

new/usr/src/man/Makefile

1

\*\*\*\*\*

1815 Thu Jul 17 00:50:38 2014

new/usr/src/man/Makefile

manpage lint.

\*\*\*\*\*

```
1 #
2 # This file and its contents are supplied under the terms of the
3 # Common Development and Distribution License ("CDDL"), version 1.0.
4 # You may only use this file in accordance with the terms of version
5 # 1.0 of the CDDL.
6 #
7 # A full copy of the text of the CDDL should have accompanied this
8 # source. A copy of the CDDL is also available via the Internet
9 # at http://www.illumos.org/license/CDDL.
10 #
```

```
12 #
13 # Copyright 2011, Richard Lowe
14 # Copyright 2013 Nexenta Systems, Inc. All rights reserved.
15 #
```

```
17 SUBDIRS=      man1      \
18               man1b     \
19               man1c     \
20               man1has   \
21               man1m     \
22               man2      \
23               man3      \
24               man3bsm   \
25               man3c     \
26               man3c_db  \
27               man3cfgadm \
28               man3computil \
29               man3contract \
30               man3cpc   \
31               man3curses \
32               man3dat   \
33               man3deviid \
34               man3devinfo \
35               man3dlpi  \
36               man3dns_sd \
37               man3elf   \
38               man3exacct \
39               man3ext   \
40               man3fcoe  \
41               man3fstyp \
42               man3gen   \
43               man3gss   \
44               man3head  \
45               man3iscsit \
46               man3kstat \
47               man3kvm   \
48               man3ldap  \
49               man3lgrp  \
50               man3lib   \
51               man3mail  \
52               man3malloc \
53               man3mp    \
54               man3mpapi \
55               man3nsl   \
56               man3nvpair \
57               man3pam   \
58               man3papi  \
59               man3perl  \
60               man3picl  \
61               man3picltree \
```

new/usr/src/man/Makefile

2

```
62               man3pool  \
63               man3proc  \
64               man3project \
65               man3resolv \
66               man3rpc   \
67               man3rsm   \
68               man3sas1  \
69               man3scf   \
70               man3sec   \
71               man3secdb \
72               man3sip   \
73               man3slp   \
74               man3socket \
75               man3stmf  \
76               man3sysevent \
77               man3tecla \
78               man3tnf   \
79               man3tsol  \
80               man3uuid  \
81               man3volmgt \
82               man3xcurses \
83               man3xnet  \
84               man4      \
85               man5      \
86               man7      \
87               man7d     \
88               man7fs    \
89               man7i     \
90               man7ipp   \
91               man7m     \
92               man7p     \
93               man9      \
94               man9e     \
95               man9f     \
96               man9p     \
97               man9s     \
```

99 .PARALLEL: \$(SUBDIRS)

```
101 all           := TARGET = all
102 clean         := TARGET = clean
103 clobber       := TARGET = clobber
104 install       := TARGET = install
105 check        := TARGET = check
```

```
107 all check clean clobber install: $(SUBDIRS)
106 all clean clobber install: $(SUBDIRS)
```

```
109 $(SUBDIRS):      FRC
110                 @cd $@; pwd; $(MAKE) $(TARGET)
```

112 FRC:

new/usr/src/man/Makefile.man

1

\*\*\*\*\*

1239 Thu Jul 17 00:50:38 2014

new/usr/src/man/Makefile.man

manpage lint.

\*\*\*\*\*

```
1 #
2 # This file and its contents are supplied under the terms of the
3 # Common Development and Distribution License ("CDDL"), version 1.0.
4 # You may only use this file in accordance with the terms of version
5 # 1.0 of the CDDL.
6 #
7 # A full copy of the text of the CDDL should have accompanied this
8 # source. A copy of the CDDL is also available via the Internet
9 # at http://www.illumos.org/license/CDDL.
10 #
```

```
12 #
13 # Copyright 2011, Richard Lowe
14 # Copyright 2013 Nexenta Systems, Inc. All rights reserved.
15 #
```

```
17 MANDOC= $(ROOT)/usr/bin/mandoc
18 ROOTMAN= $(ROOT)/usr/share/man
19 ROOTHASMAN= $(ROOT)/usr/has/man
```

```
20 FILEMODE= 0444
```

```
22 # The manual section being built, client Makefiles must set this to, for e.g.
23 # "3perl", with case matching that of the section name as installed.
24 #
25 # MANSECT=
```

```
27 MANCHECKS= $(MANFILES:%=%.check)
28 ROOTMANFILES= $(MANFILES:%=$(ROOTMAN)/man$(MANSECT)/%)
29 ROOTMANLINKS= $(MANLINKS:%=$(ROOTMAN)/man$(MANSECT)/%)
```

```
31 $(ROOTMAN)/man$(MANSECT)/% $(ROOTHASMAN)/man$(MANSECT)/%: %
32 $(INS.file)
```

```
34 $(MANCHECKS):
35 $(MANDOC) -Tlint $(@:%.check=%)
```

```
37 $(MANLINKS):
38 $(RM) $@; $(SYMLINK) $(LINKSRC) $@
```

```
40 $(ROOTMANLINKS): $(MANLINKS)
41 $(RM) $@; $(CP) -RP $(@F) $(@D)
```

```
43 all:
```

```
45 check: $(MANCHECKS)
```

```
47 clean:
```

```
49 clobber:
50 $(RM) $(MANLINKS)
```

```
52 .PARALLEL:
```

```
54 FRC:
```

new/usr/src/man/man1/head.1

1

```
*****  
      8110 Thu Jul 17 00:50:38 2014  
new/usr/src/man/man1/head.1  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

new/usr/src/man/man1/ld.1

1

```
*****  
59044 Thu Jul 17 00:50:38 2014  
new/usr/src/man/man1/ld.1  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

new/usr/src/man/man1/man.1

1

```
*****  
3098 Thu Jul 17 00:50:39 2014  
new/usr/src/man/man1/man.1  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

new/usr/src/man/man1/mandoc.1

1

```
*****  
14742 Thu Jul 17 00:50:39 2014  
new/usr/src/man/man1/mandoc.1  
manpage lint.  
*****  
_____unchanged_portion_omitted_
```

new/usr/src/man/man1/printf.1

1

```
*****  
22738 Thu Jul 17 00:50:39 2014  
new/usr/src/man/man1/printf.1  
manpage lint.  
*****  
_____unchanged_portion_omitted_
```

new/usr/src/man/man1/tar.1

1

```
*****  
34827 Thu Jul 17 00:50:39 2014  
new/usr/src/man/man1/tar.1  
manpage lint.  
*****  
_____unchanged_portion_omitted_
```



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

1

```
*****  
4623 Thu Jul 17 00:50:39 2014  
new/usr/src/man/man1m/arcstat.1m  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

\*\*\*\*\*

17974 Thu Jul 17 00:50:39 2014

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

manpage lint.

\*\*\*\*\*

```

1  \" te
2  .\" Copyright (c) 2014, Joyent, Inc. All rights Reserved.
3  .\" Copyright (c) 1992, X/Open Company Limited All Rights Reserved
4  .\" Copyright 1989 AT&T
5  .\" Portions Copyright (c) 1995, Sun Microsystems, Inc. All Rights Reserved
6  .\" Sun Microsystems, Inc. gratefully acknowledges The Open Group for permission
7  .\" http://www.opengroup.org/bookstore/.
8  .\" The Institute of Electrical and Electronics Engineers and The Open Group, ha
9  .\" This notice shall appear on any product containing this material.
10 .\" The contents of this file are subject to the terms of the Common Development
11 .\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
12 .\" When distributing Covered Code, include this CDDL HEADER in each file and in
13 .TH DD 1M "Jan 04, 2014"
14 .SH NAME
15 dd \- convert and copy a file
16 .SH SYNOPSIS
17 .LP
18 .nf
19 \fB/usr/bin/dd\fR [\fIoperand=value\fR]...
20 .fi

22 .SH DESCRIPTION
23 .sp
24 .LP
25 The \fBdd\fR utility copies the specified input file to the specified output
26 with possible conversions. The standard input and output are used by default.
27 The input and output block sizes may be specified to take advantage of raw
28 physical I/O. Sizes are specified in bytes; a number may end with \fBk\fR,
29 \fBb\fR, or \fBw\fR to specify multiplication by 1024, 512, or 2, respectively.
30 Numbers may also be separated by \fBx\fR to indicate multiplication.
31 .sp
32 .LP
33 The \fBdd\fR utility reads the input one block at a time, using the specified
34 input block size. \fBdd\fR then processes the block of data actually returned,
35 which could be smaller than the requested block size. \fBdd\fR applies any
36 conversions that have been specified and writes the resulting data to the
37 output in blocks of the specified output block size.
38 .sp
39 .LP
40 \fBcbs\fR is used only if \fBascii\fR, \fBasciib\fR, \fBbunblock\fR,
41 \fBebcdic\fR, \fBebcdicb\fR, \fBibm\fR, \fBibmb\fR, or \fBblock\fR conversion
42 is specified. In the first two cases, \fBcbs\fR characters are copied into the
43 conversion buffer, any specified character mapping is done, trailing blanks are
44 trimmed, and a \fBNEWLINE\fR is added before sending the line to output. In the
45 last three cases, characters up to \fBNEWLINE\fR are read into the conversion
46 buffer and blanks are added to make up an output record of size \fBcbs\fR.
47 \fBASCII\fR files are presumed to contain \fBNEWLINE\fR characters. If
48 \fBcbs\fR is unspecified or \fB0\fR, the \fBascii\fR, \fBasciib\fR,
49 \fBebcdic\fR, \fBebcdicb\fR, \fBibm\fR, and \fBibmb\fR options convert the
50 character set without changing the input file's block structure. The
51 \fBbunblock\fR and \fBblock\fR options become a simple file copy.
52 .sp
53 .LP
54 After completion, \fBdd\fR reports the number of whole and partial input and
55 output blocks.
56 .SH OPERANDS
57 .sp
58 .LP
59 The following operands are supported:
60 .sp
61 .ne 2

```

```

62 .na
63 \fB\bif=\fR\fIfile\fR\fR
64 .ad
65 .sp .6
66 .RS 4n
67 Specifies the input path. Standard input is the default.
68 .RE

70 .sp
71 .ne 2
72 .na
73 \fB\bof=\fR\fIfile\fR\fR
74 .ad
75 .sp .6
76 .RS 4n
77 Specifies the output path. Standard output is the default. If the
78 \fBseek=\fR\fIexpr\fR conversion is not also specified, the output file will be
79 truncated before the copy begins, unless \fBconv=notrunc\fR is specified. If
80 \fBseek=\fR\fIexpr\fR is specified, but \fBconv=notrunc\fR is not, the effect
81 of the copy will be to preserve the blocks in the output file over which
82 \fBdd\fR seeks, but no other portion of the output file will be preserved. (If
83 the size of the seek plus the size of the input file is less than the previous
84 size of the output file, the output file is shortened by the copy.)
85 .RE

87 .sp
88 .ne 2
89 .na
90 \fB\bibs=\fR\fIn\fR\fR
91 .ad
92 .sp .6
93 .RS 4n
94 Specifies the input block size in \fIn\fR bytes (default is \fB512\fR).
95 .RE

97 .sp
98 .ne 2
99 .na
100 \fB\bobs=\fR\fIn\fR\fR
101 .ad
102 .sp .6
103 .RS 4n
104 Specifies the output block size in \fIn\fR bytes (default is \fB512\fR).
105 .RE

107 .sp
108 .ne 2
109 .na
110 \fB\bbs=\fR\fIn\fR\fR
111 .ad
112 .sp .6
113 .RS 4n
114 Sets both input and output block sizes to \fIn\fR bytes, superseding \fBibs=\fR
115 and \fBobs=\fR. If no conversion other than \fBsync\fR, \fBnoerror\fR, and
116 \fBnotrunc\fR is specified, each input block is copied to the output as a
117 single block without aggregating short blocks.
118 .RE

120 .sp
121 .ne 2
122 .na
123 \fB\bcb=\fR\fIn\fR\fR
124 .ad
125 .sp .6
126 .RS 4n
127 Specifies the conversion block size for \fBblock\fR and \fBbunblock\fR in bytes

```

```

128 by \fIn\fR (default is \fB0\fR). If \fBcbs=\fR is omitted or given a value of
129 \fB0\fR, using \fBblock\fR or \fBunblock\fR produces unspecified results.
130 .sp
131 This option is used only if \fBASCII\fR or \fBBEBCDIC\fR conversion is
132 specified. For the \fBascii\fR and \fBasciiib\fR operands, the input is handled
133 as described for the \fBunblock\fR operand except that characters are converted
134 to \fBASCII\fR before the trailing \fBSPACE\fR characters are deleted. For the
135 \fBbebcidic\fR, \fBbebcidib\fR, \fBbim\b\fR, and \fBbimb\b\fR operands, the input is
136 handled as described for the \fBblock\fR operand except that the characters are
137 converted to \fBBEBCDIC\fR or IBM \fBBEBCDIC\fR after the trailing \fBSPACE\fR
138 characters are added.
139 .RE

141 .sp
142 .ne 2
143 .na
144 \fB\fBfiles=\fR\fIn\fR\fR
145 .ad
146 .sp .6
147 .RS 4n
148 Copies and concatenates \fIn\fR input files before terminating (makes sense
149 only where input is a magnetic tape or similar device).
150 .RE

152 .sp
153 .ne 2
154 .na
155 \fB\fBskip=\fR\fIn\fR\fR
156 .ad
157 .sp .6
158 .RS 4n
159 Skips \fIn\fR input blocks (using the specified input block size) before
160 starting to copy. On seekable files, the implementation reads the blocks or
161 seeks past them. On non-seekable files, the blocks are read and the data is
162 discarded.
163 .RE

165 .sp
166 .ne 2
167 .na
168 \fB\fBiseek=\fR\fIn\fR\fR
169 .ad
170 .sp .6
171 .RS 4n
172 Seeks \fIn\fR blocks from beginning of input file before copying (appropriate
173 for disk files, where \fBskip\fR can be incredibly slow).
174 .RE

176 .sp
177 .ne 2
178 .na
179 \fB\fBoseek=\fR\fIn\fR\fR
180 .ad
181 .sp .6
182 .RS 4n
183 Seeks \fIn\fR blocks from beginning of output file before copying.
184 .RE

186 .sp
187 .ne 2
188 .na
189 \fB\fBseek=\fR\fIn\fR\fR
190 .ad
191 .sp .6
192 .RS 4n
193 Skips \fIn\fR blocks (using the specified output block size) from beginning of

```

```

194 output file before copying. On non-seekable files, existing blocks are read and
195 space from the current end-of-file to the specified offset, if any, is filled
196 with null bytes. On seekable files, the implementation seeks to the specified
197 offset or reads the blocks as described for non-seekable files.
198 .RE

200 .sp
201 .ne 2
202 .na
203 \fB\fBcount=\fR\fIn\fR\fR
204 .ad
205 .sp .6
206 .RS 4n
207 Copies only \fIn\fR input blocks.
208 .RE

210 .sp
211 .ne 2
212 .na
213 \fB\fBconv=\fR\fIvalue\fR[\fB,\fR\fIvalue\fR.\|\|\|\fR
214 .ad
215 .sp .6
216 .RS 4n
217 Where \fIvalue\fRs are comma-separated symbols from the following list:
218 .sp
219 .ne 2
220 .na
221 \fB\fBascii\fR\fR
222 .ad
223 .RS 11n
224 Converts \fBBEBCDIC\fR to \fBASCII\fR.
225 .RE

227 .sp
228 .ne 2
229 .na
230 \fB\fBasciiib\fR\fR
231 .ad
232 .RS 11n
233 Converts \fBBEBCDIC\fR to \fBASCII\fR using \fBBSD\fR-compatible character
234 translations.
235 .RE

237 .sp
238 .ne 2
239 .na
240 \fB\fBbebcidic\fR\fR
241 .ad
242 .RS 11n
243 Converts \fBASCII\fR to \fBBEBCDIC\fR. If converting fixed-length \fBASCII\fR
244 records without NEWLINES, sets up a pipeline with \fBdd conv=unblock\fR
245 beforehand.
246 .RE

248 .sp
249 .ne 2
250 .na
251 \fB\fBbebcidib\fR\fR
252 .ad
253 .RS 11n
254 Converts \fBASCII\fR to \fBBEBCDIC\fR using \fBBSD\fR-compatible character
255 translations. If converting fixed-length \fBASCII\fR records without
256 \fBNEWLINE\fRs, sets up a pipeline with \fBdd conv=unblock\fR beforehand.
257 .RE

259 .sp

