346 A netgroup contains a number of hostnames. With a server configured for
 347 \fBDNS\fR or \fBLDAP\fR naming in the \fBnsswitch\fR "hosts" entry, any
 348 hostname in a netgroup must be represented as a fully qualified \fBDNS\fR or
 349 \fBLDAP\fR name.
 350 .RE

352 .sp
 353 .ne 2
 354 .na
 355 \fBdomain name suffix\fR
 356 .ad
 357 .sp .6
 358 .RS 4n
 359 To use domain membership the server must use \fBDNS\fR or \fBLDAP\fR to resolve
 360 hostnames to \fBIP\fR addresses; that is, the "hosts" entry in the
 361 \fB/etc/nsswitch.conf\fR must specify "dns" or "ldap" ahead of "nis" or
 362 "nisplus", since only \fBDNS\fR and \fBLDAP\fR return the full domain name of
 363 the host. Other name services like \fBNIS\fR or \fBNIS+\fR cannot be used to
 364 resolve hostnames on the server because when mapping an \fBIP\fR address to a
 365 hostname they do not return domain information. For example,
 366 .sp
 367 .in +2
 368 .nf
 369 NIS or LDAP 172.16.45.9 --> "myhost"
 370 .fi
 371 .in -2
 372 .sp

374 and
 375 .sp
 376 .in +2
 377 .nf
 378 DNS or LDAP 172.16.45.9 -->
 379 "myhost.mydomain.mycompany.com"
 380 .fi
 381 .in -2
 382 .sp

384 The domain name suffix is distinguished from hostnames and netgroups by a
 385 prefixed dot. For example,
 386 .sp
 387 \fBrw=.mydomain.mycompany.com\fR
 388 .sp

```

389 A single dot can be used to match a hostname with no suffix. For example,
390 .sp
391 \fBBrw=\fR
392 .sp
393 matches "mydomain" but not "mydomain.mycompany.com". This feature can be used
394 to match hosts resolved through \fBNIS\fR and \fBNIS+\fR rather than \fBDNS\fR
395 and \fBLDAP\fR.
396 .RE

398 .sp
399 .ne 2
400 .na
401 \fBnetwork\fR
402 .ad
403 .sp .6
404 .RS 4n
405 The network or subnet component is preceded by an at-sign (\fB@\fR). It can be
406 either a name or a dotted address. If a name, it is converted to a dotted
407 address by \fBgetnetbyname\fR(3SOCKET). For example,
408 .sp
409 \fB=@mynet\fR
410 .sp
411 would be equivalent to:
412 .sp
413 \fB=@172.16\fR or \fB=@172.16.0.0\fR
414 .sp
415 The network prefix assumes an octet-aligned netmask determined from the zeroth
416 octet in the low-order part of the address up to and including the high-order
417 octet, if you want to specify a single IP address (see below). In the case
418 where network prefixes are not byte-aligned, the syntax allows a mask length to
419 be specified explicitly following a slash (\fB/\fR) delimiter. For example,
420 .sp
421 \fB=@theothernet/17\fR or \fB=@172.16.132/22\fR
422 .sp
423 \&...where the mask is the number of leftmost contiguous significant bits in
424 the corresponding IP address.
425 .sp
426 When specifying individual IP addresses, use the same \fB@\fR notation
427 described above, without a netmask specification. For example:
428 .sp
429 .in +2
430 .nf
431 =@172.16.132.14
432 .fi
433 .in -2
434 .sp

436 Multiple, individual IP addresses would be specified, for example, as:
437 .sp
438 .in +2
439 .nf
440 root=@172.16.132.20:@172.16.134.20
441 .fi
442 .in -2
443 .sp

445 .RE

447 .sp
448 .LP
449 A prefixed minus sign (\fB(mi\fR) denies access to that component of
450 \fBfiaccess_list\fR. The list is searched sequentially until a match is found
451 that either grants or denies access, or until the end of the list is reached.
452 For example, if host "terra" is in the "engineering" netgroup, then
453 .sp
454 .in +2

```

```

455 .nf
456 rw=-terra:engineering
457 .fi
458 .in -2
459 .sp

461 .sp
462 .LP
463 denies access to \fBterra\fR but
464 .sp
465 .in +2
466 .nf
467 rw=engineering:-terra
468 .fi
469 .in -2
470 .sp

472 .sp
473 .LP
474 grants access to \fBterra\fR.
475 .SH OPERANDS
476 .sp
477 .LP
478 The following operands are supported:
479 .sp
480 .ne 2
481 .na
482 \fBfipathname\fR\fR
483 .ad
484 .sp .6
485 .RS 4n
486 The pathname of the file system to be shared.
487 .RE

489 .SH EXAMPLES
490 .LP
491 \fBExample 1 \fRSharing A File System With Logging Enabled
492 .sp
493 .LP
494 The following example shows the \fB/export\fR file system shared with logging
495 enabled:

497 .sp
498 .in +2
499 .nf
500 example% \fBshare -o log /export\fR
501 .fi
502 .in -2
503 .sp

505 .sp
506 .LP
507 The default global logging parameters are used since no tag identifier is
508 specified. The location of the log file, as well as the necessary logging work
509 files, is specified by the global entry in \fB/etc/nfs/nfslog.conf\fR. The
510 \fBnfslogd\fR(1M) daemon runs only if at least one file system entry in
511 \fB/etc/dfs/dfstab\fR is shared with logging enabled upon starting or rebooting
512 the system. Simply sharing a file system with logging enabled from the command
513 line does not start the \fBnfslogd\fR(1M).

515 .SH EXIT STATUS
516 .sp
517 .LP
518 The following exit values are returned:
519 .sp
520 .ne 2