```



```

392 .sp
393 .ne 2
394 .na
395 \fB\fBdsync\fR\fR
396 .ad
397 .RS 11n
398 The output file is opened with the \fBO_DSYNC\fR flag set. All data writes will
399 be synchronous. For more information on \fBO_DSYNC\fR see \fBfcntl.h\fR(3HEAD).
400 .RE

402 .sp
403 .ne 2
404 .na
405 \fB\fBsync\fR\fR
406 .ad
407 .RS 11n
408 The output file is opened with the \fBO_SYNC\fR flag set. All data and metadata
409 writes will be synchronous. For more information on \fBO_SYNC\fR see
410 \fBfcntl.h\fR(3HEAD).
411 .RE

413 .RE

413 .sp
414 .LP
415 If operands other than \fBconv=\fR and \fBoflag=\fR are specified more than once
416 the last specified \fBoperand=\fR\fIvalue\fR is used.
417 .sp
418 .LP
419 For the \fBbs=\fR, \fBcbs=\fR, \fBibs=\fR, and \fBobs=\fR operands, the
420 application must supply an expression specifying a size in bytes. The
421 expression, \fBexpr\fR, can be:
422 .RS +4
423 .TP
424 1.
425 a positive decimal number
426 .RE
427 .RS +4
428 .TP
429 2.
430 a positive decimal number followed by \fBk\fR, specifying multiplication by
431 1024
432 .RE
433 .RS +4
434 .TP
435 3.
436 a positive decimal number followed by \fBm\fR, specifying multiplication by
437 1024*1024
438 .RE
439 .RS +4
440 .TP
441 4.
442 a positive decimal number followed by \fBg\fR, specifying multiplication by
443 1024*1024*1024
444 .RE
445 .RS +4
446 .TP
447 5.
448 a positive decimal number followed by \fBt\fR, specifying multiplication by
449 1024*1024*1024*1024
450 .RE
451 .RS +4
452 .TP
453 6.
454 a positive decimal number followed by \fBp\fR, specifying multiplication by
455 1024*1024*1024*1024*1024

```

```

456 .RE
457 .RS +4
458 .TP
459 7.
460 a positive decimal number followed by \fBE\fR, specifying multiplication by
461 1024*1024*1024*1024*1024*1024
462 .RE
463 .RS +4
464 .TP
465 8.
466 a positive decimal number followed by \fBZ\fR, specifying multiplication by
467 1024*1024*1024*1024*1024*1024
468 .RE
469 .RS +4
470 .TP
471 9.
472 a positive decimal number followed by \fBb\fR, specifying multiplication by
473 512
474 .RE
475 .RS +4
476 .TP
477 10.
478 two or more positive decimal numbers (with or without \fBk\fR or \fBb\fR)
479 separated by \fBx\fR, specifying the product of the indicated values.
480 .RE
481 .sp
482 .LP
483 All of the operands will be processed before any input is read.
484 .SH SIGNALS
485 .sp
486 .LP
487 When \fBdd\fR receives either SIGINFO or SIGUSR1, \fBdd\fR will emit the current
488 input and output block counts, total bytes written, total time elapsed, and the
489 number of bytes per second to standard error. This is the same information
490 format that \fBdd\fR emits when it successfully completes. Users may send
491 SIGINFO via their terminal. The default character is ^T, see \fBstty\fR(1) for
492 more information.
493 .sp
494 .LP
495 For \fBSIGINT\fR, \fBdd\fR writes status information to standard error before
496 exiting. \fBdd\fR takes the standard action for all other signals.

498 .SH USAGE
499 .sp
500 .LP
501 See \fBlargefile\fR(5) for the description of the behavior of \fBdd\fR when
502 encountering files greater than or equal to 2 Gbyte ( 2^31 bytes).
503 .SH EXAMPLES
504 .LP
505 \fBExample 1 \fRCopying from one tape drive to another
506 .sp
507 .LP
508 The following example copies from tape drive \fB0\fR to tape drive \fB1\fR,
509 using a common historical device naming convention.

511 .sp
512 .in +2
513 .nf
514 example% \fBdd if=/dev/rmt/0h of=/dev/rmt/1h\fR
515 .fi
516 .in -2
517 .sp

519 .LP
520 \fBExample 2 \fRStripping the first 10 bytes from standard input
521 .sp

```

```

522 .LP
523 The following example strips the first 10 bytes from standard input:

525 .sp
526 .in +2
527 .nf
528 example% \fBdd ibs=10 skip=1\fR
529 .fi
530 .in -2
531 .sp

533 .LP
534 \fBExample 3 \fRReading a tape into an ASCII file
535 .sp
536 .LP
537 This example reads an \fBEBCDIC\fR tape blocked ten 80-byte \fBEBCDIC\fR card
538 images per block into the \fBASCII\fR file \fBx\fR:

540 .sp
541 .in +2
542 .nf
543 example% \fBdd if=/dev/tape of=x ibs=800 cbs=80 conv=ascii,lcase\fR
544 .fi
545 .in -2
546 .sp

548 .LP
549 \fBExample 4 \fRUsing conv=sync to write to tape
550 .sp
551 .LP
552 The following example uses \fBconv=sync\fR when writing to a tape:

554 .sp
555 .in +2
556 .nf
557 example% \fBtar cvf - . | compress | dd obs=1024k of=/dev/rmt/0 conv=sync\fR
558 .fi
559 .in -2
560 .sp

562 .SH ENVIRONMENT VARIABLES
563 .sp
564 .LP
565 See \fBenvron\fR(5) for descriptions of the following environment variables
566 that affect the execution of \fBdd\fR: \fBBLANG\fR, \fBBLC_ALL\fR,
567 \fBBLC_CTYPE\fR, \fBBLC_MESSAGES\fR, and \fBNLSPATH\fR.
568 .SH EXIT STATUS
569 .sp
570 .LP
571 The following exit values are returned:
572 .sp
573 .ne 2
574 .na
575 \fBfB0\fR
576 .ad
577 .RS 6n
578 The input file was copied successfully.
579 .RE

581 .sp
582 .ne 2
583 .na
584 \fBfB>0\fR
585 .ad
586 .RS 6n
587 An error occurred.

```

```

588 .RE

590 .sp
591 .LP
592 If an input error is detected and the \fBnoerror\fR conversion has not been
593 specified, any partial output block will be written to the output file, a
594 diagnostic message will be written, and the copy operation will be
595 discontinued. If some other error is detected, a diagnostic message will be
596 written and the copy operation will be discontinued.
597 .SH ATTRIBUTES
598 .sp
599 .LP
600 See \fBattributes\fR(5) for descriptions of the following attributes:
601 .sp

603 .sp
604 .TS
605 box;
606 c | c
607 l | l
608 ATTRIBUTE TYPE ATTRIBUTE VALUE
609 _
610 Interface Stability Standard
611 .TE

613 .SH SEE ALSO
614 .sp
615 .LP
616 \fBbcp\fR(1), \fBbsd\fR(1), \fBbtr\fR(1), \fBfcntl.h\fR(3HEAD),
617 \fBattributes\fR(5), \fBenvron\fR(5), \fBlargefile\fR(5), \fBstandards\fR(5)
618 .SH DIAGNOSTICS
619 .sp
620 .ne 2
621 .na
622 \fBfBf+p records in(out)\fR
623 .ad
624 .RS 23n
625 numbers of full and partial blocks read(written)
626 .RE

628 .SH NOTES
629 .sp
630 .LP
631 Do not use \fBdd\fR to copy files between file systems having different block
632 sizes.
633 .sp
634 .LP
635 Using a blocked device to copy a file will result in extra nulls being added
636 to the file to pad the final block to the block boundary.
637 .sp
638 .LP
639 When \fBdd\fR reads from a pipe, using the \fBbibs=X\fR and \fBbobs=Y\fR
640 operands, the output will always be blocked in chunks of size Y. When
641 \fBbbs=Z\fR is used, the output blocks will be whatever was available to be read
642 from the pipe at the time.
643 .sp
644 .LP
645 When using \fBdd\fR to copy files to a tape device, the file size must be a
646 multiple of the device sector size (for example, 512 Kbyte). To copy files of
647 arbitrary size to a tape device, use \fBtar\fR(1) or \fBcpio\fR(1).

```

\*\*\*\*\*

22252 Thu Jul 17 00:50:39 2014

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

manpage lint.

\*\*\*\*\*

```

1  \' te
2  .\ Copyright (c) 2012, Joyent, Inc. All Rights Reserved
3  .\ Copyright (c) 2013 by Delphix. All rights reserved.
4  .\ The contents of this file are subject to the terms of the Common Development
5  .\ You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
6  .\ When distributing Covered Code, include this CDDL HEADER in each file and in
7  .TH IPADM 1M "May 14, 2012"
8  .SH NAME
9  ipadm \- configure IP network interfaces and protocol properties.
10 .SH SYNOPSIS
11 .LP
12 .nf
13 \fBipadm\fR create-if [\fB-t\fR] \fIinterface\fR
14 .fi

16 .LP
17 .nf
18 \fBipadm\fR disable-if [\fB-t\fR] \fIinterface\fR
19 .fi

21 .LP
22 .nf
23 \fBipadm\fR enable-if [\fB-t\fR] \fIinterface\fR
24 .fi

26 .LP
27 .nf
28 \fBipadm\fR delete-if \fIinterface\fR
29 .fi

31 .LP
32 .nf
33 \fBipadm\fR show-if [[\fB-p\fR] \fB-o\fR \fIfield\fR[,...]] [\fIinterface\fR]
34 .fi

36 .LP
37 .nf
38 \fBipadm\fR set-ifprop [\fB-t\fR] \fB-p\fR \fIprop\fR=<\fIvalue\fR[,...]> \fB-m\
39 .fi

41 .LP
42 .nf
43 \fBipadm\fR reset-ifprop [\fB-t\fR] \fB-p\fR \fIprop\fR \fB-m\fR \fIprotocol\fR
44 .fi

46 .LP
47 .nf
48 \fBipadm\fR show-ifprop [[\fB-c\fR]\fB-o\fR \fIfield\fR[,...]] [\fB-p\fR \fIprop
49 [\fIinterface\fR]
50 .fi

52 .LP
53 .nf
54 \fBipadm\fR create-addr [\fB-t\fR] \fB-T\fR static [\fB-d\fR]
55 \fB-a\fR {\fB-l\fR|\fB-r\fR}=\fIaddr\fR[\fIprefixlen\fR],... \fIaddrobj\fR
56 .fi

58 .LP
59 .nf
60 \fBipadm\fR create-addr [\fB-t\fR] \fB-T\fR dhcp [\fB-w\fR \fIseconds\fR | forev
61 .fi

```

```

63 .LP
64 .nf
65 \fBipadm\fR create-addr [\fB-t\fR] \fB-T\fR addrconf [\fB-i\fR \fIinterface-id\f
66 [\fB-p\fR {\fB-s\fR|\fB-l\fR}=\fIstateful|stateless\fR={\fB-y\fR|\fB-n\fR},... ] \fIaddrobj\fR
67 .fi

69 .LP
70 .nf
71 \fBipadm\fR down-addr [\fB-t\fR] \fIaddrobj\fR
72 .fi

74 .LP
75 .nf
76 \fBipadm\fR up-addr [\fB-t\fR] \fIaddrobj\fR
77 .fi

79 .LP
80 .nf
81 \fBipadm\fR disable-addr [\fB-t\fR] \fIaddrobj\fR
82 .fi

84 .LP
85 .nf
86 \fBipadm\fR enable-addr [\fB-t\fR] \fIaddrobj\fR
87 .fi

89 .LP
90 .nf
91 \fBipadm\fR refresh-addr [\fB-i\fR] \fIaddrobj\fR
92 .fi

94 .LP
95 .nf
96 \fBipadm\fR delete-addr [\fB-r\fR] \fIaddrobj\fR
97 .fi

99 .LP
100 .nf
101 \fBipadm\fR show-addr [[\fB-p\fR] \fB-o\fR \fIfield\fR[,...]] [\fIaddrobj\fR]
102 .fi

104 .LP
105 .nf
106 \fBipadm\fR set-addrprop [\fB-t\fR] \fB-p\fR \fIprop\fR=<\fIvalue\fR[,...]> \fIa
107 .fi

109 .LP
110 .nf
111 \fBipadm\fR reset-addrprop [\fB-t\fR] \fB-p\fR \fIprop\fR=<\fIvalue\fR[,...]> \f
112 .fi

114 .LP
115 .nf
116 \fBipadm\fR show-addrprop [[\fB-c\fR] \fB-o\fR \fIfield\fR[,...]] [\fB-p\fR \fI
117 .fi

119 .LP
120 .nf
121 \fBipadm\fR set-prop [\fB-t\fR] \fB-p\fR \fIprop\fR[+|-]=<\fIvalue\fR[,...]> \fI
122 .fi

124 .LP
125 .nf
126 \fBipadm\fR reset-prop [\fB-t\fR] \fB-p\fR \fIprop\fR \fIprotocol\fR
127 .fi

```

```

129 .LP
130 .nf
131 \fBipadm\fR show-prop [[\fB-c\fR] \fB-o\fR \fIifield\fR[,...]] [\fB-p\fR \fIprop\
132 .fi

134 .SH DESCRIPTION
135 .sp
136 .LP

138 The \fBipadm\fR command is a stable replacement for the \fBifconfig\fR(1M) and
139 \fBndd\fR(1M) commands. It is used to create IP interfaces and to configure IP
140 addresses on those interfaces. It is also used to get, set or reset properties
141 on interfaces, addresses and protocols.
142 .LP
143 For subcommands that take an \fIaddrobj\fR, the \fIaddrobj\fR specifies a
144 unique address on the system. It is made up of two parts, delimited by a '/'.
145 The first part is the name of the interface and the second part is a string up
146 to 32 characters long. For example, "lo0/v4" is a loopback interface
147 addrobj name.
148 .LP
149 For subcommands that take a \fIprotocol\fR, this can be one of
150 the following values: ip, ipv4, ipv6, icmp, tcp, sctp or udp.

152 .SH SUBCOMMANDS
153 .sp
154 .LP
155 The following subcommands are supported:
156 .sp
157 .ne 2
158 .na
159 \fBifbcreate-if\fR [\fB-t\fR] \fIinterface\fR\fR
160 .ad
161 .sp .6
162 .RS 4n
163 The \fBifbcreate-if\fR subcommand is used to create an IP interface that will
164 handle both IPv4 and IPv6 packets. The interface will be enabled as part of
165 the creation process. The IPv4 interface will have the address 0.0.0.0.
166 The IPv6 interface will have the address ::.
167 .sp
168 The \fB-t\fR option (also \fB--temporary\fR) means
169 that the creation is temporary and will not be persistent across reboots.
170 .sp

172 .RE

174 .sp
175 .ne 2
176 .na
177 \fBifbdisable-if\fR [\fB-t\fR] \fIinterface\fR\fR
178 .ad
179 .sp .6
180 .RS 4n
181 The \fBifbdisable-if\fR subcommand is used to disable an IP interface.
182 .sp
183 The \fB-t\fR option (also \fB--temporary\fR) means
184 that the disable is temporary and will not be persistent across reboots.
185 .sp

187 .RE

189 .sp
190 .ne 2
191 .na
192 \fBifbenable-if\fR [\fB-t\fR] \fIinterface\fR\fR
193 .ad

```

```

194 .sp .6
195 .RS 4n
196 The \fBifbenable-if\fR subcommand is used to enable an IP interface.
197 .sp
198 The \fB-t\fR option (also \fB--temporary\fR) means
199 that the enable is temporary and will not be persistent across reboots.
200 .sp

202 .RE

204 .sp
205 .ne 2
206 .na
207 \fBifbdelete-if\fR \fIinterface\fR\fR
208 .ad
209 .sp .6
210 .RS 4n
211 The \fBifbdelete-if\fR subcommand is used to permanently delete an IP interface.
212 .sp

214 .RE

216 .sp
217 .ne 2
218 .na
219 \fBifbshow-if\fR [[\fB-p\fR] \fB-o\fR \fIifield\fR[,...]] [\fIinterface\fR]\fR
220 .ad
221 .sp .6
222 .RS 4n
223 The \fBifbshow-if\fR subcommand is used to show the current IP interface
224 configuration.
225 .sp
226 The \fB-p\fR option (also \fB--parsable\fR) prints
227 the output in a parsable format.
228 .sp
229 The \fB-o\fR option (also \fB--output\fR) is used
230 to select which fields will be shown. The field value can be one of the
231 following names:
232 .sp
233 .ne 2
234 .na
235 .RS 4n
236 \fBIFNAME\fR
237 .ad
238 .RS 4n
239 Display all fields
240 .RE

242 .sp
243 .ne 2
244 .na
245 \fBIFNAME\fR
246 .ad
247 .RS 4n
248 The name of the interface
249 .RE

251 .sp
252 .ne 2
253 .na
254 \fBIFSTATE\fR
255 .ad
256 .RS 4n
257 The state can be one of the following values:
258 .sp
259 .ne 2

```



```

260 .na
261 .RS 4n
262 ok - resources for the interface have been allocated
263 .sp
264 offline - the interface is offline
265 .sp
266 failed - the interface's datalink is down
267 .sp
268 down - the interface is down
269 .sp
270 disabled - the interface is disabled
271 .RE
272 .RE

274 .sp
275 .ne 2
276 .na
277 \fBCURRENT\fR
278 .ad
279 .RS 4n
280 A set of single character flags indicating the following:
281 .sp
282 .ne 2
283 .na
284 .RS 4n
285 b - broadcast (mutually exclusive with 'p')
286 .br
287 m - multicast
288 .br
289 p - point-to-point (mutually exclusive with 'b')
290 .br
291 v - virtual interface
292 .br
293 I - IPMP
294 .br
295 s - IPMP standby
296 .br
297 i - IPMP inactive
298 .br
299 V - VRRP
300 .br
301 a - VRRP accept mode
302 .br
303 4 - IPv4
304 .br
305 6 - IPv6
306 .RE
307 .RE

309 .sp
310 .ne 2
311 .na
312 \fBPERSISTENT\fR
313 .ad
314 .RS 4n
315 A set of single character flags showing what configuration will be used the
316 next time the interface is enabled:
317 .sp
318 .ne 2
319 .na
320 .RS 4n
321 s - IPMP standby
322 .br
323 4 - IPv4
324 .br
325 6 - IPv6

```

```

326 .RE
327 .RE
328 .RE

330 .RE

332 .sp
333 .ne 2
334 .na
335 \fB\fBset-ifprop\fR [\fB-t\fR] \fB-p\fR \fIprop\fR=<\fIvalue\fR[,...]> \fB-m\fR
336 .ad
337 .sp .6
338 .RS 4n
339 The \fBset-ifprop\fR subcommand is used to set a property's value(s) on the IP
340 interface.
341 .sp
342 The \fB-t\fR option (also \fB--temporary\fR) means
343 that the setting is temporary and will not be persistent across reboots.
344 .sp
345 The \fB-p\fR option (also \fB--prop\fR) specifies the property name and
346 value(s). The property name can be one of the following:
347 .sp
348 .ne 2
349 .na

351 .RS 4n

353 \fBarp\fR
354 .ad
355 .RS 4n
356 Enables ("on") or disables ("off") ARP.
357 .RE

359 .sp
360 .ne 2
361 .na
362 \fBexchange_routes\fR
363 .ad
364 .RS 4n
365 Enables ("on") or disables ("off") the exchange of routing data.
366 .RE

368 .sp
369 .ne 2
370 .na
371 \fBforwarding\fR
372 .ad
373 .RS 4n
374 Enables ("on") or disables ("off") IP forwarding.
375 .RE

377 .sp
378 .ne 2
379 .na
380 \fBmetric\fR
381 .ad
382 .RS 4n
383 Set the routing metric to the numeric value. The value is treated as extra
384 hops to the destination.
385 .RE

387 .sp
388 .ne 2
389 .na
390 \fBmtu\fR
391 .ad

```

```

392 .RS 4n
393 Set the maximum transmission unit to the numeric value.
394 .RE

396 .sp
397 .ne 2
398 .na
399 \fBnud\fR
400 .ad
401 .RS 4n
402 Enables ("on") or disables ("off") neighbor unreachability detection.
403 .RE

405 .sp
406 .ne 2
407 .na
408 \fBusesrc\fR
409 .ad
410 .RS 4n
411 Indicates which interface to use for source address selection. A value
412 "none" may also be used.
413 .RE
414 .RE

416 .sp
417 The \fB-m\fR option (also \fB--module\fR) specifies which protocol
418 the setting applies to.
419 .sp

421 .RE
422 .RE

423 .sp
424 .ne 2
425 .na
426 \fB\breset-ifprop\fR [\fB-t\fR] \fB-p\fR \fIprop\fR \fB-m\fR \fIprotocol\fR \fI
427 .ad
428 .sp .6
429 .RS 4n
430 The \fBreset-ifprop\fR subcommand is used to reset an IP interface's property
431 value to the default.
432 .sp
433 The \fB-t\fR option (also \fB--temporary\fR) means
434 that the disable is temporary and will not be persistent across reboots.
435 .sp
436 The \fB-p\fR option (also \fB--prop\fR) specifies the property name.
437 See the \fBset-ifprop\fR subcommand for the list of property names.
438 .sp
439 The \fB-m\fR option (also \fB--module\fR) specifies which protocol
440 the setting applies to.
441 .sp

443 .RE

445 .sp
446 .ne 2
447 .na
448 \fB\bshow-ifprop\fR [[\fB-c\fR]\fB-o\fR \fIfield\fR[,...]] [\fB-p\fR \fIprop\fR
449 [\fIinterface\fR]\fR
450 .ad
451 .sp .6
452 .RS 4n
453 The \fBshow-ifprop\fR subcommand is used to display the property values
454 for one or all of the IP interfaces.
455 .sp
456 The \fB-c\fR option (also \fB--parsable\fR) prints

```

```

457 the output in a parsable format.
458 .sp
459 The \fB-o\fR option (also \fB--output\fR) is used
460 to select which fields will be shown. The field value can be one of the
461 following names:
462 .sp
463 .ne 2
464 .na
465 .RS 4n
466 \fBALL\fR
467 .ad
468 .RS 4n
469 Display all fields
470 .RE

472 .sp
473 .ne 2
474 .na
475 \fBIFNAME\fR
476 .ad
477 .RS 4n
478 The name of the interface
479 .RE

481 .sp
482 .ne 2
483 .na
484 \fBPROPERTY\fR
485 .ad
486 .RS 4n
487 The name of the property
488 .RE

490 .sp
491 .ne 2
492 .na
493 \fBPROTO\fR
494 .ad
495 .RS 4n
496 The name of the protocol
497 .RE

499 .sp
500 .ne 2
501 .na
502 \fBPERM\fR
503 .ad
504 .RS 4n
505 If the property is readable ("r") and/or writable ("w").
506 .RE

508 .sp
509 .ne 2
510 .na
511 \fBCURRENT\fR
512 .ad
513 .RS 4n
514 The value of the property
515 .RE

517 .sp
518 .ne 2
519 .na
520 \fBPERSISTENT\fR
521 .ad
522 .RS 4n

```

523 The persistent value of the property  
524 .RE

526 .sp  
527 .ne 2  
528 .na  
529 \fBDEFAULT\fR  
530 .ad  
531 .RS 4n  
532 The default value of the property  
533 .RE

535 .sp  
536 .ne 2  
537 .na  
538 \fBPOSSIBLE\fR  
539 .ad  
540 .RS 4n  
541 The possible values for the property  
542 .RE  
543 .RE

545 .sp  
546 The \fB-p\fR option (also \fB--prop\fR) is used  
547 to specify which properties to display. See the \fBset-ifprop\fR  
548 subcommand for the list of property names.  
549 .sp  
550 The \fB-m\fR option (also \fB--module\fR) specifies which protocol  
551 to display.  
552 .sp

554 .RE

556 .sp  
557 .ne 2  
558 .na  
559 \fBfbcreate-addr\fR [\fB-t\fR] \fB-T\fR static [\fB-d\fR] \\  
560 \fB-a\fR {local|remote}=\fIaddr\fR[/\fIprefixlen\fR],... \fIaddrobj\fR\fR  
561 .br  
562 \fBfbcreate-addr\fR [\fB-t\fR] \fB-T\fR dhcp [\fB-w\fR \fIseconds\fR | forever  
563 .br  
564 \fBfbcreate-addr\fR [\fB-t\fR] \fB-T\fR addrconf [\fB-i\fR \fIinterface-id\fR]  
565 [\fB-p\fR {stateful|stateless}={yes|no},...] \fIaddrobj\fR\fR  
566 .ad  
567 .sp .6  
568 .RS 4n  
569 The \fBfbcreate-addr\fR subcommand is used to set an address on an IP interface.  
570 The address will be enabled but can disabled using the \fBfbdisable-addr\fR  
571 subcommand. This subcommand has three different forms, depending on the  
572 value of the \fB-T\fR option.  
573 .sp  
574 The \fB-t\fR option (also \fB--temporary\fR) means  
575 that the address is temporary and will not be persistent across reboots.  
576 .sp  
577 The \fB-T\fR static option creates a static addrobj. This takes the following  
578 options:  
579 .RS 4n

581 The \fB-d\fR option (also \fB--down\fR) means the address is down.  
582 .sp  
583 The \fB-a\fR option (also \fB--address\fR) specifies the address.  
584 The "local" or "remote" prefix can be used for a point-to-point interface.  
585 In this case, both addresses must be given.  
586 Otherwise, the equal sign ("=") should be omitted and the address should be  
587 provided by itself and with no second address.  
588 .sp

590 .RE

592 The \fB-T\fR dhcp option causes the address to be obtained via DHCP.  
593 This takes the following options:  
594 .RS 4n

596 The \fB-w\fR option (also \fB--wait\fR) gives the time, in seconds,  
597 that the command should wait to obtain an address.  
598 .sp

600 .RE

602 The \fB-T\fR addrconf option creates an auto-configured address.  
603 This takes the following options:  
604 .RS 4n

606 The \fB-i\fR option (also \fB--interface-id\fR) gives the interface ID to  
607 be used.  
608 .sp  
609 The \fB-p\fR option (also \fB--prop\fR) indicates which method of  
610 auto-configuration should be used.  
611 .sp

613 .RE  
614 .RE

616 .sp  
617 .ne 2  
618 .na  
619 \fBfbdown-addr\fR [\fB-t\fR] \fIaddrobj\fR\fR  
620 .ad  
621 .sp .6  
622 .RS 4n  
623 The \fBfbdown-addr\fR subcommand is used to down the address. This will  
624 stop packets from being sent or received.  
625 .sp  
626 The \fB-t\fR option (also \fB--temporary\fR) means  
627 that the down is temporary and will not be persistent across reboots.  
628 .sp

630 .RE

632 .sp  
633 .ne 2  
634 .na  
635 \fBfbup-addr\fR [\fB-t\fR] \fIaddrobj\fR\fR  
636 .ad  
637 .sp .6  
638 .RS 4n  
639 The \fBfbup-addr\fR subcommand is used to up the address. This will  
640 enable packets to be sent and received.  
641 .sp  
642 The \fB-t\fR option (also \fB--temporary\fR) means  
643 that the up is temporary and will not be persistent across reboots.  
644 .sp

646 .RE

648 .sp  
649 .ne 2  
650 .na  
651 \fBfbdisable-addr\fR [\fB-t\fR] \fIaddrobj\fR\fR  
652 .ad  
653 .sp .6  
654 .RS 4n

```

655 The \fBdisable-addr\fR subcommand is used to disable the address.
656 .sp
657 The \fB-t\fR option (also \fB--temporary\fR) means
658 that the disable is temporary and will not be persistent across reboots.
659 .sp
661 .RE

663 .sp
664 .ne 2
665 .na
666 \fB\fBenable-addr\fR [\fB-t\fR] \fIaddrobj\fR\fR
667 .ad
668 .sp .6
669 .RS 4n
670 The \fBenable-addr\fR subcommand is used to enable the address.
671 .sp
672 The \fB-t\fR option (also \fB--temporary\fR) means
673 that the enable is temporary and will not be persistent across reboots.
674 .sp
676 .RE

678 .sp
679 .ne 2
680 .na
681 \fB\fBrefresh-addr\fR [\fB-i\fR] \fIaddrobj\fR\fR
682 .ad
683 .sp .6
684 .RS 4n
685 The \fBrefresh-addr\fR subcommand is used to extend the lease for DHCP
686 addresses. It also restarts duplicate address detection for Static addresses.
687 .sp
688 The \fB-i\fR option (also \fB--inform\fR) means
689 that the network configuration will be obtained from DHCP without taking
690 a lease on the address.
691 .sp
693 .RE

695 .sp
696 .ne 2
697 .na
698 \fB\fBdelete-addr\fR [\fB-r\fR] \fIaddrobj\fR\fR
699 .ad
700 .sp .6
701 .RS 4n
702 The \fBdelete-addr\fR subcommand deletes the given address.
703 .sp
704 The \fB-r\fR option (also \fB--release\fR) is used for DHCP-assigned
705 addresses to indicate that the address should be released.
706 .sp
708 .RE

710 .sp
711 .ne 2
712 .na
713 \fB\fBshow-addr\fR [[\fB-p\fR] \fB-o\fR \fIifield\fR[,...]] [\fIaddrobj\fR]\fR
714 .ad
715 .sp .6
716 .RS 4n
717 The \fBshow-addr\fR subcommand is used to show the current address properties.
718 .sp
719 The \fB-p\fR option (also \fB--parsable\fR) prints
720 the output in a parsable format.

```

```

721 .sp
722 The \fB-o\fR option (also \fB--output\fR) is used
723 to select which fields will be shown. The field value can be one of the
724 following names:
725 .sp
726 .ne 2
727 .na
728 .RS 4n
729 \fBALL\fR
730 .ad
731 .RS 4n
732 Display all fields
733 .RE

735 .sp
736 .ne 2
737 .na
738 \fBADDROBJ\fR
739 .ad
740 .RS 4n
741 The name of the address
742 .RE

744 .sp
745 .ne 2
746 .na
747 \fBTYPE\fR
748 .ad
749 .RS 4n
750 The type of the address. It can be "static", "dhcp" or "addrconf".
751 .RE

753 .sp
754 .ne 2
755 .na
756 \fBSTATE\fR
757 .ad
758 .RS 4n
759 The state of the address. It can be one of the following values:
760 .sp
761 .ne 2
762 .na
763 .RS 4n
764 disabled s see the \fBdisable-addr\fR subcommand
765 .sp
766 down - see the \fBdown-addr\fR subcommand
767 .sp
768 duplicate - the address is a duplicate
769 .sp
770 inaccessible - the interface for this address has failed
771 .sp
772 ok - the address is up
773 .sp
774 tentative - duplicate address detection in progress
775 .RE
776 .RE

778 .sp
779 .ne 2
780 .na
781 \fBCURRENT\fR
782 .ad
783 .RS 4n
784 A set of single character flags indicating the following:
785 .sp
786 .ne 2

```

```

787 .na
788 .RS 4n
789 U - up
790 .br
791 u - unnumbered (matches another local address)
792 .br
793 p - private, not advertised to routing
794 .br
795 t - temporary IPv6 address
796 .br
797 d - deprecated (not used for outgoing packets)
798 .RE
799 .RE

801 .sp
802 .ne 2
803 .na
804 \fBPERSISTENT\fR
805 .ad
806 .RS 4n
807 A set of single character flags showing the configuration which will be used
808 when the address is enabled.
809 .sp
810 .ne 2
811 .na
812 .RS 4n
813 U - up
814 .br
815 p - private, not advertised to routing
816 .br
817 d - deprecated (not used for outgoing packets)
818 .RE
819 .RE

821 .sp
822 .ne 2
823 .na
824 \fBADDR\fR
825 .ad
826 .RS 4n
827 The address
828 .RE
829 .RE

831 .RE

833 .sp
834 .ne 2
835 .na
836 \fB\fBset-addrprop\fR [\fB-t\fR] \fB-p\fR \fIprop\fR=<\fIvalue\fR[,...]> \fIaddr
837 .ad
838 .sp .6
839 .RS 4n
840 The \fBset-addrprop\fR subcommand is used to set a property's value(s) on the
841 addrobj.
842 .sp
843 The \fB-t\fR option (also \fB--temporary\fR) means
844 that the setting is temporary and will not be persistent across reboots.
845 .sp
846 The \fB-p\fR option (also \fB--prop\fR) specifies the property name and
847 value(s). The property name can be one of the following:
848 .sp
849 .ne 2
850 .na

852 .RS 4n

```

```

854 \fBbroadcast\fR
855 .ad
856 .RS 4n
857 The broadcast address (read-only)
858 .RE

860 .sp
861 .ne 2
862 .na
863 \fBdeprecated\fR
864 .ad
865 .RS 4n
866 The address should not be used to send packets but can still receive packets.
867 Can be "on" or "off".
868 .RE

870 .sp
871 .ne 2
872 .na
873 \fBprefixlen\fR
874 .ad
875 .RS 4n
876 The number of bits in the IPv4 netmask or IPv6 prefix.
877 .RE

879 .sp
880 .ne 2
881 .na
882 \fBprivate\fR
883 .ad
884 .RS 4n
885 The address is not advertised to routing.
886 Can be "on" or "off".
887 .RE

889 .sp
890 .ne 2
891 .na
892 \fBtransmit\fR
893 .ad
894 .RS 4n
895 Packets can be transmitted.
896 Can be "on" or "off".
897 .RE

899 .sp
900 .ne 2
901 .na
902 \fBzone\fR
903 .ad
904 .RS 4n
905 The zone the addrobj is in.
906 .RE

908 .RE
909 .RE

911 .sp
912 .ne 2
913 .na
914 \fB\fBreset-addrprop\fR [\fB-t\fR] \fB-p\fR \fIprop\fR \fIaddrobj\fR\fR
915 .ad
916 .sp .6
917 .RS 4n
918 The \fBreset-addrprop\fR subcommand is used to reset an addrobj's property

```

```

919 value to the default.
920 .sp
921 The \fB-t\fR option (also \fB--temporary\fR) means
922 that the disable is temporary and will not be persistent across reboots.
923 .sp
924 The \fB-p\fR option (also \fB--prop\fR) specifies the property name.
925 See the \fBset-addrprop\fR subcommand for the list of property names.
926 .sp

928 .RE

930 .sp
931 .ne 2
932 .na
933 \fB\fBshow-addrprop\fR [[\fB-c\fR]\fB-o\fR \fIfield\fR[,...]] [\fB-p\fR \fIprop\
934 .ad
935 .sp .6
936 .RS 4n
937 The \fBshow-addrprop\fR subcommand is used to display the property values
938 for one or all of the addrobjs.
939 .sp
940 The \fB-c\fR option (also \fB--parsable\fR) prints
941 the output in a parsable format.
942 .sp
943 The \fB-o\fR option (also \fB--output\fR) is used
944 to select which fields will be shown. The field value can be one of the
945 following names:
946 .sp
947 .ne 2
948 .na
949 .RS 4n
950 \fBALL\fR
951 .ad
952 .RS 4n
953 Display all fields
954 .RE

956 .sp
957 .ne 2
958 .na
959 \fBADDROBJ\fR
960 .ad
961 .RS 4n
962 The name of the addrobj
963 .RE

965 .sp
966 .ne 2
967 .na
968 \fBPROPERTY\fR
969 .ad
970 .RS 4n
971 The name of the property
972 .RE

974 .sp
975 .ne 2
976 .na
977 \fBPERM\fR
978 .ad
979 .RS 4n
980 If the property is readable ("r") and/or writable ("w").
981 .RE

983 .sp
984 .ne 2