```

```

521 .na
522 \fB\FB0\fR\fR
523 .ad
524 .sp .6
525 .RS 4n
526 Successful completion.
527 .RE

529 .sp
530 .ne 2
531 .na
532 \fB\FB>0\fR\fR
533 .ad
534 .sp .6
535 .RS 4n
536 An error occurred.
537 .RE

539 .SH FILES
540 .sp
541 .ne 2
542 .na
543 \fB\FB/etc/dfs/fstypes\fR\fR
544 .ad
545 .sp .6
546 .RS 4n
547 list of system types, \fBNFS\fR by default
548 .RE

550 .sp
551 .ne 2
552 .na
553 \fB\FB/etc/dfs/sharetab\fR\fR
554 .ad
555 .sp .6
556 .RS 4n
557 system record of shared file systems
558 .RE

560 .sp
561 .ne 2
562 .na
563 \fB\FB/etc/nfs/nfslogtab\fR\fR
564 .ad
565 .sp .6
566 .RS 4n
567 system record of logged file systems
568 .RE

570 .sp
571 .ne 2
572 .na
573 \fB\FB/etc/nfs/nfslog.conf\fR\fR
574 .ad
575 .sp .6
576 .RS 4n
577 logging configuration file
578 .RE

580 .SH SEE ALSO
581 .sp
582 .LP
583 \fBmount\fR(1M), \fBmountd\fR(1M), \fBbnfsd\fR(1M), \fBbnfslogd\fR(1M),
584 \fBbshare\fR(1M), \fBbunshare\fR(1M), \fBbgetnetbyname\fR(3SOCKET),
585 \fBbnfslog.conf\fR(4), \fBbnetgroup\fR(4), \fBbattributes\fR(5), \fBbnfssec\fR(5)
586 .SH NOTES

```

```

587 .sp
588 .LP
589 If the \fBsec=\fR option is presented at least once, all uses of the
590 \fBwindow=,\fR \fBrw,\fR \fBro,\fR \fBrw=,\fR \fBro=\fR and \fBroot=\fR options
591 must come \fBafter\fR the first \fBsec=\fR option. If the \fBsec=\fR option is
592 not presented, then \fBsec=\fR\fR is implied.
593 .sp
594 .LP
595 If one or more explicit \fBsec=\fR options are presented, \fR\fR must appear
596 in one of the options mode lists for accessing using the \fBAUTH_SYS\fR
597 security mode to be allowed. For example:
598 .sp
599 .in +2
600 .nf
601 \fBshare\fR \fB-F\fR \fBnfs /var\fR
602 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBsec=sys /var\fR
603 .fi
604 .in -2
605 .sp

607 .sp
608 .LP
609 grants read-write access to any host using \fBAUTH_SYS,\fR but
610 .sp
611 .in +2
612 .nf
613 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBsec=dh /var\fR
614 .fi
615 .in -2
616 .sp

618 .sp
619 .LP
620 grants no access to clients that use \fBAUTH_SYS.\fR
621 .sp
622 .LP
623 Unlike previous implementations of \fBshare_nfs\fR, access checking for the
624 \fBwindow=, rw, ro, rw=,\fR and \fBro=\fR options is done per \fBNFS\fR
625 request, instead of per mount request.
626 .sp
627 .LP
628 Combining multiple security modes can be a security hole in situations where
629 the \fBro=\fR and \fBrw=\fR options are used to control access to weaker
630 security modes. In this example,
631 .sp
632 .in +2
633 .nf
634 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBsec=dh,rw,sec=sys,rw=hosta /var\fR
635 .fi
636 .in -2
637 .sp

639 .sp
640 .LP
641 an intruder can forge the IP address for \fBhosta\fR (albeit on each \fBNFS\fR
642 request) to side-step the stronger controls of \fBAUTH_DES.\fR Something like:
643 .sp
644 .in +2
645 .nf
646 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBsec=dh,rw,sec=sys,ro /var\fR
647 .fi
648 .in -2
649 .sp

651 .sp
652 .LP

```

```

653 is safer, because any client (intruder or legitimate) that avoids
654 \fBAUTH_DES\fR only gets read-only access. In general, multiple security modes
655 per \fBshare\fR command should only be used in situations where the clients
656 using more secure modes get stronger access than clients using less secure
657 modes.
658 .sp
659 .LP
660 If \fBBrw=\fR and \fBro=\fR options are specified in the same \fBsec=\fR
661 clause, and a client is in both lists, the order of the two options determines
662 the access the client gets. If client \fBhosta\fR is in two netgroups -
663 \fBgroup1\fR and \fBgroup2\fR - in this example, the client would get read-only
664 access:
665 .sp
666 .in +2
667 .nf
668 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBro=group1,rw=group2 /var\fR
669 .fi
670 .in -2
671 .sp

673 .sp
674 .LP
675 In this example \fBhosta\fR would get read-write access:
676 .sp
677 .in +2
678 .nf
679 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBrw=group2,ro=group1 /var\fR
680 .fi
681 .in -2
682 .sp

684 .sp
685 .LP
686 If within a \fBsec=\fR clause, both the \fBro\fR and \fBBrw=\fR options are
687 specified, for compatibility, the order of the options rule is not enforced.
688 All hosts would get read-only access, with the exception to those in the
689 read-write list. Likewise, if the \fBro=\fR and \fBBrw\fR options are specified,
690 all hosts get read-write access with the exceptions of those in the read-only
691 list.
692 .sp
693 .LP
694 The \fBro=\fR and \fBBrw=\fR options are guaranteed to work over \fBUDP\fR and
695 \fBTCP\fR but may not work over other transport providers.
696 .sp
697 .LP
698 The \fBroot=\fR option with \fBAUTH_SYS\fR is guaranteed to work over \fBUDP\fR
699 and \fBTCP\fR but may not work over other transport providers.
700 .sp
701 .LP
702 The \fBroot=\fR option with \fBAUTH_DES\fR is guaranteed to work over any
703 transport provider.
704 .sp
705 .LP
706 There are no interactions between the \fBroot=\fR option and the \fBBrw, ro,
707 rw,\fR and \fBro=\fR options. Putting a host in the \fBroot\fR list does not
708 override the semantics of the other options. The access the host gets is the
709 same as when the \fBroot=\fR options is absent. For example, the following
710 \fBshare\fR command denies access to \fBhostb\fR
711 .sp
712 .in +2
713 .nf
714 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBro=hosta,root=hostb /var\fR
715 .fi
716 .in -2
717 .sp