```

```

985 .na
986 \fBCURRENT\fR
987 .ad
988 .RS 4n
989 The value of the property
990 .RE

992 .sp
993 .ne 2
994 .na
995 \fBPERSISTENT\fR
996 .ad
997 .RS 4n
998 The persistent value of the property
999 .RE

1001 .sp
1002 .ne 2
1003 .na
1004 \fBDEFAULT\fR
1005 .ad
1006 .RS 4n
1007 The default value of the property
1008 .RE

1010 .sp
1011 .ne 2
1012 .na
1013 \fBPOSSIBLE\fR
1014 .ad
1015 .RS 4n
1016 The possible values for the property
1017 .RE
1018 .RE

1020 .sp
1021 The \fB-p\fR option (also \fB--prop\fR) is used
1022 to specify which properties to display. See the \fBset-addrprop\fR
1023 subcommand for the list of property names.
1024 .sp

1026 .RE

1028 .sp
1029 .ne 2
1030 .na
1031 \fBset-prop\fR [\fB-t\fR] \fB-p\fR \fIprop\fR[+|-]=<\fIvalue\fR[,...]> \fIpro
1032 .ad
1033 .sp .6
1034 .RS 4n
1035 The \fBset-prop\fR subcommand is used to set a property's value(s) on the
1036 protocol.
1037 .sp
1038 The \fB-t\fR option (also \fB--temporary\fR) means
1039 that the setting is temporary and will not be persistent across reboots.
1040 .sp
1041 The \fB-p\fR option (also \fB--prop\fR) specifies the property name and
1042 value(s). The optional [+|-] syntax can be used to add/remove values from the
1043 current list of values on the property.
1044 The property name can be one of the following:
1045 .sp
1046 .ne 2
1047 .na

1049 .RS 4n

```

```

1051 \fBecn\fR
1052 .ad
1053 .RS 4n
1054 Explicit congestion control (TCP-only)
1055 Can be "never", "passive" or "active".
1056 .RE

1058 \fBextra_priv_ports\fR
1059 .ad
1060 .RS 4n
1061 Additional privileged ports (SCTP, TCP or UDP)
1062 .RE

1064 \fBforwarding\fR
1065 .ad
1066 .RS 4n
1067 Packet forwarding is enabled.
1068 Can be "on" or "off".
1069 .RE

1071 \fBhoplimit\fR
1072 .ad
1073 .RS 4n
1074 The IPv6 hoplimit.
1075 .RE

1077 \fBlargest_anon_port\fR
1078 .ad
1079 .RS 4n
1080 Largest ephemeral port (SCTP, TCP or UDP)
1081 .RE

1083 \fBmax_buf\fR
1084 .ad
1085 .RS 4n
1086 Maximum receive or send buffer size (ICMP, SCTP, TCP, or UDP). This also
1087 sets the upper limit for the \fBrecv_buf\fB and \fBsend_buf\fB properties.
1088 .RE

1090 \fBrecv_buf\fR
1091 .ad
1092 .RS 4n
1093 Default receive buffer size (ICMP, SCTP, TCP, or UDP). The maximum value for
1094 this property is controlled by the \fBmax_buf\fR property.
1095 .RE

1097 \fBsack\fR
1098 .ad
1099 .RS 4n
1100 Selective acknowledgement (TCP).
1101 Can be "active", "passive" or "never".
1102 .RE

1104 \fBsend_buf\fR
1105 .ad
1106 .RS 4n
1107 Default send buffer size (ICMP, SCTP, TCP, or UDP). The maximum value for
1108 this property is controlled by the \fBmax_buf\fR property.
1109 .RE

1111 \fBsmallest_anon_port\fR
1112 .ad
1113 .RS 4n
1114 Smallest ephemeral port (SCTP, TCP or UDP)
1115 .RE

```

```

1117 \fBsmallest_nonpriv_port\fR
1118 .ad
1119 .RS 4n
1120 Smallest non-privileged port (SCTP, TCP or UDP)
1121 .RE

1123 \fBttl\fR
1124 .ad
1125 .RS 4n
1126 The IPv4 time-to-live.
1127 .RE

1129 .RE
1130 .RE

1132 .sp
1133 .ne 2
1134 .na
1135 \fB\fBreset-prop\fR [\fB-t\fR] \fB-p\fR \fIprop\fR \fIprotocol\fR\fR
1136 .ad
1137 .sp .6
1138 .RS 4n
1139 The \fBreset-prop\fR subcommand is used to reset a protocol's property
1140 value to the default.
1141 .sp
1142 The \fB-t\fR option (also \fB--temporary\fR) means
1143 that the disable is temporary and will not be persistent across reboots.
1144 .sp
1145 The \fB-p\fR option (also \fB--prop\fR) specifies the property name.
1146 See the \fBset-prop\fR subcommand for the list of property names.
1147 .sp

1149 .RE

1151 .sp
1152 .ne 2
1153 .na
1154 \fB\fBshow-prop\fR [[\fB-c\fR]\fB-o\fR \fIfield\fR[,...]] [\fB-p\fR \fIprop\fR,.
1155 .ad
1156 .sp .6
1157 .RS 4n
1158 The \fBshow-prop\fR subcommand is used to display the property values
1159 for one or all of the protocols.
1160 .sp
1161 The \fB-c\fR option (also \fB--parsable\fR) prints
1162 the output in a parsable format.
1163 .sp
1164 The \fB-o\fR option (also \fB--output\fR) is used
1165 to select which fields will be shown. The field value can be one of the
1166 following names:
1167 .sp
1168 .ne 2
1169 .na
1170 .RS 4n
1171 \fBALL\fR
1172 .ad
1173 .RS 4n
1174 Display all fields
1175 .RE

1177 .sp
1178 .ne 2
1179 .na
1180 \fBPROTO\fR
1181 .ad
1182 .RS 4n

```

```

1183 The name of the protocol
1184 .RE

1186 .sp
1187 .ne 2
1188 .na
1189 \fBPROPERTY\fR
1190 .ad
1191 .RS 4n
1192 The name of the property
1193 .RE

1195 .sp
1196 .ne 2
1197 .na
1198 \fBPERM\fR
1199 .ad
1200 .RS 4n
1201 If the property is readable ("r") and/or writable ("w").
1202 .RE

1204 .sp
1205 .ne 2
1206 .na
1207 \fBCURRENT\fR
1208 .ad
1209 .RS 4n
1210 The value of the property
1211 .RE

1213 .sp
1214 .ne 2
1215 .na
1216 \fBPERSISTENT\fR
1217 .ad
1218 .RS 4n
1219 The persistent value of the property
1220 .RE

1222 .sp
1223 .ne 2
1224 .na
1225 \fBDEFAULT\fR
1226 .ad
1227 .RS 4n
1228 The default value of the property
1229 .RE

1231 .sp
1232 .ne 2
1233 .na
1234 \fBPOSSIBLE\fR
1235 .ad
1236 .RS 4n
1237 The possible values for the property
1238 .RE
1239 .RE

1241 .sp
1242 The \fB-p\fR option (also \fB--prop\fR) is used
1243 to specify which properties to display. See the \fBset-prop\fR
1244 subcommand for the list of property names.
1245 .sp

1247 .RE

```

```

1249 .SH SEE ALSO
1250 .sp
1251 .LP
1252 \fBifconfig\fR(1M), \fBdladm\fR(1M), \fBndd\fR(1M), \fBzonecfg\fR(1M),
1253 \fBarp\fR(1M), \fBcfgadm\fR(1M), \fBifmpadm\fR(1M), \fBnsswitch.conf\fR(4),
1254 and \fBdhcp\fR(5).

```



\*\*\*\*\*

18373 Thu Jul 17 00:50:40 2014

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

manpage lint.

\*\*\*\*\*

```

1 \" te
2.\" Copyright 2013 Nexenta Systems, Inc. All rights reserved.
3.\" Copyright (c) 2008, 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 LOFIADM 1M \"Aug 28, 2013\"
8.SH NAME
9 lofiadm \- administer files available as block devices through lofi
10.SH SYNOPSIS
11.LP
12.nf
13 \fBlofiadm\fR [\fB-r\fR] \fB-a\fR \fIfile\fR [\fIdevice\fR]
14.fi

16.LP
17.nf
18 \fBlofiadm\fR [\fB-r\fR] \fB-c\fR \fIcrypto_algorithm\fR \fB-a\fR \fIfile\fR [\f
19.fi

21.LP
22.nf
23 \fBlofiadm\fR [\fB-r\fR] \fB-c\fR \fIcrypto_algorithm\fR \fB-k\fR \fIraw_key_fil
24.fi

26.LP
27.nf
28 \fBlofiadm\fR [\fB-r\fR] \fB-c\fR \fIcrypto_algorithm\fR \fB-T\fR \fItoken_key\f
29.fi

31.LP
32.nf
33 \fBlofiadm\fR [\fB-r\fR] \fB-c\fR \fIcrypto_algorithm\fR \fB-T\fR \fItoken_key\f
34 \fB-k\fR \fIwrapped_key_file\fR \fB-a\fR \fIfile\fR [\fIdevice\fR]
35.fi

37.LP
38.nf
39 \fBlofiadm\fR [\fB-r\fR] \fB-c\fR \fIcrypto_algorithm\fR \fB-e\fR \fB-a\fR \fIffi
40.fi

42.LP
43.nf
44 \fBlofiadm\fR \fB-C\fR \fIalgorithm\fR [\fB-s\fR \fIsegment_size\fR] \fIfile\fR
45.fi

47.LP
48.nf
49 \fBlofiadm\fR \fB-d\fR \fIfile\fR | \fIdevice\fR
50.fi

52.LP
53.nf
54 \fBlofiadm\fR \fB-U\fR \fIfile\fR
55.fi

57.LP
58.nf
59 \fBlofiadm\fR [ \fIfile\fR | \fIdevice\fR]
60.fi

```

62 .SH DESCRIPTION

63 .sp

64 .LP

65 \fBlofiadm\fR administers \fBblofi\fR, the loopback file driver. \fBblofi\fR  
66 allows a file to be associated with a block device. That file can then be  
67 accessed through the block device. This is useful when the file contains an  
68 image of some filesystem (such as a floppy or \fBCD-ROM\fR image), because the  
69 block device can then be used with the normal system utilities for mounting,  
70 checking or repairing filesystems. See \fBfsck\fR(1M) and \fBmount\fR(1M).

71 .sp

72 .LP

73 Use \fBlofiadm\fR to add a file as a loopback device, remove such an

74 association, or print information about the current associations.

75 .sp

76 .LP

77 Encryption and compression options are mutually exclusive on the command line.  
78 Further, an encrypted file cannot be compressed later, nor can a compressed  
79 file be encrypted later.

81 In the global zone, \fBlofiadm\fR can be used on both the global

82 zone devices and all devices owned by other non-global zones on the system.

83 .sp

84 .LP

84 .SH OPTIONS

85 .sp

86 .LP

87 The following options are supported:

88 .sp

89 .ne 2

90 .na

91 \fBlofiadm\fR \fB-a\fR \fIfile\fR [\fIdevice\fR]\fR

92 .ad

93 .sp .6

94 .RS 4n

95 Add \fIfile\fR as a block device.

96 .sp

97 If \fIdevice\fR is not specified, an available device is picked.

98 .sp

99 If \fIdevice\fR is specified, \fBlofiadm\fR attempts to assign it to  
100 \fIfile\fR. \fIdevice\fR must be available or \fBlofiadm\fR will fail. The  
101 ability to specify a device is provided for use in scripts that wish to  
102 reestablish a particular set of associations.

103 .RE

105 .sp

106 .ne 2

107 .na

108 \fBlofiadm\fR \fB-C\fR {\fIgzip\fR | \fIgzip-N\fR | \fIlzma\fR}\fR

109 .ad

110 .sp .6

111 .RS 4n

112 Compress the file with the specified compression algorithm.

113 .sp

114 The \fBgzip\fR compression algorithm uses the same compression as the

115 open-source \fBgzip\fR command. You can specify the \fBgzip\fR level by using

116 the value \fBgzip-*IN*\fR where *IN* is 6 (fast) or 9 (best compression

117 ratio). Currently, \fBgzip\fR, without a number, is equivalent to \fBgzip-6\fR

118 (which is also the default for the \fBgzip\fR command).

119 .sp

120 \fIlzma\fR stands for the LZMA (Lempel-Ziv-Markov) compression algorithm.

121 .sp

122 Note that you cannot write to a compressed file, nor can you mount a compressed

123 file read/write.

124 .RE

126 .sp

```

127 .ne 2
128 .na
129 \fB\fB-d\fR \fIfile\fR | \fIdevice\fR\fR
130 .ad
131 .sp .6
132 .RS 4n
133 Remove an association by \fIfile\fR or \fIdevice\fR name, if the associated
134 block device is not busy, and deallocates the block device.
135 .RE

137 .sp
138 .ne 2
139 .na
140 \fB\fB-r\fR
141 .ad
142 .sp .6
143 .RS 4n
144 If the \fB-r\fR option is specified before the \fB-a\fR option, the
145 \fIdevice\fR will be opened read-only.
146 .RE

148 .sp
149 .ne 2
150 .na
151 \fB\fB-s\fR \fIsegment_size\fR\fR
152 .ad
153 .sp .6
154 .RS 4n
155 The segment size to use to divide the file being compressed. \fIsegment_size\fR
156 can be an integer multiple of 512.
157 .RE

159 .sp
160 .ne 2
161 .na
162 \fB\fB-U\fR \fIfile\fR\fR
163 .ad
164 .sp .6
165 .RS 4n
166 Uncompress a compressed file.
167 .RE

169 .sp
170 .LP
171 The following options are used when the file is encrypted:
172 .sp
173 .ne 2
174 .na
175 \fB\fB-c\fR \fIcrypto_algorithm\fR\fR
176 .ad
177 .sp .6
178 .RS 4n
179 Select the encryption algorithm. The algorithm must be specified when
180 encryption is enabled because the algorithm is not stored in the disk image.
181 .sp
182 If none of \fB-e\fR, \fB-k\fR, or \fB-T\fR is specified, \fB-lofiadm\fR prompts
183 for a passphrase, with a minimum length of eight characters, to be entered .
184 The passphrase is used to derive a symmetric encryption key using PKCS#5 PBKD2.
185 .RE

187 .sp
188 .ne 2
189 .na
190 \fB\fB-k\fR \fIraw_key_file\fR | \fIwrapped_key_file\fR\fR
191 .ad
192 .sp .6

```

```

193 .RS 4n
194 Path to raw or wrapped symmetric encryption key. If a PKCS#11 object is also
195 given with the \fB-T\fR option, then the key is wrapped by that object. If
196 \fB-T\fR is not specified, the key is used raw.
197 .RE

199 .sp
200 .ne 2
201 .na
202 \fB\fB-T\fR \fItoken_key\fR\fR
203 .ad
204 .sp .6
205 .RS 4n
206 The key in a PKCS#11 token to use for the encryption or for unwrapping the key
207 file.
208 .sp
209 If \fB-k\fR is also specified, \fB-T\fR identifies the unwrapping key, which
210 must be an RSA private key.
211 .RE

213 .sp
214 .ne 2
215 .na
216 \fB\fB-e\fR\fR
217 .ad
218 .sp .6
219 .RS 4n
220 Generate an ephemeral symmetric encryption key.
221 .RE

223 .SH OPERANDS
224 .sp
225 .LP
226 The following operands are supported:
227 .sp
228 .ne 2
229 .na
230 \fB\fIcrypto_algorithm\fR\fR
231 .ad
232 .sp .6
233 .RS 4n
234 One of: \fBaes-128-cbc\fR, \fBaes-192-cbc\fR, \fBaes-256-cbc\fR,
235 \fBdes3-cbc\fR, \fBblowfish-cbc\fR.
236 .RE

238 .sp
239 .ne 2
240 .na
241 \fB\fIdevice\fR\fR
242 .ad
243 .sp .6
244 .RS 4n
245 Display the file name associated with the block device \fIdevice\fR.
246 .sp
247 Without arguments, print a list of the current associations. Filenames must be
248 valid absolute pathnames.
249 .sp
250 When a file is added, it is opened for reading or writing by root. Any
251 restrictions apply (such as restricted root access over \fBNFS\fR). The file is
252 held open until the association is removed. It is not actually accessed until
253 the block device is used, so it will never be written to if the block device is
254 only opened read-only.

256 Note that the filename may appear as "?" if it is not possible to resolve the
257 path in the current context (for example, if it's an NFS path in a non-global
258 zone).