```

```

719 .sp
720 .LP
721 The following gives read-only permissions to \fBhostb\fR
722 .sp
723 .in +2
724 .nf
725 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBro=hostb,root=hostb /var\fR
726 .fi
727 .in -2
728 .sp

730 .sp
731 .LP
732 The following gives read-write permissions to \fBhostb\fR
733 .sp
734 .in +2
735 .nf
736 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBro=hosta,rw=hostb,root=hostb /var\fR
737 .fi
738 .in -2
739 .sp

741 .sp
742 .LP
743 If the file system being shared is a symbolic link to a valid pathname, the
744 canonical path (the path which the symbolic link follows) are shared. For
745 example, if \fB/export/foo\fR is a symbolic link to \fB/export/bar\fR
746 (\fB/export/foo -> /export/bar\fR), the following \fBshare\fR command results
747 in \fB/export/bar\fR as the shared pathname (and not \fB/export/foo\fR).
748 .sp
749 .in +2
750 .nf
751 \fBexample# share\fR \fB-F\fR \fBnfs /export/foo\fR
752 .fi
753 .in -2
754 .sp

756 .sp
757 .LP
758 An \fB NFS mount of \fBserver:/export/foo\fR results in
759 \fBserver:/export/bar\fR really being mounted.
760 .sp
761 .LP
762 This line in the \fB/etc/dfs/dfstab\fR file shares the \fB/disk\fR file system
763 read-only at boot time:
764 .sp
765 .in +2
766 .nf
767 \fBshare\fR \fB-F\fR \fBnfs\fR \fB-o\fR \fBro /disk\fR
768 .fi
769 .in -2
770 .sp

772 .sp
773 .LP
774 The same command entered from the command line does not share the \fB/disk\fR
775 file system unless there is at least one file system entry in the
776 \fB/etc/dfs/dfstab\fR file. The \fBmountd(1M)\fR and \fBnfsd(1M)\fR daemons
777 only run if there is a file system entry in \fB/etc/dfs/dfstab\fR when starting
778 or rebooting the system.
779 .sp
780 .LP
781 The \fBmountd(1M)\fR process allows the processing of a path name the contains
782 a symbolic link. This allows the processing of paths that are not themselves
783 explicitly shared with \fBshare_nfs\fR. For example, \fB/export/foo\fR might be
784 a symbolic link that refers to \fB/export/bar\fR which has been specifically

```

785 shared. When the client mounts \fB/export/foo\fR the \fBmountd\fR processing
786 follows the symbolic link and responds with the \fB/export/bar\fR. The NFS
787 Version 4 protocol does not use the \fBmountd\fR processing and the client's
788 use of \fB/export/foo\fR does not work as it does with NFS Version 2 and
789 Version 3 and the client receives an error when attempting to mount
790 \fB/export/foo\fR.

6810 Sat Mar 22 14:25:25 2014
new/usr/src/man/man3nsl/xdr_admin.3nsl

4398 Extra spaces in man pages

Reviewed by: Marcel Telka <marcel@telka.sk>

```

1  \" te
2  \" Copyright 1989 AT&T Copyright (c) 1997, Sun Microsystems, Inc. All Rights
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 XDR_ADMIN 3NSL \"Mar 17, 2014\"
6  .TH XDR_ADMIN 3NSL \"Dec 30, 1996\"
7  .SH NAME
8  xdr_admin, xdr_control, xdr_getpos, xdr_inline, xdrrec_endofrecord, xdrrec_eof,
9  xdrrec_readbytes, xdrrec_skiprecord, xdr_setpos, xdr_sizeof - library routines
10 for external data representation
11 .SH DESCRIPTION
12 .sp
13 .LP
14 \fBXDR\fR library routines allow C programmers to describe arbitrary data
15 structures in a machine-independent fashion. Protocols such as remote procedure
16 calls (RPC) use these routines to describe the format of the data.
17 .sp
18 .LP
19 These routines deal specifically with the management of the \fBXDR\fR stream.
20 .SS \"Routines\"
21 .sp
22 .LP
23 See \fBrpc\fR(3NSL) for the definition of the \fBXDR\fR data structure. Note
24 that any buffers passed to the \fBXDR\fR routines must be properly aligned. It
25 is suggested either that \fBmalloc\fR(3C) be used to allocate these buffers, or
26 that the programmer insure that the buffer address is divisible evenly by
27 four.
28 .sp
29 .LP
30 \fB#include <rpc/xdr.h>\fR
31 .sp
32 .ne 2
33 .na
34 \fB\fBbool_t xdr_control(XDR *\fR\fIxdrs\fR\fB, int\fR\fI req\fR\fB, void
34 \fB\fBbool_t xdr_control( XDR *\fR\fIxdrs\fR\fB, int\fR\fI req\fR\fB, void
35 *\fR\fIinfo\fR\fB);\fR
36 .ad
37 .sp .6
38 .RS 4n
39 A function macro to change or retrieve various information about an \fBXDR\fR
40 stream. \fIreq\fR indicates the type of operation and \fIinfo\fR is a pointer
41 to the information. The supported values of \fIreq\fR is
42 \fBXDR_GET_BYTES_AVAIL\fR and its argument type is \fBxdr_bytesrec *\fR. They
43 return the number of bytes left unconsumed in the stream and a flag indicating
44 whether or not this is the last fragment.
45 .RE

47 .sp
48 .ne 2
49 .na
50 \fB\fBuint_t xdr_getpos(const XDR *\fR\fIxdrs\fR\fB);\fR
51 .ad
52 .sp .6
53 .RS 4n
54 A macro that invokes the get-position routine associated with the \fBXDR\fR
55 stream, \fIxdrs\fR. The routine returns an unsigned integer, which indicates
56 the position of the \fBXDR\fR byte stream. A desirable feature of \fBXDR\fR
57 streams is that simple arithmetic works with this number, although the
58 \fBXDR\fR stream instances need not guarantee this. Therefore, applications

```

59 written for portability should not depend on this feature.
60 .RE

```

62 .sp
63 .ne 2
64 .na
65 \fB\fBlong *xdr_inline(XDR *\fR\fIxdrs\fR\fB, const int \fR\fIlen\fR\fB);\fR
66 .ad
67 .sp .6
68 .RS 4n
69 A macro that invokes the in-line routine associated with the \fBXDR\fR stream,
70 \fIxdrs\fR. The routine returns a pointer to a contiguous piece of the stream's
71 buffer; \fIlen\fR is the byte length of the desired buffer. Note: pointer is
72 cast to \fBlong *\fR.
73 .sp
74 Warning: \fBxdr_inline()\fR may return \fBNULL\fR (\fB0\fR) if it cannot
75 allocate a contiguous piece of a buffer. Therefore the behavior may vary among
76 stream instances; it exists for the sake of efficiency, and applications
77 written for portability should not depend on this feature.
78 .RE

```

```

80 .sp
81 .ne 2
82 .na
83 \fB\fBbool_t xdrrec_endofrecord(XDR *xdrs, int \fR\fIsendnow\fR\fB);\fR
84 .ad
85 .sp .6
86 .RS 4n
87 This routine can be invoked only on streams created by \fBxdrrec_create()\fR.
88 See \fBxdr_create\fR(3NSL). The data in the output buffer is marked as a
89 completed record, and the output buffer is optionally written out if
90 \fIsendnow\fR is non-zero. This routine returns \fBTRUE\fR if it succeeds,
91 \fBFALSE\fR otherwise.
92 .RE

```

```

94 .sp
95 .ne 2
96 .na
97 \fB\fBbool_t xdrrec_eof(XDR *\fR\fIxdrs\fR\fB);\fR
98 .ad
99 .sp .6
100 .RS 4n
101 This routine can be invoked only on streams created by \fBxdrrec_create()\fR.
102 After consuming the rest of the current record in the stream, this routine
103 returns \fBTRUE\fR if there is no more data in the stream's input buffer. It
104 returns \fBFALSE\fR if there is additional data in the stream's input buffer.
105 .RE

```

```

107 .sp
108 .ne 2
109 .na
110 \fB\fBint xdrrec_readbytes(XDR *\fR\fIxdrs\fR\fB, caddr_t\fR\fI addr\fR\fB,
111 uint_t\fR\fI nbytes\fR\fB);\fR
112 .ad
113 .sp .6
114 .RS 4n
115 This routine can be invoked only on streams created by \fBxdrrec_create()\fR.
116 It attempts to read \fInbytes\fR bytes from the \fBXDR\fR stream into the
117 buffer pointed to by \fIaddr\fR. Upon success this routine returns the number
118 of bytes read. Upon failure, it returns \fB(mii)\fR. A return value of
119 \fB0\fR indicates an end of record.
120 .RE

```

```

122 .sp
123 .ne 2
124 .na

```

```
125 \fB\fBbool_t xdrrec_skiprecord(XDR *\fR\fIxdrs\fR\fB);\fR\fR
126 .ad
127 .sp .6
128 .RS 4n
129 This routine can be invoked only on streams created by \fBxdrrec_create()\fR.
130 See \fBxdr_create\fR(3NSL). It tells the \fBXHR\fR implementation that the
131 rest of the current record in the stream's input buffer should be discarded.
132 This routine returns \fBTRUE\fR if it succeeds, \fBFALSE\fR otherwise.
133 .RE

135 .sp
136 .ne 2
137 .na
138 \fB\fBbool_t xdr_setpos(XDR *\fR\fIxdrs\fR\fB, const uint_t
139 \fR\fIpos\fR\fB);\fR\fR
140 .ad
141 .sp .6
142 .RS 4n
143 A macro that invokes the set position routine associated with the \fBXHR\fR
144 stream \fIxdrs\fR. The parameter \fIpos\fR is a position value obtained from
145 \fBxdr_getpos()\fR. This routine returns \fBTRUE\fR if the \fBXHR\fR stream was
146 repositioned, and \fBFALSE\fR otherwise.
147 .sp
148 Warning: it is difficult to reposition some types of \fBXHR\fR streams, so this
149 routine may fail with one type of stream and succeed with another. Therefore,
150 applications written for portability should not depend on this feature.
151 .RE

153 .sp
154 .ne 2
155 .na
156 \fB\fBunsigned long xdr_sizeof(xdrproc_t \fR\fIfunc\fR\fB, void
157 *\fR\fIdata\fR\fB);\fR\fR
158 .ad
159 .sp .6
160 .RS 4n
161 This routine returns the number of bytes required to encode \fIdata\fR using
162 the \fBXHR\fR filter function \fIfunc\fR, excluding potential overhead such as
163 \fBRPC\fR headers or record markers. \fB0\fR is returned on error. This
164 information might be used to select between transport protocols, or to
165 determine the buffer size for various lower levels of \fBRPC\fR client and
166 server creation routines, or to allocate storage when \fBXHR\fR is used
167 outside of the \fBRPC\fR subsystem.
168 .RE

170 .SH ATTRIBUTES
171 .sp
172 .LP
173 See \fBattributes\fR(5) for descriptions of the following attributes:
174 .sp

176 .sp
177 .TS
178 box;
179 c | c
180 l | l .
181 ATTRIBUTE TYPE    ATTRIBUTE VALUE
182 _
183 MT-Level          Safe
184 .TE

186 .SH SEE ALSO
187 .sp
188 .LP
189 \fBmalloc\fR(3C), \fBbrpc\fR(3NSL), \fBxdr_complex\fR(3NSL),
190 \fBxdr_create\fR(3NSL), \fBxdr_simple\fR(3NSL), \fBattributes\fR(5)
```