```

```

259 .RE

261 .sp
262 .ne 2
263 .na
264 \fB\fIfile\fR\fR
265 .ad
266 .sp .6
267 .RS 4n
268 Display the block device associated with \fIfile\fR.
269 .RE

271 .sp
272 .ne 2
273 .na
274 \fB\fIraw_key_file\fR\fR
275 .ad
276 .sp .6
277 .RS 4n
278 Path to a file of the appropriate length, in bits, to use as a raw symmetric
279 encryption key.
280 .RE

282 .sp
283 .ne 2
284 .na
285 \fB\fItoken_key\fR\fR
286 .ad
287 .sp .6
288 .RS 4n
289 PKCS#11 token object in the format:
290 .sp
291 .in +2
292 .nf
293 \fItoken_name\fR:\fImanufacturer_id\fR:\fIserial_number\fR:\fIkey_label\fR
294 .fi
295 .in -2
296 .sp

298 All but the key label are optional and can be empty. For example, to specify a
299 token object with only its key label \fBMylofiKey\fR, use:
300 .sp
301 .in +2
302 .nf
303 -T ::MylofiKey
304 .fi
305 .in -2
306 .sp

308 .RE

310 .sp
311 .ne 2
312 .na
313 \fB\fIwrapped_key_file\fR\fR
314 .ad
315 .sp .6
316 .RS 4n
317 Path to file containing a symmetric encryption key wrapped by the RSA private
318 key specified by \fB-T\fR.
319 .RE

321 .SH EXAMPLES
322 .LP
323 \fBExample 1\fR Mounting an Existing CD-ROM Image
324 .sp

```

```

325 .LP
326 You should ensure that Solaris understands the image before creating the
327 \fBCD\fR. \fBblofi\fR allows you to mount the image and see if it works.

329 .sp
330 .LP
331 This example mounts an existing \fBCD-ROM\fR image (\fBsparc.iso\fR), of the
332 \fBRed Hat 6.0 CD\fR which was downloaded from the Internet. It was created
333 with the \fBmkisofs\fR utility from the Internet.

335 .sp
336 .LP
337 Use \fBblofiadm\fR to attach a block device to it:

339 .sp
340 .in +2
341 .nf
342 # \fBblofiadm -a /home/mike_s/RH6.0/sparc.iso\fR
343 /dev/lofi/1
344 .fi
345 .in -2
346 .sp

348 .sp
349 .LP
350 \fBblofiadm\fR picks the device and prints the device name to the standard
351 output. You can run \fBblofiadm\fR again by issuing the following command:

353 .sp
354 .in +2
355 .nf
356 # \fBblofiadm\fR
357 Block Device      File                      Options
358 /dev/lofi/1      /home/mike_s/RH6.0/sparc.iso  -
359 .fi
360 .in -2
361 .sp

363 .sp
364 .LP
365 Or, you can give it one name and ask for the other, by issuing the following
366 command:

368 .sp
369 .in +2
370 .nf
371 # \fBblofiadm /dev/lofi/1\fR
372 /home/mike_s/RH6.0/sparc.iso
373 .fi
374 .in -2
375 .sp

377 .sp
378 .LP
379 Use the \fBmount\fR command to mount the image:

381 .sp
382 .in +2
383 .nf
384 # \fBmount -F hsfs -o ro /dev/lofi/1 /mnt\fR
385 .fi
386 .in -2
387 .sp

389 .sp
390 .LP

```

```

391 Check to ensure that Solaris understands the image:

393 .sp
394 .in +2
395 .nf
396 # \fBdf -k /mnt\fR
397 Filesystem          kbytes    used    avail capacity  Mounted on
398 /dev/lofi/1         512418    512418    0    100%    /mnt
399 # \fBls /mnt\fR
400 \&./                RedHat/    doc/       ls-lR      rr_moved/
401 \&../                TRANS.TBL  dosutils/  ls-lR.gz   sbin@
402 \&.buildlog         bin@       etc@        misc/      tmp/
403 COPYING             boot/      images/     mnt/       usr@
404 README              boot.cat*  kernels/    modules/
405 RPM-PGP-KEY         dev@      lib@        proc/
406 .fi
407 .in -2
408 .sp

410 .sp
411 .LP
412 Solaris can mount the CD-ROM image, and understand the filenames. The image was
413 created properly, and you can now create the \fBCD-ROM\fR with confidence.

415 .sp
416 .LP
417 As a final step, unmount and detach the images:

419 .sp
420 .in +2
421 .nf
422 # \fBumount /mnt\fR
423 # \fBlofiadm -d /dev/lofi/1\fR
424 # \fBlofiadm\fR
425 Block Device          File          Options
426 .fi
427 .in -2
428 .sp

430 .LP
431 \fBExample 2 \fRMounting a Floppy Image
432 .sp
433 .LP
434 This is similar to the first example.

436 .sp
437 .LP
438 Using \fBlofi\fR to help you mount files that contain floppy images is helpful
439 if a floppy disk contains a file that you need, but the machine which you are
440 on does not have a floppy drive. It is also helpful if you do not want to take
441 the time to use the \fBdd\fR command to copy the image to a floppy.

443 .sp
444 .LP
445 This is an example of getting to \fB MDB \fR floppy for Solaris on an x86
446 platform:

448 .sp
449 .in +2
450 .nf
451 # \fBlofiadm -a /export/s28/MDB_s28x_wos/latest/boot.3\fR
452 /dev/lofi/1
453 # \fBmount -F pcfs /dev/lofi/1 /mnt\fR
454 # \fBls /mnt\fR
455 \&./                COMMENT.BAT*  RC.D/       SOLARIS.MAP*
456 \&../                IDENT*        REPLACE.BAT*  X/

```

```

457 APPEND.BAT*    MAKEDIR.BAT*    SOLARIS/
458 # \fBumount /mnt\fR
459 # \fBlofiadm -d /export/s28/MDB_s28x_wos/latest/boot.3\fR
460 .fi
461 .in -2
462 .sp

464 .LP
465 \fBExample 3 \fRMaking a \fBUFS\fR Filesystem on a File
466 .sp
467 .LP
468 Making a \fBUFS\fR filesystem on a file can be useful, particularly if a test
469 suite requires a scratch filesystem. It can be painful (or annoying) to have to
470 repartition a disk just for the test suite, but you do not have to. You can
471 \fBnewfs\fR a file with \fBlofi\fR

473 .sp
474 .LP
475 Create the file:

477 .sp
478 .in +2
479 .nf
480 # \fBmkfile 35m /export/home/test\fR
481 .fi
482 .in -2
483 .sp

485 .sp
486 .LP
487 Attach it to a block device. You also get the character device that \fBnewfs\fR
488 requires, so \fBnewfs\fR that:

490 .sp
491 .in +2
492 .nf
493 # \fBlofiadm -a /export/home/test\fR
494 /dev/lofi/1
495 # \fBnewfs /dev/rlofi/1\fR
496 newfs: construct a new file system /dev/rlofi/1: (y/n)? \fBy\fR
497 /dev/rlofi/1: 71638 sectors in 119 cylinders of 1 tracks, 602 sectors
498 35.0MB in 8 cyl groups (16 c/g, 4.70MB/g, 2240 i/g)
499 super-block backups (for fsck -F ufs -o b=#) at:
500 32, 9664, 19296, 28928, 38560, 48192, 57824, 67456,
501 .fi
502 .in -2
503 .sp

505 .sp
506 .LP
507 Note that \fBUfs\fR might not be able to use the entire file. Mount and use the
508 filesystem:

510 .sp
511 .in +2
512 .nf
513 # \fBmount /dev/lofi/1 /mnt\fR
514 # \fBdf -k /mnt\fR
515 Filesystem          kbytes    used    avail capacity  Mounted on
516 /dev/lofi/1         33455     9    30101    1%    /mnt
517 # \fBls /mnt\fR
518 \&./                ../          lost+found/
519 # \fBumount /mnt\fR
520 # \fBlofiadm -d /dev/lofi/1\fR
521 .fi
522 .in -2

```

```

523 .sp
525 .LP
526 \fBExample 4 \fRCreating a PC (FAT) File System on a Unix File
527 .sp
528 .LP
529 The following series of commands creates a \fBFAT\fR file system on a Unix
530 file. The file is associated with a block device created by \fBblofiadm\fR.

```

```

532 .sp
533 .in +2
534 .nf
535 # \fBmkfile 10M /export/test/testfs\fR
536 # \fBblofiadm -a /export/test/testfs\fR
537 /dev/lofi/1
538 \fBNote use of\fR rlofi\fB, not\fR lofi\fB, in following command.\fR
539 # \fBmkfs -F pcfs -o nofdisk,size=20480 /dev/rlofi/1\fR
540 \fBConstruct a new FAT file system on /dev/rlofi/1: (y/n)?\fR y
541 # \fBmount -F pcfs /dev/lofi/1 /mnt\fR
542 # \fBcd /mnt\fR
543 # \fBdf -k .\fR
544 Filesystem          kbytes    used   avail capacity  Mounted on
545 /dev/lofi/1         10142      0   10142     0%   /mnt
546 .fi
547 .in -2
548 .sp

```

```

550 .LP
551 \fBExample 5 \fRCompressing an Existing CD-ROM Image
552 .sp
553 .LP
554 The following example illustrates compressing an existing CD-ROM image
555 (\fBsolaris.iso\fR), verifying that the image is compressed, and then
556 uncompressing it.

```

```

558 .sp
559 .in +2
560 .nf
561 # \fBblofiadm -C gzip /export/home/solaris.iso\fR
562 .fi
563 .in -2
564 .sp

```

```

566 .sp
567 .LP
568 Use \fBblofiadm\fR to attach a block device to it:

```

```

570 .sp
571 .in +2
572 .nf
573 # \fBblofiadm -a /export/home/solaris.iso\fR
574 /dev/lofi/1
575 .fi
576 .in -2
577 .sp

```

```

579 .sp
580 .LP
581 Check if the mapped image is compressed:

```

```

583 .sp
584 .in +2
585 .nf
586 # \fBblofiadm\fR
587 Block Device      File          Options
588 /dev/lofi/1       /export/home/solaris.iso  Compressed(gzip)

```

```

589 /dev/lofi/2          /export/home/regular.iso      -
590 .fi
591 .in -2
592 .sp

```

```

594 .sp
595 .LP
596 Unmap the compressed image and uncompress it:

```

```

598 .sp
599 .in +2
600 .nf
601 # \fBblofiadm -d /dev/lofi/1\fR
602 # \fBblofiadm -U /export/home/solaris.iso\fR
603 .fi
604 .in -2
605 .sp

```

```

607 .LP
608 \fBExample 6 \fRCreating an Encrypted UFS File System on a File
609 .sp
610 .LP
611 This example is similar to the example of making a UFS filesystem on a file,
612 above.

```

```

614 .sp
615 .LP
616 Create the file:

```

```

618 .sp
619 .in +2
620 .nf
621 # \fBmkfile 35m /export/home/test\fR
622 .fi
623 .in -2
624 .sp

```

```

626 .sp
627 .LP
628 Attach the file to a block device and specify that the file image is encrypted.
629 As a result of this command, you obtain the character device, which is
630 subsequently used by \fBnewfs\fR:

```

```

632 .sp
633 .in +2
634 .nf
635 # \fBblofiadm -c aes-256-cbc -a /export/home/secrets\fR
636 Enter passphrase: \fBMy-M0th3r:l0v3s_m3+4lw4ys!\fR          (\fBnot echoed\fR)
637 Re-enter passphrase: \fBMy-M0th3r:l0v3s_m3+4lw4ys!\fR          (\fBnot echoed\fR)
638 /dev/lofi/1

```

```

640 # \fBnewfs /dev/rlofi/1\fR
641 newfs: construct a new file system /dev/rlofi/1: (y/n)? \fBy\fR
642 /dev/rlofi/1: 71638 sectors in 119 cylinders of 1 tracks, 602 sectors
643 35.0MB in 8 cyl groups (16 c/g, 4.70MB/g, 2240 i/g)
644 super-block backups (for fsck -F ufs -o b=#) at:
645 32, 9664, 19296, 28928, 38560, 48192, 57824, 67456,
646 .fi
647 .in -2
648 .sp

```

```

650 .sp
651 .LP
652 The mapped file system shows that encryption is enabled:

```

```

654 .sp

```

```

655 .in +2
656 .nf
657 # \fBlofiadm\fR
658 Block Device      File      Options
659 /dev/lofi/1      /export/home/secrets  Encrypted
660 .fi
661 .in -2
662 .sp

664 .sp
665 .LP
666 Mount and use the filesystem:

668 .sp
669 .in +2
670 .nf
671 # \fBmount /dev/lofi/1 /mnt\fR
672 # \fBcp moms_secret_*_recipe /mnt\fR
673 # \fBls /mnt\fR
674 \&./      moms_secret_cookie_recipe  moms_secret_soup_recipe
675 \&../      moms_secret_fudge_recipe    moms_secret_stuffing_recipe
676 lost+found/ moms_secret_meatloaf_recipe  moms_secret_waffle_recipe
677 # \fBumount /mnt\fR
678 # \fBlofiadm -d /dev/lofi/1\fR
679 .fi
680 .in -2
681 .sp

683 .sp
684 .LP
685 Subsequent attempts to map the filesystem with the wrong key or the wrong
686 encryption algorithm will fail:

688 .sp
689 .in +2
690 .nf
691 # \fBlofiadm -c blowfish-cbc -a /export/home/secrets\fR
692 Enter passphrase: \fBmommy\fR      (\fInot echoed\fR)
693 Re-enter passphrase: \fBmommy\fR    (\fInot echoed\fR)
694 lofiadm: could not map file /root/lofi: Invalid argument
695 # \fBlofiadm\fR
696 Block Device      File      Options
697 #
698 .fi
699 .in -2
700 .sp

702 .sp
703 .LP
704 Attempts to map the filesystem without encryption will succeed, however
705 attempts to mount and use the filesystem will fail:

707 .sp
708 .in +2
709 .nf
710 # \fBlofiadm -a /export/home/secrets\fR
711 /dev/lofi/1
712 # \fBlofiadm\fR
713 Block Device      File      Options
714 /dev/lofi/1      /export/home/secrets  -
715 # \fBmount /dev/lofi/1 /mnt\fR
716 mount: /dev/lofi/1 is not this fstype
717 #
718 .fi
719 .in -2
720 .sp

```

```

722 .SH ENVIRONMENT VARIABLES
723 .sp
724 .LP
725 See \fBenvron\fR(5) for descriptions of the following environment variables
726 that affect the execution of \fBlofiadm\fR: \fBLC_CTYPE\fR, \fBLC_MESSAGES\fR
727 and \fBNLSPATH\fR.
728 .SH EXIT STATUS
729 .sp
730 .LP
731 The following exit values are returned:
732 .sp
733 .ne 2
734 .na
735 \fB0\fR
736 .ad
737 .sp .6
738 .RS 4n
739 Successful completion.
740 .RE

742 .sp
743 .ne 2
744 .na
745 \fB>0\fR
746 .ad
747 .sp .6
748 .RS 4n
749 An error occurred.
750 .RE

752 .SH SEE ALSO
753 .sp
754 .LP
755 \fBfsck\fR(1M), \fBmount\fR(1M), \fBmount_ufs\fR(1M), \fBnewfs\fR(1M),
756 \fBattributes\fR(5), \fBlofi\fR(7D), \fBlofs\fR(7FS)
757 .SH NOTES
758 .sp
759 .LP
760 Just as you would not directly access a disk device that has mounted file
761 systems, you should not access a file associated with a block device except
762 through the \fBlofi\fR file driver. It might also be appropriate to ensure that
763 the file has appropriate permissions to prevent such access.
764 .sp
765 .LP
766 The abilities of \fBlofiadm\fR, and who can use them, are controlled by the
767 permissions of \fB/dev/lofi/1\fR. Read-access allows query operations, such as
768 listing all the associations. Write-access is required to do any state-changing
769 operations, like adding an association. As shipped, \fB/dev/lofi/1\fR is owned
770 by \fBroot\fR, in group \fBsys\fR, and mode \fB0644\fR, so all users can do
771 query operations but only root can change anything. The administrator can give
772 users write-access, allowing them to add or delete associations, but that is
773 very likely a security hole and should probably only be given to a trusted
774 group.
775 .sp
776 .LP
777 When mounting a filesystem image, take care to use appropriate mount options.
778 In particular, the \fBnosuid\fR mount option might be appropriate for \fBUBFS\fR
779 images whose origin is unknown. Also, some options might not be useful or
780 appropriate, like \fBlogging\fR or \fBforcedirectio\fR for \fBUBFS\fR. For
781 compatibility purposes, a raw device is also exported along with the block
782 device. For example, \fBnewfs\fR(1M) requires one.
783 .sp
784 .LP
785 The output of \fBlofiadm\fR (without arguments) might change in future
786 releases.

```

new/usr/src/man/man1m/root\_archive.1m

1

```
*****  
2466 Thu Jul 17 00:50:40 2014  
new/usr/src/man/man1m/root_archive.1m  
manpage lint.  
*****  
_____unchanged_portion_omitted_
```

\*\*\*\*\*

43446 Thu Jul 17 00:50:40 2014

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

manpage lint.