19590 Sat Mar 22 14:25:26 2014

new/usr/src/man/man7p/tcp.7p

4398 Extra spaces in man pages

Reviewed by: Marcel Telka <marcel@telka.sk>

```

1  \" te
2  \" Copyright (c) 2006, Sun Microsystems, Inc. All Rights Reserved.
3  \" Copyright (c) 2011 Nexenta Systems, Inc. All rights reserved.
4  \" Copyright 1989 AT&T
5  \" The contents of this file are subject to the terms of the Common Development
6  \" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
7  \" When distributing Covered Code, include this CDDL HEADER in each file and in
8  .TH TCP 7P \"Mar 17, 2014\"
9  .TH TCP 7P \"Jun 30, 2006\"
10 .SH NAME
11 .SH SYNOPSIS
12 .LP
13 .nf
14 \fB#include <sys/socket.h>\fR
15 .fi

17 .LP
18 .nf
19 \fB#include <netinet/in.h>\fR
20 .fi

22 .LP
23 .nf
24 \fBs = socket(AF_INET, SOCK_STREAM, 0);\fR
25 .fi

27 .LP
28 .nf
29 \fBs = socket(AF_INET6, SOCK_STREAM, 0);\fR
30 .fi

32 .LP
33 .nf
34 \fBt = t_open(\"/dev/tcp\", O_RDWR);\fR
35 .fi

37 .LP
38 .nf
39 \fBt = t_open(\"/dev/tcp6\", O_RDWR);\fR
40 .fi

42 .SH DESCRIPTION
43 .sp
44 .LP
45 \fBTCP\fR is the virtual circuit protocol of the Internet protocol family. It
46 provides reliable, flow-controlled, in order, two-way transmission of data. It
47 is a byte-stream protocol layered above the Internet Protocol (\fBIP\fR), or
48 the Internet Protocol Version 6 (\fBIPv6\fR), the Internet protocol family's
49 internetwork datagram delivery protocol.
50 .sp
51 .LP
52 Programs can access \fBTCP\fR using the socket interface as a \fBFSOCK_STREAM\fR
53 socket type, or using the Transport Level Interface (\fBTLI\fR) where it
54 supports the connection-oriented (\fBT_COTS_ORD\fR) service type.
55 .sp
56 .LP
57 \fBTCP\fR uses \fBIP\fR's host-level addressing and adds its own per-host
58 collection of "port addresses." The endpoints of a \fBTCP\fR connection are
59 identified by the combination of an \fBIP\fR or IPv6 address and a \fBTCP\fR

```

```

60 port number. Although other protocols, such as the User Datagram Protocol
61 (UDP), may use the same host and port address format, the port space of these
62 protocols is distinct. See \fBinet\fR(7P) and \fBinet6\fR(7P) for details on
63 the common aspects of addressing in the Internet protocol family.
64 .sp
65 .LP
66 Sockets utilizing \fBTCP\fR are either "active" or "passive." Active sockets
67 initiate connections to passive sockets. Both types of sockets must have their
68 local \fBIP\fR or IPv6 address and \fBTCP\fR port number bound with the
69 \fBbind\fR(3SOCKET) system call after the socket is created. By default,
70 \fBTCP\fR sockets are active. A passive socket is created by calling the
71 \fBlisten\fR(3SOCKET) system call after binding the socket with \fBbind()\fR.
72 This establishes a queueing parameter for the passive socket. After this,
73 connections to the passive socket can be received with the
74 \fBaccept\fR(3SOCKET) system call. Active sockets use the
75 \fBconnect\fR(3SOCKET) call after binding to initiate connections.
76 .sp
77 .LP
78 By using the special value \fBINADDR_ANY\fR with \fBIP\fR, or the unspecified
79 address (all zeroes) with IPv6, the local \fBIP\fR address can be left
80 unspecified in the \fBbind()\fR call by either active or passive \fBTCP\fR
81 sockets. This feature is usually used if the local address is either unknown or
82 irrelevant. If left unspecified, the local \fBIP\fR or IPv6 address will be
83 bound at connection time to the address of the network interface used to
84 service the connection.
85 .sp
86 .LP
87 Note that no two TCP sockets can be bound to the same port unless the bound IP
88 addresses are different. IPv4 \fBINADDR_ANY\fR and IPv6 unspecified addresses
89 compare as equal to any IPv4 or IPv6 address. For example, if a socket is bound
90 to \fBINADDR_ANY\fR or unspecified address and port X, no other socket can bind
91 to port X, regardless of the binding address. This special consideration of
92 \fBINADDR_ANY\fR and unspecified address can be changed using the socket option
93 \fBBSO_REUSEADDR\fR. If \fBBSO_REUSEADDR\fR is set on a socket doing a bind, IPv4
94 \fBINADDR_ANY\fR and IPv6 unspecified address do not compare as equal to any IP
95 address. This means that as long as the two sockets are not both bound to
96 \fBINADDR_ANY\fR/unspecified address or the same IP address, the two sockets
97 can be bound to the same port.
98 .sp
99 .LP
100 If an application does not want to allow another socket using the
101 If an application does not want to allow another socket using the
102 \fBBSO_REUSEADDR\fR option to bind to a port its socket is bound to, the
103 application can set the socket level option \fBBSO_EXCLBIND\fR on a socket. The
104 option values of 0 and 1 mean enabling and disabling the option respectively.
105 Once this option is enabled on a socket, no other socket can be bound to the
106 same port.
107 .sp
108 .LP
109 Once a connection has been established, data can be exchanged using the
110 \fBread\fR(2) and \fBwrite\fR(2) system calls.
111 .sp
112 .LP
113 Under most circumstances, \fBTCP\fR sends data when it is presented. When
114 outstanding data has not yet been acknowledged, \fBTCP\fR gathers small amounts
115 of output to be sent in a single packet once an acknowledgement has been
116 received. For a small number of clients, such as window systems that send a
117 stream of mouse events which receive no replies, this packetization may cause
118 significant delays. To circumvent this problem, \fBTCP\fR provides a
119 socket-level boolean option, \fBTCP_NODELAY\fR. \fBTCP_NODELAY\fR is defined in
120 \fB<netinet/tcp.h>\fR, and is set with \fBsetsockopt\fR(3SOCKET) and tested
121 with \fBgetsockopt\fR(3SOCKET). The option level for the \fBsetsockopt()\fR
122 call is the protocol number for \fBTCP\fR available from
123 \fBgetprotobyname\fR(3SOCKET).
124 .sp
125 .LP