\*\*\*\*\*

```

1  \' te
2  .\" Copyright (c) 2004, 2009 Sun Microsystems, Inc. All Rights Reserved.
3  .\" Copyright 2013 Joyent, 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 limitati
6  .\" fields enclosed by brackets \"[]\" replaced with your own identifying informat
7  .TH ZONECFG 1M \"Feb 28, 2014\"
8  .SH NAME
9  zonecfg \- set up zone configuration
10 .SH SYNOPSIS
11 .LP
12 .nf
13 \fBzonecfg\fR \fB-z\fR \fIzonename\fR
14 .fi

16 .LP
17 .nf
18 \fBzonecfg\fR \fB-z\fR \fIzonename\fR \fIsubcommand\fR
19 .fi

21 .LP
22 .nf
23 \fBzonecfg\fR \fB-z\fR \fIzonename\fR \fB-f\fR \fIcommand_file\fR
24 .fi

26 .LP
27 .nf
28 \fBzonecfg\fR help
29 .fi

31 .SH DESCRIPTION
32 .sp
33 .LP
34 The \fBzonecfg\fR utility creates and modifies the configuration of a zone.
35 Zone configuration consists of a number of resources and properties.
36 .sp
37 .LP
38 To simplify the user interface, \fBzonecfg\fR uses the concept of a scope. The
39 default scope is global.
40 .sp
41 .LP
42 The following synopsis of the \fBzonecfg\fR command is for interactive usage:
43 .sp
44 .in +2
45 .nf
46 zonecfg \fB-z\fR \fIzonename subcommand\fR
47 .fi
48 .in -2
49 .sp

51 .sp
52 .LP
53 Parameters changed through \fBzonecfg\fR do not affect a running zone. The zone
54 must be rebooted for the changes to take effect.
55 .sp
56 .LP
57 In addition to creating and modifying a zone, the \fBzonecfg\fR utility can
58 also be used to persistently specify the resource management settings for the
59 global zone.
60 .sp
61 .LP

```

```

62 In the following text, "rctl" is used as an abbreviation for "resource
63 control". See \fBresource_controls\fR(5).
64 .sp
65 .LP
66 Every zone is configured with an associated brand. The brand determines the
67 user-level environment used within the zone, as well as various behaviors for
68 the zone when it is installed, boots, or is shutdown. Once a zone has been
69 installed the brand cannot be changed. The default brand is determined by the
70 installed distribution in the global zone. Some brands do not support all of
71 the \fBzonecfg\fR properties and resources. See the brand-specific man page for
72 more details on each brand. For an overview of brands, see the \fBbrands\fR(5)
73 man page.
74 .SS "Resources"
75 .sp
76 .LP
77 The following resource types are supported:
78 .sp
79 .ne 2
80 .na
81 \fB\battr\fR
82 .ad
83 .sp .6
84 .RS 4n
85 Generic attribute.
86 .RE

88 .sp
89 .ne 2
90 .na
91 \fB\b capped-cpu\fR
92 .ad
93 .sp .6
94 .RS 4n
95 Limits for CPU usage.
96 .RE

98 .sp
99 .ne 2
100 .na
101 \fB\b capped-memory\fR
102 .ad
103 .sp .6
104 .RS 4n
105 Limits for physical, swap, and locked memory.
106 .RE

108 .sp
109 .ne 2
110 .na
111 \fB\b dataset\fR
112 .ad
113 .sp .6
114 .RS 4n
115 \fBZFS\fR dataset.
116 .RE

118 .sp
119 .ne 2
120 .na
121 \fB\b dedicated-cpu\fR
122 .ad
123 .sp .6
124 .RS 4n
125 Subset of the system's processors dedicated to this zone while it is running.
126 .RE

```



```

128 .sp
129 .ne 2
130 .na
131 \fB\fBdevice\fR\fR
132 .ad
133 .sp .6
134 .RS 4n
135 Device.
136 .RE

138 .sp
139 .ne 2
140 .na
141 \fB\fBfs\fR\fR
142 .ad
143 .sp .6
144 .RS 4n
145 file-system
146 .RE

148 .sp
149 .ne 2
150 .na
151 \fB\fBnet\fR\fR
152 .ad
153 .sp .6
154 .RS 4n
155 Network interface.
156 .RE

158 .sp
159 .ne 2
160 .na
161 \fB\fBbrctl\fR\fR
162 .ad
163 .sp .6
164 .RS 4n
165 Resource control.
166 .RE

168 .SS "Properties"
169 .sp
170 .LP
171 Each resource type has one or more properties. There are also some global
172 properties, that is, properties of the configuration as a whole, rather than of
173 some particular resource.
174 .sp
175 .LP
176 The following properties are supported:
177 .sp
178 .ne 2
179 .na
180 \fB(global)\fR
181 .ad
182 .sp .6
183 .RS 4n
184 \fBzone\fR
185 .RE

187 .sp
188 .ne 2
189 .na
190 \fB(global)\fR
191 .ad
192 .sp .6
193 .RS 4n

```

```

194 \fBzonepath\fR
195 .RE

197 .sp
198 .ne 2
199 .na
200 \fB(global)\fR
201 .ad
202 .sp .6
203 .RS 4n
204 \fBautoboot\fR
205 .RE

207 .sp
208 .ne 2
209 .na
210 \fB(global)\fR
211 .ad
212 .sp .6
213 .RS 4n
214 \fBbootargs\fR
215 .RE

217 .sp
218 .ne 2
219 .na
220 \fB(global)\fR
221 .ad
222 .sp .6
223 .RS 4n
224 \fBpool\fR
225 .RE

227 .sp
228 .ne 2
229 .na
230 \fB(global)\fR
231 .ad
232 .sp .6
233 .RS 4n
234 \fBlimitpriv\fR
235 .RE

237 .sp
238 .ne 2
239 .na
240 \fB(global)\fR
241 .ad
242 .sp .6
243 .RS 4n
244 \fBbrand\fR
245 .RE

247 .sp
248 .ne 2
249 .na
250 \fB(global)\fR
251 .ad
252 .sp .6
253 .RS 4n
254 \fBcpu-shares\fR
255 .RE

257 .sp
258 .ne 2
259 .na

```

```

260 \fB(global)\fR
261 .ad
262 .sp .6
263 .RS 4n
264 \fBhostid\fR
265 .RE

267 .sp
268 .ne 2
269 .na
270 \fB(global)\fR
271 .ad
272 .sp .6
273 .RS 4n
274 \fBmax-lwps\fR
275 .RE

277 .sp
278 .ne 2
279 .na
280 \fB(global)\fR
281 .ad
282 .sp .6
283 .RS 4n
284 \fBmax-msg-ids\fR
285 .RE

287 .sp
288 .ne 2
289 .na
290 \fB(global)\fR
291 .ad
292 .sp .6
293 .RS 4n
294 \fBmax-sem-ids\fR
295 .RE

297 .sp
298 .ne 2
299 .na
300 \fB(global)\fR
301 .ad
302 .sp .6
303 .RS 4n
304 \fBmax-shm-ids\fR
305 .RE

307 .sp
308 .ne 2
309 .na
310 \fB(global)\fR
311 .ad
312 .sp .6
313 .RS 4n
314 \fBmax-shm-memory\fR
315 .RE

317 .sp
318 .ne 2
319 .na
320 \fB(global)\fR
321 .ad
322 .sp .6
323 .RS 4n
324 \fBscheding-class\fR
325 .RE

```

```

327 .sp
328 .ne 2
329 .na
330 .B (global)
331 .ad
332 .sp .6
333 .RS 4n
334 .B fs-allowed
335 .RE

337 .sp
338 .ne 2
339 .na
340 \fB\fBs\fR\fR
341 .ad
342 .sp .6
343 .RS 4n
344 \fBdir\fR, \fBspecial\fR, \fBraw\fR, \fBtype\fR, \fBoptions\fR
345 .RE

347 .sp
348 .ne 2
349 .na
350 \fB\fBnet\fR\fR
351 .ad
352 .sp .6
353 .RS 4n
354 \fBaddress\fR, \fBphysical\fR, \fBdefrouter\fR
355 .RE

357 .sp
358 .ne 2
359 .na
360 \fB\fBdevice\fR\fR
361 .ad
362 .sp .6
363 .RS 4n
364 \fBmatch\fR
365 .RE

367 .sp
368 .ne 2
369 .na
370 \fB\fBrc1\fR\fR
371 .ad
372 .sp .6
373 .RS 4n
374 \fBname\fR, \fBvalue\fR
375 .RE

377 .sp
378 .ne 2
379 .na
380 \fB\fBattr\fR\fR
381 .ad
382 .sp .6
383 .RS 4n
384 \fBname\fR, \fBtype\fR, \fBvalue\fR
385 .RE

387 .sp
388 .ne 2
389 .na
390 \fB\fBdataset\fR\fR
391 .ad

```

```

392 .sp .6
393 .RS 4n
394 \fBname\fR
395 .RE

397 .sp
398 .ne 2
399 .na
400 \fB\fBdedicated-cpu\fR\fR
401 .ad
402 .sp .6
403 .RS 4n
404 \fBncpus\fR, \fBimportance\fR
405 .RE

407 .sp
408 .ne 2
409 .na
410 \fB\fBcapped-memory\fR\fR
411 .ad
412 .sp .6
413 .RS 4n
414 \fBphysical\fR, \fBswap\fR, \fBblocked\fR
415 .RE

417 .sp
418 .ne 2
419 .na
420 \fB\fBcapped-cpu\fR\fR
421 .ad
422 .sp .6
423 .RS 4n
424 \fBncpus\fR
425 .RE

427 .sp
428 .LP
429 As for the property values which are paired with these names, they are either
430 simple, complex, or lists. The type allowed is property-specific. Simple values
431 are strings, optionally enclosed within quotation marks. Complex values have
432 the syntax:
433 .sp
434 .in +2
435 .nf
436 (<\fIname\fR>=<\fIvalue\fR>,<\fIname\fR>=<\fIvalue\fR>,...)
437 .fi
438 .in -2
439 .sp

441 .sp
442 .LP
443 where each <\fIvalue\fR> is simple, and the <\fIname\fR> strings are unique
444 within a given property. Lists have the syntax:
445 .sp
446 .in +2
447 .nf
448 [<\fIvalue\fR>,...]
449 .fi
450 .in -2
451 .sp

453 .sp
454 .LP
455 where each <\fIvalue\fR> is either simple or complex. A list of a single value
456 (either simple or complex) is equivalent to specifying that value without the
457 list syntax. That is, "foo" is equivalent to "[foo]". A list can be empty

```

```

458 (denoted by "[ ]").
459 .sp
460 .LP
461 In interpreting property values, \fBzonecfg\fR accepts regular expressions as
462 specified in \fBfnmatch\fR(5). See \fBEXAMPLES\fR.
463 .sp
464 .LP
465 The property types are described as follows:
466 .sp
467 .ne 2
468 .na
469 \fBglobal: \fBzonename\fR\fR
470 .ad
471 .sp .6
472 .RS 4n
473 The name of the zone.
474 .RE

476 .sp
477 .ne 2
478 .na
479 \fBglobal: \fBzonepath\fR\fR
480 .ad
481 .sp .6
482 .RS 4n
483 Path to zone's file system.
484 .RE

486 .sp
487 .ne 2
488 .na
489 \fBglobal: \fBautoboot\fR\fR
490 .ad
491 .sp .6
492 .RS 4n
493 Boolean indicating that a zone should be booted automatically at system boot.
494 Note that if the zones service is disabled, the zone will not autoboot,
495 regardless of the setting of this property. You enable the zones service with a
496 \fBsvcadm\fR command, such as:
497 .sp
498 .in +2
499 .nf
500 # \fBsvcadm enable svc:/system/zones:default\fR
501 .fi
502 .in -2
503 .sp

505 Replace \fBenable\fR with \fBdisable\fR to disable the zones service. See
506 \fBsvcadm\fR(1M).
507 .RE

509 .sp
510 .ne 2
511 .na
512 \fBglobal: \fBbootargs\fR\fR
513 .ad
514 .sp .6
515 .RS 4n
516 Arguments (options) to be passed to the zone bootup, unless options are
517 supplied to the "\fBzoneadm boot\fR" command, in which case those take
518 precedence. The valid arguments are described in \fBzoneadm\fR(1M).
519 .RE

521 .sp
522 .ne 2
523 .na

```

```

524 \fBglobal: \fBpool\fR\fR
525 .ad
526 .sp .6
527 .RS 4n
528 Name of the resource pool that this zone must be bound to when booted. This
529 property is incompatible with the \fBdedicated-cpu\fR resource.
530 .RE

532 .sp
533 .ne 2
534 .na
535 \fBglobal: \fBlimitpriv\fR\fR
536 .ad
537 .sp .6
538 .RS 4n
539 The maximum set of privileges any process in this zone can obtain. The property
540 should consist of a comma-separated privilege set specification as described in
541 \fBpriv_str_to_set\fR(3C). Privileges can be excluded from the resulting set by
542 preceding their names with a dash (-) or an exclamation point (!). The special
543 privilege string "zone" is not supported in this context. If the special string
544 "default" occurs as the first token in the property, it expands into a safe set
545 of privileges that preserve the resource and security isolation described in
546 \fBzones\fR(5). A missing or empty property is equivalent to this same set of
547 safe privileges.
548 .sp
549 The system administrator must take extreme care when configuring privileges for
550 a zone. Some privileges cannot be excluded through this mechanism as they are
551 required in order to boot a zone. In addition, there are certain privileges
552 which cannot be given to a zone as doing so would allow processes inside a zone
553 to unduly affect processes in other zones. \fBzoneadm\fR(1M) indicates when an
554 invalid privilege has been added or removed from a zone's privilege set when an
555 attempt is made to either "boot" or "ready" the zone.
556 .sp
557 See \fBprivileges\fR(5) for a description of privileges. The command "\fBppriv
558 -l\fR" (see \fBppriv\fR(1)) produces a list of all Solaris privileges. You can
559 specify privileges as they are displayed by \fBppriv\fR. In
560 \fBprivileges\fR(5), privileges are listed in the form
561 PRIV_\fIprivilege_name\fR. For example, the privilege \fIisys_time\fR, as you
562 would specify it in this property, is listed in \fBprivileges\fR(5) as
563 \fBPRIV_SYS_TIME\fR.
564 .RE

566 .sp
567 .ne 2
568 .na
569 \fBglobal: \fBbrand\fR\fR
570 .ad
571 .sp .6
572 .RS 4n
573 The zone's brand type.
574 .RE

576 .sp
577 .ne 2
578 .na
579 \fBglobal: \fBip-type\fR\fR
580 .ad
581 .sp .6
582 .RS 4n
583 A zone can either share the IP instance with the global zone, which is the
584 default, or have its own exclusive instance of IP.
585 .sp
586 This property takes the values \fBshared\fR and \fBexclusive\fR.
587 .RE

589 .sp

```

```

590 .ne 2
591 .na
592 \fBglobal: \fBhostid\fR\fR
593 .ad
594 .sp .6
595 .RS 4n
596 A zone can emulate a 32-bit host identifier to ease system consolidation. A
597 zone's \fBhostid\fR property is empty by default, meaning that the zone does
598 not emulate a host identifier. Zone host identifiers must be hexadecimal values
599 between 0 and FFFFFFFE. A \fB0x\fR or \fB0X\fR prefix is optional. Both
600 uppercase and lowercase hexadecimal digits are acceptable.
601 .RE

603 .sp
604 .ne 2
605 .na
606 \fBfbfs: dir, special, raw, type, options\fR
607 .ad
608 .sp .6
609 .RS 4n
610 Values needed to determine how, where, and so forth to mount file systems. See
611 \fBmount\fR(1M), \fBmount\fR(2), \fBfsck\fR(1M), and \fBfstab\fR(4).
612 .RE

614 .sp
615 .ne 2
616 .na
617 \fBfbnet: address, physical, defrouter\fR
618 .ad
619 .sp .6
620 .RS 4n
621 The network address and physical interface name of the network interface. The
622 network address is one of:
623 .RS +4
624 .TP
625 .ie t \(\bu
626 .el o
627 a valid IPv4 address, optionally followed by "\fB/\fR" and a prefix length;
628 .RE
629 .RS +4
630 .TP
631 .ie t \(\bu
632 .el o
633 a valid IPv6 address, which must be followed by "\fB/\fR" and a prefix length;
634 .RE
635 .RS +4
636 .TP
637 .ie t \(\bu
638 .el o
639 a host name which resolves to an IPv4 address.
640 .RE
641 Note that host names that resolve to IPv6 addresses are not supported.
642 .sp
643 The physical interface name is the network interface name.
644 .sp
645 The default router is specified similarly to the network address except that it
646 must not be followed by a \fB/\fR (slash) and a network prefix length.
647 .sp
648 A zone can be configured to be either exclusive-IP or shared-IP. For a
649 shared-IP zone, you must set both the physical and address properties; setting
650 the default router is optional. The interface specified in the physical
651 property must be plumbed in the global zone prior to booting the non-global
652 zone. However, if the interface is not used by the global zone, it should be
653 configured \fBdown\fR in the global zone, and the default router for the
654 interface should be specified here.
655 .sp

```

656 For an exclusive-IP zone, the physical property must be set and the address and  
 657 default router properties cannot be set.  
 658 .RE

660 .sp  
 661 .ne 2  
 662 .na  
 663 \fB\fBdevice\fR: match\fR  
 664 .ad  
 665 .sp .6  
 666 .RS 4n  
 667 Device name to match.  
 668 .RE

670 .sp  
 671 .ne 2  
 672 .na  
 673 \fB\fBrcctl\fR: name, value\fR  
 674 .ad  
 675 .sp .6  
 676 .RS 4n  
 677 The name and \fIpriv\fR/\fIlimit\fR/\fIaction\fR triple of a resource control.  
 678 See \fBprctl\fR(1) and \fBrcctladm\fR(1M). The preferred way to set rctl values  
 679 is to use the global property name associated with a specific rctl.  
 680 .RE