```

125 For some applications, it may be desirable for TCP not to send out data unless
 126 a full TCP segment can be sent. To enable this behavior, an application can use
 127 the `\fBTCP_CORK\fR` socket option. When `\fBTCP_CORK\fR` is set with a non-zero
 128 value, TCP sends out a full TCP segment only. When `\fBTCP_CORK\fR` is set to
 129 zero after it has been enabled, all buffered data is sent out (as permitted by
 130 the peer's receive window and the current congestion window). `\fBTCP_CORK\fR` is
 131 defined in `<\fBnetinet/tcp.h\fR>`, and is set with `\fBsetsockopt\fR(3SOCKET)`
 132 and tested with `\fBgetsockopt\fR(3SOCKET)`. The option level for the
 133 `\fBsetsockopt()` `\fR` call is the protocol number for TCP, available from
 134 `\fBgetprotobyname\fR(3SOCKET)`.
 135 .sp
 136 .LP
 137 The `SO_RCVBUF` socket level option can be used to control the window that TCP
 138 advertises to the peer. IP level options may also be used with TCP. See
 139 `\fBip\fR(7P)` and `\fBip6\fR(7P)`.
 140 .sp
 141 .LP
 142 Another socket level option, `\fBSO_RCVBUF,\fR` can be used to control the window
 143 that `\fBTCP\fR` advertises to the peer. `\fBIP\fR` level options may also be used
 144 with `\fBTCP,\fR` See `\fBip\fR(7P)` and `\fBip6\fR(7P)`.
 145 .sp
 146 .LP
 147 `\fBTCP\fR` provides an urgent data mechanism, which may be invoked using the
 148 out-of-band provisions of `\fBsend\fR(3SOCKET)`. The caller may mark one byte as
 149 "urgent" with the `\fBMSG_OOB\fR` flag to `\fBsend\fR(3SOCKET)`. This sets an
 150 "urgent pointer" pointing to this byte in the `\fBTCP\fR` stream. The receiver on
 151 the other side of the stream is notified of the urgent data by a `\fBSIGURG\fR`
 152 signal. The `\fBSIOCATMARK\fR \fBiocctl\fR(2) request returns a value indicating
 153 whether the stream is at the urgent mark. Because the system never returns data
 154 across the urgent mark in a single \fBread\fR(2) call, it is possible to
 155 advance to the urgent data in a simple loop which reads data, testing the
 156 socket with the \fBSIOCATMARK\fR \fBiocctl() \fR request, until it reaches the
 157 mark.
 158 .sp
 159 .LP
 160 Incoming connection requests that include an \fBIP\fR source route option are
 161 noted, and the reverse source route is used in responding.
 162 .sp
 163 .LP
 164 A checksum over all data helps \fBTCP\fR implement reliability. Using a
 165 window-based flow control mechanism that makes use of positive
 166 acknowledgements, sequence numbers, and a retransmission strategy, \fBTCP\fR
 167 can usually recover when datagrams are damaged, delayed, duplicated or
 168 delivered out of order by the underlying communication medium.
 169 .sp
 170 .LP
 171 If the local \fBTCP\fR receives no acknowledgements from its peer for a period
 172 of time, (for example, if the remote machine crashes), the connection is closed
 173 and an error is returned.
 174 .sp
 175 .LP
 176 TCP follows the congestion control algorithm described in \fIRFC 2581\fR, and
 177 also supports the initial congestion window (cwnd) changes in \fIRFC 3390\fR.
 178 The initial cwnd calculation can be overridden by the socket option
 179 TCP_INIT_CWND. An application can use this option to set the initial cwnd to a
 180 specified number of TCP segments. This applies to the cases when the connection
 181 first starts and restarts after an idle period. The process must have the
 182 PRIV_SYS_NET_CONFIG privilege if it wants to specify a number greater than that
 183 calculated by \fIRFC 3390\fR.
 184 .sp
 185 .LP
 186 SunOS supports \fBTCP\fR Extensions for High Performance (\fIRFC 1323\fR) which
 187 includes the window scale and time stamp options, and Protection Against Wrap
 188 Around Sequence Numbers (PAWS). SunOS also supports Selective Acknowledgment
 189 (SACK) capabilities (RFC 2018) and Explicit Congestion Notification (ECN)
 190 mechanism (\fIRFC 3168\fR).`