682 .sp  
 683 .ne 2  
 684 .na  
 685 \fB\fBattr\fR: name, type, value\fR  
 686 .ad  
 687 .sp .6  
 688 .RS 4n  
 689 The name, type and value of a generic attribute. The \fBtype\fR must be one of  
 690 \fBint\fR, \fBuint\fR, \fBboolean\fR or \fBstring\fR, and the value must be of  
 691 that type. \fBuint\fR means unsigned , that is, a non-negative integer.  
 692 .RE

694 .sp  
 695 .ne 2  
 696 .na  
 697 \fB\fBdataset\fR: name\fR  
 698 .ad  
 699 .sp .6  
 700 .RS 4n  
 701 The name of a \fBZFS\fR dataset to be accessed from within the zone. See  
 702 \fBzfs\fR(1M).  
 703 .RE

705 .sp  
 706 .ne 2  
 707 .na  
 708 \fBglobal: \fBcpu-shares\fR\fR  
 709 .ad  
 710 .sp .6  
 711 .RS 4n  
 712 The number of Fair Share Scheduler (FSS) shares to allocate to this zone. This  
 713 property is incompatible with the \fBdedicated-cpu\fR resource. This property  
 714 is the preferred way to set the \fBzone.cpu-shares\fR rctl.  
 715 .RE

717 .sp  
 718 .ne 2  
 719 .na  
 720 \fBglobal: \fBmax-lwps\fR\fR  
 721 .ad

722 .sp .6  
 723 .RS 4n  
 724 The maximum number of LWPs simultaneously available to this zone. This property  
 725 is the preferred way to set the \fBzone.max-lwps\fR rctl.  
 726 .RE

728 .sp  
 729 .ne 2  
 730 .na  
 731 \fBglobal: \fBmax-msg-ids\fR\fR  
 732 .ad  
 733 .sp .6  
 734 .RS 4n  
 735 The maximum number of message queue IDs allowed for this zone. This property is  
 736 the preferred way to set the \fBzone.max-msg-ids\fR rctl.  
 737 .RE

739 .sp  
 740 .ne 2  
 741 .na  
 742 \fBglobal: \fBmax-sem-ids\fR\fR  
 743 .ad  
 744 .sp .6  
 745 .RS 4n  
 746 The maximum number of semaphore IDs allowed for this zone. This property is the  
 747 preferred way to set the \fBzone.max-sem-ids\fR rctl.  
 748 .RE

750 .sp  
 751 .ne 2  
 752 .na  
 753 \fBglobal: \fBmax-shm-ids\fR\fR  
 754 .ad  
 755 .sp .6  
 756 .RS 4n  
 757 The maximum number of shared memory IDs allowed for this zone. This property is  
 758 the preferred way to set the \fBzone.max-shm-ids\fR rctl.  
 759 .RE

761 .sp  
 762 .ne 2  
 763 .na  
 764 \fBglobal: \fBmax-shm-memory\fR\fR  
 765 .ad  
 766 .sp .6  
 767 .RS 4n  
 768 The maximum amount of shared memory allowed for this zone. This property is the  
 769 preferred way to set the \fBzone.max-shm-memory\fR rctl. A scale (K, M, G, T)  
 770 can be applied to the value for this number (for example, 1M is one megabyte).  
 771 .RE

773 .sp  
 774 .ne 2  
 775 .na  
 776 \fBglobal: \fBsched-class\fR\fR  
 777 .ad  
 778 .sp .6  
 779 .RS 4n  
 780 Specifies the scheduling class used for processes running in a zone. When this  
 781 property is not specified, the scheduling class is established as follows:  
 782 .RS +4  
 783 .TP  
 784 .ie t \(\bu  
 785 .el o  
 786 If the \fBcpu-shares\fR property or equivalent rctl is set, the scheduling  
 787 class FSS is used.

```

788 .RE
789 .RS +4
790 .TP
791 .ie t \(bu
792 .el o
793 If neither \fBcpu-shares\fR nor the equivalent rctl is set and the zone's pool
794 property references a pool that has a default scheduling class, that class is
795 used.
796 .RE
797 .RS +4
798 .TP
799 .ie t \(bu
800 .el o
801 Under any other conditions, the system default scheduling class is used.
802 .RE
803 .RE

```

```

807 .sp
808 .ne 2
809 .na
810 \fB\fBdedicated-cpu\fR: ncpus, importance\fR
811 .ad
812 .sp .6
813 .RS 4n
814 The number of CPUs that should be assigned for this zone's exclusive use. The
815 zone will create a pool and processor set when it boots. See \fBpooladm\fR(1M)
816 and \fBpoolcfg\fR(1M) for more information on resource pools. The \fBncpu\fR
817 property can specify a single value or a range (for example, 1-4) of
818 processors. The \fBimportance\fR property is optional; if set, it will specify
819 the \fBpset.importance\fR value for use by \fBpoold\fR(1M). If this resource is
820 used, there must be enough free processors to allocate to this zone when it
821 boots or the zone will not boot. The processors assigned to this zone will not
822 be available for the use of the global zone or other zones. This resource is
823 incompatible with both the \fBpool\fR and \fBcpu-shares\fR properties. Only a
824 single instance of this resource can be added to the zone.
825 .RE

```

```

827 .sp
828 .ne 2
829 .na
830 \fB\fBcapped-memory\fR: physical, swap, locked\fR
831 .ad
832 .sp .6
833 .RS 4n
834 The caps on the memory that can be used by this zone. A scale (K, M, G, T) can
835 be applied to the value for each of these numbers (for example, 1M is one
836 megabyte). Each of these properties is optional but at least one property must
837 be set when adding this resource. Only a single instance of this resource can
838 be added to the zone. The \fBphysical\fR property sets the \fBmax-rss\fR for
839 this zone. This will be enforced by \fBrcapd\fR(1M) running in the global zone.
840 The \fBswap\fR property is the preferred way to set the \fBzone.max-swap\fR
841 rctl. The \fBblocked\fR property is the preferred way to set the
842 \fBzone.max-locked-memory\fR rctl.
843 .RE

```

```

845 .sp
846 .ne 2
847 .na
848 \fB\fBcapped-cpu\fR: ncpus\fR
849 .ad
850 .sp .6
851 .RS 4n
852 Sets a limit on the amount of CPU time that can be used by a zone. The unit
853 used translates to the percentage of a single CPU that can be used by all user

```

```

854 threads in a zone, expressed as a fraction (for example, \fB&.75\fR) or a
855 mixed number (whole number and fraction, for example, \fB1.25\fR). An
856 \fBncpu\fR value of \fB1\fR means 100% of a CPU, a value of \fB1.25\fR means
857 125%, \fB&.75\fR mean 75%, and so forth. When projects within a capped zone
858 have their own caps, the minimum value takes precedence.
859 .sp
860 The \fBcapped-cpu\fR property is an alias for \fBzone.cpu-cap\fR resource
861 control and is related to the \fBzone.cpu-cap\fR resource control. See
862 \fBresource_controls\fR(5).
863 .RE

```

```

865 .sp
866 .ne 2
867 .mk
868 .na
869 \fBglobal: \fBfs-allowed\fR
870 .ad
871 .sp .6
872 .RS 4n
873 A comma-separated list of additional filesystems that may be mounted within
874 the zone; for example "ufs,pcfs". By default, only hfs(7fs) and network
875 filesystems can be mounted. If the first entry in the list is "-" then
876 after "-" then only those filesystems can be mounted.

```

```

878 This property does not apply to filesystems mounted into the zone via "add fs"
879 or "add dataset".

```

```

881 WARNING: allowing filesystem mounts other than the default may allow the zone
882 administrator to compromise the system with a malicious filesystem image, and
883 is not supported.
884 .RE

```

```

886 .sp
887 .LP
888 The following table summarizes resources, property-names, and types:
889 .sp
890 .in +2
891 .nf
892 resource          property-name  type
893 (global)          zonename      simple
894 (global)          zonepath      simple
895 (global)          autoboot      simple
896 (global)          bootargs      simple
897 (global)          pool          simple
898 (global)          limitpriv     simple
899 (global)          brand         simple
900 (global)          ip-type       simple
901 (global)          hostid        simple
902 (global)          cpu-shares    simple
903 (global)          max-lwps      simple
904 (global)          max-msg-ids   simple
905 (global)          max-sem-ids   simple
906 (global)          max-shm-ids   simple
907 (global)          max-shm-memory simple
908 (global)          scheduling-class simple
909 fs                dir           simple
910                  special       simple
911                  raw           simple
912                  type          simple
913                  options       list of simple
914 net                address       simple
915                  physical      simple
916 device            match         simple
917 rctl              name          simple
918                  value         list of complex

```

```

919 attr          name          simple
920              type          simple
921              value         simple
922 dataset        name          simple
923 dedicated-cpu ncpus          simple or range
924              importance    simple

926 capped-memory physical    simple with scale
927              swap          simple with scale
928              locked        simple with scale

930 capped-cpu    ncpus          simple
931 .fi
932 .in -2
933 .sp

935 .sp
936 .LP
937 To further specify things, the breakdown of the complex property "value" of the
938 "rctl" resource type, it consists of three name/value pairs, the names being
939 "priv", "limit" and "action", each of which takes a simple value. The "name"
940 property of an "attr" resource is syntactically restricted in a fashion similar
941 but not identical to zone names: it must begin with an alphanumeric, and can
942 contain alphanumerics plus the hyphen (\fB-\fR), underscore (\fB_\fR), and dot
943 (\fB&.\fR) characters. Attribute names beginning with "zone" are reserved for
944 use by the system. Finally, the "autoboot" global property must have a value of
945 "true" or "false".
946 .SS "Using Kernel Statistics to Monitor CPU Caps"
947 .sp
948 .LP
949 Using the kernel statistics (\fBkstat\fR(3KSTAT)) module \fBcaps\fR, the system
950 maintains information for all capped projects and zones. You can access this
951 information by reading kernel statistics (\fBkstat\fR(3KSTAT)), specifying
952 \fBcaps\fR as the \fBkstat\fR module name. The following command displays
953 kernel statistics for all active CPU caps:
954 .sp
955 .in +2
956 .nf
957 # \fBkstat caps::'/cpucaps/'\fR
958 .fi
959 .in -2
960 .sp

962 .sp
963 .LP
964 A \fBkstat\fR(1M) command running in a zone displays only CPU caps relevant for
965 that zone and for projects in that zone. See \fBEXAMPLES\fR.
966 .sp
967 .LP
968 The following are cap-related arguments for use with \fBkstat\fR(1M):
969 .sp
970 .ne 2
971 .na
972 \fB\fBcaps\fR\fR
973 .ad
974 .sp .6
975 .RS 4n
976 The \fBkstat\fR module.
977 .RE

979 .sp
980 .ne 2
981 .na
982 \fB\fBproject_caps\fR or \fB\fBzone_caps\fR\fR
983 .ad
984 .sp .6

```

```

985 .RS 4n
986 \fBkstat\fR class, for use with the \fBkstat\fR \fB-c\fR option.
987 .RE

989 .sp
990 .ne 2
991 .na
992 \fB\fBcpu_caps_project\fR or \fB\fBcpu_caps_zone\fR\fR
993 .ad
994 .sp .6
995 .RS 4n
996 \fBkstat\fR name, for use with the \fBkstat\fR \fB-n\fR option. \fB\fR is the
997 project or zone identifier.
998 .RE

1000 .sp
1001 .LP
1002 The following fields are displayed in response to a \fBkstat\fR(1M) command
1003 requesting statistics for all CPU caps.
1004 .sp
1005 .ne 2
1006 .na
1007 \fB\fBmodule\fR\fR
1008 .ad
1009 .sp .6
1010 .RS 4n
1011 In this usage of \fBkstat\fR, this field will have the value \fBcaps\fR.
1012 .RE

1014 .sp
1015 .ne 2
1016 .na
1017 \fB\fBname\fR\fR
1018 .ad
1019 .sp .6
1020 .RS 4n
1021 As described above, \fBcpu_caps_project\fR or
1022 \fBcpu_caps_zone\fR
1023 .RE

1025 .sp
1026 .ne 2
1027 .na
1028 \fB\fBabove_sec\fR\fR
1029 .ad
1030 .sp .6
1031 .RS 4n
1032 Total time, in seconds, spent above the cap.
1033 .RE

1035 .sp
1036 .ne 2
1037 .na
1038 \fB\fBbelow_sec\fR\fR
1039 .ad
1040 .sp .6
1041 .RS 4n
1042 Total time, in seconds, spent below the cap.
1043 .RE

1045 .sp
1046 .ne 2
1047 .na
1048 \fB\fBmaxusage\fR\fR
1049 .ad
1050 .sp .6

```

```

1051 .RS 4n
1052 Maximum observed CPU usage.
1053 .RE

1055 .sp
1056 .ne 2
1057 .na
1058 \fB\fBwait\fR\fR
1059 .ad
1060 .sp .6
1061 .RS 4n
1062 Number of threads on cap wait queue.
1063 .RE

1065 .sp
1066 .ne 2
1067 .na
1068 \fB\fBusage\fR\fR
1069 .ad
1070 .sp .6
1071 .RS 4n
1072 Current aggregated CPU usage for all threads belonging to a capped project or
1073 zone, in terms of a percentage of a single CPU.
1074 .RE

1076 .sp
1077 .ne 2
1078 .na
1079 \fB\fBvalue\fR\fR
1080 .ad
1081 .sp .6
1082 .RS 4n
1083 The cap value, in terms of a percentage of a single CPU.
1084 .RE

1086 .sp
1087 .ne 2
1088 .na
1089 \fB\fBzonename\fR\fR
1090 .ad
1091 .sp .6
1092 .RS 4n
1093 Name of the zone for which statistics are displayed.
1094 .RE

1096 .sp
1097 .LP
1098 See \fBEXAMPLES\fR for sample output from a \fBkstat\fR command.
1099 .SH OPTIONS
1100 .sp
1101 .LP
1102 The following options are supported:
1103 .sp
1104 .ne 2
1105 .na
1106 \fB\fB-f\fR \fIcommand_file\fR\fR
1107 .ad
1108 .sp .6
1109 .RS 4n
1110 Specify the name of \fBzonecfg\fR command file. \fIcommand_file\fR is a text
1111 file of \fBzonecfg\fR subcommands, one per line.
1112 .RE

1114 .sp
1115 .ne 2
1116 .na

```

```

1117 \fB\fB-z\fR \fIzonename\fR\fR
1118 .ad
1119 .sp .6
1120 .RS 4n
1121 Specify the name of a zone. Zone names are case sensitive. Zone names must
1122 begin with an alphanumeric character and can contain alphanumeric characters,
1123 the underscore (\fB_\fR) the hyphen (\fB-\fR), and the dot (\fB&.\fR). The
1124 name \fBglobal\fR and all names beginning with \fBSUNW\fR are reserved and
1125 cannot be used.
1126 .RE

1128 .SH SUBCOMMANDS
1129 .sp
1130 .LP
1131 You can use the \fBadd\fR and \fBselect\fR subcommands to select a specific
1132 resource, at which point the scope changes to that resource. The \fBend\fR and
1133 \fBcancel\fR subcommands are used to complete the resource specification, at
1134 which time the scope is reverted back to global. Certain subcommands, such as
1135 \fBadd\fR, \fBremove\fR and \fBset\fR, have different semantics in each scope.
1136 .sp
1137 .LP
1138 \fBzonecfg\fR supports a semicolon-separated list of subcommands. For example:
1139 .sp
1140 .in +2
1141 .nf
1142 # \fBzonecfg -z myzone "add net; set physical=myvnic; end"\fR
1143 .fi
1144 .in -2
1145 .sp

1147 .sp
1148 .LP
1149 Subcommands which can result in destructive actions or loss of work have an
1150 \fB-F\fR option to force the action. If input is from a terminal device, the
1151 user is prompted when appropriate if such a command is given without the
1152 \fB-F\fR option otherwise, if such a command is given without the \fB-F\fR
1153 option, the action is disallowed, with a diagnostic message written to standard
1154 error.
1155 .sp
1156 .LP
1157 The following subcommands are supported:
1158 .sp
1159 .ne 2
1160 .na
1161 \fB\fBadd\fR \fIresource-type\fR (global scope)\fR
1162 .ad
1163 .br
1164 .na
1165 \fB\fBadd\fR \fIproperty-name property-value\fR (resource scope)\fR
1166 .ad
1167 .sp .6
1168 .RS 4n
1169 In the global scope, begin the specification for a given resource type. The
1170 scope is changed to that resource type.
1171 .sp
1172 In the resource scope, add a property of the given name with the given value.
1173 The syntax for property values varies with different property types. In
1174 general, it is a simple value or a list of simple values enclosed in square
1175 brackets, separated by commas (\fB[foo,bar,baz]\fR). See \fBPROPERTIES\fR.
1176 .RE

1178 .sp
1179 .ne 2
1180 .na
1181 \fB\fBcancel\fR\fR
1182 .ad

```



1183 .sp .6  
 1184 .RS 4n  
 1185 End the resource specification and reset scope to global. Abandons any  
 1186 partially specified resources. \fBcancel\fR is only applicable in the resource  
 1187 scope.  
 1188 .RE