191 .sp
 192 .LP
 193 Turn on the window scale option in one of the following ways:
 194 .RS +4
 195 .TP
 196 .ie t \(\bu
 197 .el o
 198 An application can set `\fBSO_SNDBUF\fR` or `\fBSO_RCVBUF\fR` size in the
 199 `\fBsetsockopt()` `\fR` option to be larger than 64K. This must be done `\fIbefore\fR`
 200 the program calls `\fBlisten()` `\fR` or `\fBconnect()` `\fR`, because the window scale
 201 option is negotiated when the connection is established. Once the connection
 202 has been made, it is too late to increase the send or receive window beyond the
 203 default `\fBTCP\fR` limit of 64K.
 204 .RE
 205 .RS +4
 206 .TP
 207 .ie t \(\bu
 208 .el o
 209 For all applications, use `\fBndd\fR(1M)` to modify the configuration parameter
 210 `\fBtcp_wscale_always\fR`. If `\fBtcp_wscale_always\fR` is set to `\fB1\fR`, the
 211 window scale option will always be set when connecting to a remote system. If
 212 `\fBtcp_wscale_always\fR` is `\fB0,\fR` the window scale option will be set only if
 213 the user has requested a send or receive window larger than 64K. The default
 214 value of `\fBtcp_wscale_always\fR` is `\fB1\fR`.
 215 .RE
 216 .RS +4
 217 .TP
 218 .ie t \(\bu
 219 .el o
 220 Regardless of the value of `\fBtcp_wscale_always\fR`, the window scale option
 221 will always be included in a connect acknowledgement if the connecting system
 222 has used the option.
 223 .RE
 224 .sp
 225 .LP
 226 Turn on `\fBSACK\fR` capabilities in the following way:
 227 .RS +4
 228 .TP
 229 .ie t \(\bu
 230 .el o
 231 Use `\fBndd\fR` to modify the configuration parameter `\fBtcp_sack_permitted\fR`.
 232 If `\fBtcp_sack_permitted\fR` is set to `\fB0\fR`, `\fBTCP\fR` will not accept
 233 `\fBSACK\fR` or send out `\fBSACK\fR` information. If `\fBtcp_sack_permitted\fR` is
 234 set to `\fB1\fR`, `\fBTCP\fR` will not initiate a connection with `\fBSACK\fR`
 235 permitted option in the `\fBSYN\fR` segment, but will respond with `\fBSACK\fR`
 236 permitted option in the `\fBSYN|ACK\fR` segment if an incoming connection request
 237 has the `\fBSACK \fR` permitted option. This means that `\fBTCP\fR` will only
 238 accept `\fBSACK\fR` information if the other side of the connection also accepts
 239 `\fBSACK\fR` information. If `\fBtcp_sack_permitted\fR` is set to `\fB2\fR`, it will
 240 both initiate and accept connections with `\fBSACK\fR` information. The default
 241 for `\fBtcp_sack_permitted\fR` is `\fB2\fR` (active enabled).
 242 .RE
 243 .sp
 244 .LP
 245 Turn on `\fBTCP ECN\fR` mechanism in the following way:
 246 .RS +4
 247 .TP
 248 .ie t \(\bu
 249 .el o
 250 Use `\fBndd\fR` to modify the configuration parameter `\fBtcp_ecn_permitted\fR`. If
 251 `\fBtcp_ecn_permitted\fR` is set to `\fB0\fR`, `\fBTCP\fR` will not negotiate with a
 252 peer that supports `\fBECN\fR` mechanism. If `\fBtcp_ecn_permitted\fR` is set to
 253 `\fB1\fR` when initiating a connection, TCP will not tell a peer that it supports
 254 ECN mechanism. However, it will tell a peer that it supports `\fBECN\fR`
 255 mechanism when accepting a new incoming connection request if the peer
 256 indicates that it supports `\fBECN\fR` mechanism in the SYN segment. If

257 tcp_ecn_permitted is set to 2, in addition to negotiating with a peer on ECN
 258 mechanism when accepting connections, TCP will indicate in the outgoing SYN
 259 segment that it supports \fBECN\fR mechanism when \fBTCP\fR makes active
 260 outgoing connections. The default for \fBtcp_ecn_permitted\fR is 1.

261 .RE
 262 .sp
 263 .LP
 264 Turn on the time stamp option in the following way:
 265 .RS +4
 266 .TP
 267 .ie t \ (bu
 268 .el o
 269 Use \fBndd\fR to modify the configuration parameter \fBtcp_tstamp_always\fR. If
 270 \fBtcp_tstamp_always\fR is \fB1\fR, the time stamp option will always be set
 271 when connecting to a remote machine. If \fBtcp_tstamp_always\fR is \fB0\fR, the
 272 timestamp option will not be set when connecting to a remote system. The
 273 default for \fBtcp_tstamp_always\fR is \fB0\fR.

274 .RE
 275 .RS +4
 276 .TP
 277 .ie t \ (bu
 278 .el o
 279 Regardless of the value of \fBtcp_tstamp_always\fR, the time stamp option will
 280 always be included in a connect acknowledgement (and all succeeding packets) if
 281 the connecting system has used the time stamp option.

282 .RE
 283 .sp
 284 .LP
 285 Use the following procedure to turn on the time stamp option only when the
 286 window scale option is in effect:
 287 .RS +4
 288 .TP
 289 .ie t \ (bu
 290 .el o
 291 Use \fBndd\fR to modify the configuration parameter \fBtcp_tstamp_if_wscale\fR.
 292 Setting \fBtcp_tstamp_if_wscale\fR to \fB1\fR will cause the time stamp option
 293 to be set when connecting to a remote system, if the window scale option has
 294 been set. If \fBtcp_tstamp_if_wscale\fR is \fB0\fR, the time stamp option will
 295 not be set when connecting to a remote system. The default for
 296 \fBtcp_tstamp_if_wscale\fR is \fB1\fR.

297 .RE
 298 .sp
 299 .LP
 300 Protection Against Wrap Around Sequence Numbers (PAWS) is always used when the
 301 time stamp option is set.

302 .sp
 303 .LP
 304 SunOS also supports multiple methods of generating initial sequence numbers.
 305 One of these methods is the improved technique suggested in \fBRFC 1948\fR. We
 306 \fBHIGLY\fR recommend that you set sequence number generation parameters as
 307 close to boot time as possible. This prevents sequence number problems on
 308 connections that use the same connection-ID as ones that used a different
 309 sequence number generation. The \fBsvc:/network/initial:default\fR service
 310 configures the initial sequence number generation. The service reads the value
 311 contained in the configuration file \fB/etc/default/inetinit\fR to determine
 312 which method to use.

313 .sp
 314 .LP
 315 The \fB/etc/default/inetinit\fR file is an unstable interface, and may change
 316 in future releases.

317 .sp
 318 .LP
 319 \fBTCP\fR may be configured to report some information on connections that
 320 terminate by means of an \fBRST\fR packet. By default, no logging is done. If
 321 the \fBndd\fR(1M) parameter \fBtcp_trace\fR is set to 1, then trace data is
 322 collected for all new connections established after that time.

323 .sp
 324 .LP
 325 The trace data consists of the \fBTCP\fR headers and \fBIP\fR source and
 326 destination addresses of the last few packets sent in each direction before RST
 327 occurred. Those packets are logged in a series of \fBstrlog\fR(9F) calls. This
 328 trace facility has a very low overhead, and so is superior to such utilities as
 329 \fBsnoop\fR(1M) for non-intrusive debugging for connections terminating by
 330 means of an \fBRST\fR.

331 .sp
 332 .LP
 333 SunOS supports the keep-alive mechanism described in \fBIRFC 1122\fR. It is
 334 enabled using the socket option SO_KEEPALIVE. When enabled, the first
 335 keep-alive probe is sent out after a TCP is idle for two hours. If the peer does
 336 not respond to the probe within eight minutes, the TCP connection is aborted.
 337 You can alter the interval for sending out the first probe using the socket
 338 option TCP_KEEPALIVE_THRESHOLD. The option value is an unsigned integer in
 339 milliseconds. The system default is controlled by the TCP ndd parameter
 340 tcp_keepalive_interval. The minimum value is ten seconds. The maximum is ten
 341 days, while the default is two hours. If you receive no response to the probe,
 342 you can use the TCP_KEEPALIVE_ABORT_THRESHOLD socket option to change the time
 343 threshold for aborting a TCP connection. The option value is an unsigned
 344 integer in milliseconds. The value zero indicates that TCP should never time
 345 out and abort the connection when probing. The system default is controlled by
 346 the TCP ndd parameter tcp_keepalive_abort_interval. The default is eight
 347 minutes.

348 .sp
 349 .LP
 350 socket options TCP_KEEPIIDLE, TCP_KEEPCNT and TCP_KEEPINTVL are also supported
 351 for compatibility with other Unix Flavors. TCP_KEEPIIDLE option specifies the
 352 interval in seconds for sending out the first keep-alive probe. TCP_KEEPCNT
 353 specifies the number of keep-alive probes to be sent before aborting the
 354 connection in the event of no response from peer. TCP_KEEPINTVL specifies the
 355 interval in seconds between successive keep-alive probes.

356 .SH SEE ALSO
 357 .sp
 358 .LP
 359 \fBsvcs\fR(1), \fBndd\fR(1M), \fBioctl\fR(2), \fBread\fR(2), \fBsvcadm\fR(1M),
 360 \fBwrite\fR(2), \fBaccept\fR(3SOCKET), \fBbind\fR(3SOCKET),
 361 \fBconnect\fR(3SOCKET), \fBgetprotobyname\fR(3SOCKET),
 362 \fBgetsockopt\fR(3SOCKET), \fBlisten\fR(3SOCKET), \fBsend\fR(3SOCKET),
 363 \fBsmf\fR(5), \fBinet\fR(7P), \fBinet6\fR(7P), \fBip\fR(7P), \fBip6\fR(7P)

364 .sp
 365 .LP
 366 Ramakrishnan, K., Floyd, S., Black, D., RFC 3168, \fBThe Addition of Explicit
 367 Congestion Notification (ECN) to IP\fR, September 2001.

368 .sp
 369 .LP
 370 Mathias, M. and Haahdavi, J. Pittsburgh Supercomputing Center; Ford, S. Lawrence
 371 Berkeley National Laboratory; Romanow, A. Sun Microsystems, Inc. \fBIRFC 2018,
 372 TCP Selective Acknowledgement Options\fR, October 1996.

373 .sp
 374 .LP
 375 Bellovin, S., \fBIRFC 1948, Defending Against Sequence Number Attacks\fR, May
 376 1996.

377 .sp
 378 .LP
 379 Jacobson, V., Braden, R., and Borman, D., \fBIRFC 1323, TCP Extensions for High
 380 Performance\fR, May 1992.

381 .sp
 382 .LP
 383 Postel, Jon, \fBIRFC 793, Transmission Control Protocol - DARPA Internet Program
 384 Protocol Specification\fR, Network Information Center, SRI International, Menlo
 385 Park, CA., September 1981.

386 .SH DIAGNOSTICS
 387 .sp
 388 .LP

```
389 A socket operation may fail if:
390 .sp
391 .ne 2
392 .na
393 \fB\fBEISCONN\fR\fR
394 .ad
395 .RS 17n
396 A \fBconnect()\fR operation was attempted on a socket on which a
397 \fBconnect()\fR operation had already been performed.
398 .RE

400 .sp
401 .ne 2
402 .na
403 \fB\fBETIMEDOUT\fR\fR
404 .ad
405 .RS 17n
406 A connection was dropped due to excessive retransmissions.
407 .RE

409 .sp
410 .ne 2
411 .na
412 \fB\fBECNNRESET\fR\fR
413 .ad
414 .RS 17n
415 The remote peer forced the connection to be closed (usually because the remote
416 machine has lost state information about the connection due to a crash).
417 .RE

419 .sp
420 .ne 2
421 .na
422 \fB\fBECNNREFUSED\fR\fR
423 .ad
424 .RS 17n
425 The remote peer actively refused connection establishment (usually because no
426 process is listening to the port).
427 .RE

429 .sp
430 .ne 2
431 .na
432 \fB\fBEADDRINUSE\fR\fR
433 .ad
434 .RS 17n
435 A \fBbind()\fR operation was attempted on a socket with a network address/port
436 pair that has already been bound to another socket.
437 .RE

439 .sp
440 .ne 2
441 .na
442 \fB\fBEADDRNOTAVAIL\fR\fR
443 .ad
444 .RS 17n
445 A \fBbind()\fR operation was attempted on a socket with a network address for
446 which no network interface exists.
447 .RE

449 .sp
450 .ne 2
451 .na
452 \fB\fBEACCES\fR\fR
453 .ad
454 .RS 17n
```

```
455 A \fBbind()\fR operation was attempted with a "reserved" port number and the
456 effective user \fBID\fR of the process was not the privileged user.
457 .RE

459 .sp
460 .ne 2
461 .na
462 \fB\fBENOBUFS\fR\fR
463 .ad
464 .RS 17n
465 The system ran out of memory for internal data structures.
466 .RE

468 .SH NOTES
469 .sp
470 .LP
471 The \fBtcp\fR service is managed by the service management facility,
472 \fBsmf(5)\fR, under the service identifier:
473 .sp
474 .in +2
475 .nf
476 svc:/network/initial:default
477 .fi
478 .in -2
479 .sp

481 .sp
482 .LP
483 Administrative actions on this service, such as enabling, disabling, or
484 requesting restart, can be performed using \fBsvcadm\fR(1M). The service's
485 status can be queried using the \fBsvcs\fR(1) command.
```