1190 .sp  
 1191 .ne 2  
 1192 .na  
 1193 \fB\fBclear\fR \fIproperty-name\fR\fR  
 1194 .ad  
 1195 .sp .6  
 1196 .RS 4n  
 1197 Clear the value for the property.  
 1198 .RE

1200 .sp  
 1201 .ne 2  
 1202 .na  
 1203 \fB\fBcommit\fR\fR  
 1204 .ad  
 1205 .sp .6  
 1206 .RS 4n  
 1207 Commit the current configuration from memory to stable storage. The  
 1208 configuration must be committed to be used by \fBzoneadm\fR. Until the  
 1209 in-memory configuration is committed, you can remove changes with the  
 1210 \fBbrevert\fR subcommand. The \fBcommit\fR operation is attempted automatically  
 1211 upon completion of a \fBzonecfg\fR session. Since a configuration must be  
 1212 correct to be committed, this operation automatically does a verify.  
 1213 .RE

1215 .sp  
 1216 .ne 2  
 1217 .na  
 1218 \fB\fBcreate [\fR\fB-F\fR\fB] [\fR \fB-a\fR \fIpath\fR | \fB-b\fR \fB]\fR  
 1219 \fB-t\fR \fItemplate\fR\fB]\fR  
 1220 .ad  
 1221 .sp .6  
 1222 .RS 4n  
 1223 Create an in-memory configuration for the specified zone. Use \fBcreate\fR to  
 1224 begin to configure a new zone. See \fBcommit\fR for saving this to stable  
 1225 storage.  
 1226 .sp  
 1227 If you are overwriting an existing configuration, specify the \fB-F\fR option  
 1228 to force the action. Specify the \fB-t\fR \fItemplate\fR option to create a  
 1229 configuration identical to \fItemplate\fR, where \fItemplate\fR is the name of  
 1230 a configured zone.  
 1231 .sp  
 1232 Use the \fB-a\fR \fIpath\fR option to facilitate configuring a detached zone on  
 1233 a new host. The \fIpath\fR parameter is the zonepath location of a detached  
 1234 zone that has been moved on to this new host. Once the detached zone is  
 1235 configured, it should be installed using the "\fBzoneadm attach\fR" command  
 1236 (see \fBzoneadm\fR(1M)). All validation of the new zone happens during the  
 1237 \fBattach\fR process, not during zone configuration.  
 1238 .sp  
 1239 Use the \fB-b\fR option to create a blank configuration. Without arguments,  
 1240 \fBcreate\fR applies the Sun default settings.  
 1241 .RE

1243 .sp  
 1244 .ne 2  
 1245 .na  
 1246 \fB\fBdelete [\fR\fB-F\fR\fB]\fR  
 1247 .ad  
 1248 .sp .6

1249 .RS 4n  
 1250 Delete the specified configuration from memory and stable storage. This action  
 1251 is instantaneous, no commit is necessary. A deleted configuration cannot be  
 1252 reverted.  
 1253 .sp  
 1254 Specify the \fB-F\fR option to force the action.  
 1255 .RE

1257 .sp  
 1258 .ne 2  
 1259 .na  
 1260 \fB\fBend\fR  
 1261 .ad  
 1262 .sp .6  
 1263 .RS 4n  
 1264 End the resource specification. This subcommand is only applicable in the  
 1265 resource scope. \fBzonecfg\fR checks to make sure the current resource is  
 1266 completely specified. If so, it is added to the in-memory configuration (see  
 1267 \fBcommit\fR for saving this to stable storage) and the scope reverts to  
 1268 global. If the specification is incomplete, it issues an appropriate error  
 1269 message.  
 1270 .RE

1272 .sp  
 1273 .ne 2  
 1274 .na  
 1275 \fB\fBexport [\fR\fB-f\fR \fIoutput-file\fR\fB]\fR  
 1276 .ad  
 1277 .sp .6  
 1278 .RS 4n  
 1279 Print configuration to standard output. Use the \fB-f\fR option to print the  
 1280 configuration to \fIoutput-file\fR. This option produces output in a form  
 1281 suitable for use in a command file.  
 1282 .RE

1284 .sp  
 1285 .ne 2  
 1286 .na  
 1287 \fB\fBhelp [usage] [\fIsubcommand\fR] [syntax] [\fR\fIcommand-name\fR\fB]\fR  
 1288 .ad  
 1289 .sp .6  
 1290 .RS 4n  
 1291 Print general help or help about given topic.  
 1292 .RE

1294 .sp  
 1295 .ne 2  
 1296 .na  
 1297 \fB\fBinfo zonename | zonepath | autoboot | brand | pool | limitpriv\fR  
 1298 .ad  
 1299 .br  
 1300 .na  
 1301 \fB\fBinfo [\fR\fIresource-type\fR  
 1302 \fB[\fR\fIproperty-name\fR\fB=\fR\fIproperty-value\fR\fB\*]\fR  
 1303 .ad  
 1304 .sp .6  
 1305 .RS 4n  
 1306 Display information about the current configuration. If \fIresource-type\fR is  
 1307 specified, displays only information about resources of the relevant type. If  
 1308 any \fIproperty-name\fR value pairs are specified, displays only information  
 1309 about resources meeting the given criteria. In the resource scope, any  
 1310 arguments are ignored, and \fBinfo\fR displays information about the resource  
 1311 which is currently being added or modified.  
 1312 .RE

1314 .sp

```

1315 .ne 2
1316 .na
1317 \fB\fBremove\fR \fIresource-type\fR \fB{\fR\fIproperty-name\fR \fB=\fR\fIproperty
1318 -value\fR \fB}\fR(global scope)\fR
1319 .ad
1320 .sp .6
1321 .RS 4n
1322 In the global scope, removes the specified resource. The \fB[]\fR syntax means
1323 0 or more of whatever is inside the square braces. If you want only to remove a
1324 single instance of the resource, you must specify enough property name-value
1325 pairs for the resource to be uniquely identified. If no property name-value
1326 pairs are specified, all instances will be removed. If there is more than one
1327 pair is specified, a confirmation is required, unless you use the \fB-F\fR
1328 option.
1329 .RE

1331 .sp
1332 .ne 2
1333 .na
1334 \fB\fBselect\fR \fIresource-type\fR
1335 \fB{\fR\fIproperty-name\fR \fB=\fR\fIproperty-value\fR \fB}\fR\fR
1336 .ad
1337 .sp .6
1338 .RS 4n
1339 Select the resource of the given type which matches the given
1340 \fIproperty-name\fR \fIproperty-value\fR pair criteria, for modification. This
1341 subcommand is applicable only in the global scope. The scope is changed to that
1342 resource type. The \fB{}\fR syntax means 1 or more of whatever is inside the
1343 curly braces. You must specify enough \fIproperty -name property-value\fR pairs
1344 for the resource to be uniquely identified.
1345 .RE

1347 .sp
1348 .ne 2
1349 .na
1350 \fB\fBset\fR \fIproperty-name\fR \fB=\fR \fIproperty\fR \fB-\fR \fIvalue\fR \fR
1351 .ad
1352 .sp .6
1353 .RS 4n
1354 Set a given property name to the given value. Some properties (for example,
1355 \fBzonepath\fR and \fBzonename\fR) are global while others are
1356 resource-specific. This subcommand is applicable in both the global and
1357 resource scopes.
1358 .RE

1360 .sp
1361 .ne 2
1362 .na
1363 \fB\fBverify\fR \fR
1364 .ad
1365 .sp .6
1366 .RS 4n
1367 Verify the current configuration for correctness.
1368 .RS +4
1369 .TP
1370 .ie t \(\bu
1371 .el o
1372 All resources have all of their required properties specified.
1373 .RE
1374 .RS +4
1375 .TP
1376 .ie t \(\bu
1377 .el o
1378 A \fBzonepath\fR is specified.
1379 .RE
1380 .RE

```

```

1382 .sp
1383 .ne 2
1384 .na
1385 \fB\fBvert\fR \fB{\fR \fB-F\fR \fB}\fR \fR
1386 .ad
1387 .sp .6
1388 .RS 4n
1389 Revert the configuration back to the last committed state. The \fB-F\fR option
1390 can be used to force the action.
1391 .RE

1393 .sp
1394 .ne 2
1395 .na
1396 \fB\fBexit [\fR \fB-F\fR \fB]\fR \fR
1397 .ad
1398 .sp .6
1399 .RS 4n
1400 Exit the \fBzonecfg\fR session. A commit is automatically attempted if needed.
1401 You can also use an \fBEOF\fR character to exit \fBzonecfg\fR. The \fB-F\fR
1402 option can be used to force the action.
1403 .RE

1405 .SH EXAMPLES
1406 .LP
1407 \fBExample 1 \fRCreating the Environment for a New Zone
1408 .sp
1409 .LP
1410 In the following example, \fBzonecfg\fR creates the environment for a new zone.
1411 \fB/usr/local\fR is loopback mounted from the global zone into
1412 \fB/opt/local\fR. \fB/opt/sfw\fR is loopback mounted from the global zone,
1413 three logical network interfaces are added, and a limit on the number of
1414 fair-share scheduler (FSS) CPU shares for a zone is set using the \fBbrctl\fR
1415 resource type. The example also shows how to select a given resource for
1416 modification.

1418 .sp
1419 .in +2
1420 .nf
1421 example# \fBzonecfg -z myzone3\fR
1422 my-zone3: No such zone configured
1423 Use 'create' to begin configuring a new zone.
1424 zonecfg:myzone3> \fBcreate\fR
1425 zonecfg:myzone3> \fBset zonepath=/export/home/my-zone3\fR
1426 zonecfg:myzone3> \fBset autoboot=true\fR
1427 zonecfg:myzone3> \fBadd fs\fR
1428 zonecfg:myzone3:fs> \fBset dir=/usr/local\fR
1429 zonecfg:myzone3:fs> \fBset special=/opt/local\fR
1430 zonecfg:myzone3:fs> \fBset type=lofs\fR
1431 zonecfg:myzone3:fs> \fBadd options [ro,nodevices]\fR
1432 zonecfg:myzone3:fs> \fBend\fR
1433 zonecfg:myzone3> \fBadd fs\fR
1434 zonecfg:myzone3:fs> \fBset dir=/mnt\fR
1435 zonecfg:myzone3:fs> \fBset special=/dev/dsk/c0t0d0s7\fR
1436 zonecfg:myzone3:fs> \fBset raw=/dev/rdisk/c0t0d0s7\fR
1437 zonecfg:myzone3:fs> \fBset type=ufs\fR
1438 zonecfg:myzone3:fs> \fBend\fR
1439 zonecfg:myzone3> \fBadd net\fR
1440 zonecfg:myzone3:net> \fBset address=192.168.0.1/24\fR
1441 zonecfg:myzone3:net> \fBset physical=eri0\fR
1442 zonecfg:myzone3:net> \fBend\fR
1443 zonecfg:myzone3> \fBadd net\fR
1444 zonecfg:myzone3:net> \fBset address=192.168.1.2/24\fR
1445 zonecfg:myzone3:net> \fBset physical=eri0\fR
1446 zonecfg:myzone3:net> \fBend\fR

```

```

1447 zonecfg:myzone3> \fBadd net\fR
1448 zonecfg:myzone3:net> \fBset address=192.168.2.3/24\fR
1449 zonecfg:myzone3:net> \fBset physical=eri0\fR
1450 zonecfg:myzone3:net> \fBend\fR
1451 zonecfg:my-zone3> \fBset cpu-shares=5\fR
1452 zonecfg:my-zone3> \fBadd capped-memory\fR
1453 zonecfg:my-zone3:capped-memory> \fBset physical=50m\fR
1454 zonecfg:my-zone3:capped-memory> \fBset swap=100m\fR
1455 zonecfg:my-zone3:capped-memory> \fBend\fR
1456 zonecfg:myzone3> \fBexit\fR
1457 .fi
1458 .in -2
1459 .sp

1461 .LP
1462 \fBExample 2 \fRCreating a Non-Native Zone
1463 .sp
1464 .LP
1465 The following example creates a new Linux zone:

1467 .sp
1468 .in +2
1469 .nf
1470 example# \fBzonecfg -z lxzone\fR
1471 lxzone: No such zone configured
1472 Use 'create' to begin configuring a new zone
1473 zonecfg:lxzone> \fBcreate -t SUNWlx\fR
1474 zonecfg:lxzone> \fBset zonepath=/export/zones/lxzone\fR
1475 zonecfg:lxzone> \fBset autoboot=true\fR
1476 zonecfg:lxzone> \fBexit\fR
1477 .fi
1478 .in -2
1479 .sp

1481 .LP
1482 \fBExample 3 \fRCreating an Exclusive-IP Zone
1483 .sp
1484 .LP
1485 The following example creates a zone that is granted exclusive access to
1486 \fBbge1\fR and \fBbge33000\fR and that is isolated at the IP layer from the
1487 other zones configured on the system.

1489 .sp
1490 .LP
1491 The IP addresses and routing is configured inside the new zone using
1492 \fBsysidtool\fR(1M).

1494 .sp
1495 .in +2
1496 .nf
1497 example# \fBzonecfg -z excl\fR
1498 excl: No such zone configured
1499 Use 'create' to begin configuring a new zone
1500 zonecfg:excl> \fBcreate\fR
1501 zonecfg:excl> \fBset zonepath=/export/zones/excl\fR
1502 zonecfg:excl> \fBset ip-type=exclusive\fR
1503 zonecfg:excl> \fBadd net\fR
1504 zonecfg:excl:net> \fBset physical=bge1\fR
1505 zonecfg:excl:net> \fBend\fR
1506 zonecfg:excl> \fBadd net\fR
1507 zonecfg:excl:net> \fBset physical=bge33000\fR
1508 zonecfg:excl:net> \fBend\fR
1509 zonecfg:excl> \fBexit\fR
1510 .fi
1511 .in -2
1512 .sp

```

```

1514 .LP
1515 \fBExample 4 \fRAssociating a Zone with a Resource Pool
1516 .sp
1517 .LP
1518 The following example shows how to associate an existing zone with an existing
1519 resource pool:

1521 .sp
1522 .in +2
1523 .nf
1524 example# \fBzonecfg -z myzone\fR
1525 zonecfg:myzone> \fBset pool=mypool\fR
1526 zonecfg:myzone> \fBexit\fR
1527 .fi
1528 .in -2
1529 .sp

1531 .sp
1532 .LP
1533 For more information about resource pools, see \fBpooladm\fR(1M) and
1534 \fBpoolcfg\fR(1M).

1536 .LP
1537 \fBExample 5 \fRChanging the Name of a Zone
1538 .sp
1539 .LP
1540 The following example shows how to change the name of an existing zone:

1542 .sp
1543 .in +2
1544 .nf
1545 example# \fBzonecfg -z myzone\fR
1546 zonecfg:myzone> \fBset zonename=myzone2\fR
1547 zonecfg:myzone2> \fBexit\fR
1548 .fi
1549 .in -2
1550 .sp

1552 .LP
1553 \fBExample 6 \fRChanging the Privilege Set of a Zone
1554 .sp
1555 .LP
1556 The following example shows how to change the set of privileges an existing
1557 zone's processes will be limited to the next time the zone is booted. In this
1558 particular case, the privilege set will be the standard safe set of privileges
1559 a zone normally has along with the privilege to change the system date and
1560 time:

1562 .sp
1563 .in +2
1564 .nf
1565 example# \fBzonecfg -z myzone\fR
1566 zonecfg:myzone> \fBset limitpriv="default,sys_time"\fR
1567 zonecfg:myzone2> \fBexit\fR
1568 .fi
1569 .in -2
1570 .sp

1572 .LP
1573 \fBExample 7 \fRSetting the \fBzone.cpu-shares\fR Property for the Global Zone
1574 .sp
1575 .LP
1576 The following command sets the \fBzone.cpu-shares\fR property for the global
1577 zone:

```

```

1579 .sp
1580 .in +2
1581 .nf
1582 example# \fBzonecfg -z global\fR
1583 zonecfg:global> \fBset cpu-shares=5\fR
1584 zonecfg:global> \fBexit\fR
1585 .fi
1586 .in -2
1587 .sp

1589 .LP
1590 \fBExample 8 \fRUsing Pattern Matching
1591 .sp
1592 .LP
1593 The following commands illustrate \fBzonecfg\fR support for pattern matching.
1594 In the zone \fBflexlm\fR, enter:

1596 .sp
1597 .in +2
1598 .nf
1599 zonecfg:flexlm> \fBadd device\fR
1600 zonecfg:flexlm:device> \fBset match="/dev/cua/a00[2-5]"\fR
1601 zonecfg:flexlm:device> \fBend\fR
1602 .fi
1603 .in -2
1604 .sp

1606 .sp
1607 .LP
1608 In the global zone, enter:

1610 .sp
1611 .in +2
1612 .nf
1613 global# \fBls /dev/cua\fR
1614 a      a000 a001 a002 a003 a004 a005 a006 a007 b
1615 .fi
1616 .in -2
1617 .sp

1619 .sp
1620 .LP
1621 In the zone \fBflexlm\fR, enter:

1623 .sp
1624 .in +2
1625 .nf
1626 flexlm# \fBls /dev/cua\fR
1627 a002 a003 a004 a005
1628 .fi
1629 .in -2
1630 .sp

1632 .LP
1633 \fBExample 9 \fRSetting a Cap for a Zone to Three CPUs
1634 .sp
1635 .LP
1636 The following sequence uses the \fBzonecfg\fR command to set the CPU cap for a
1637 zone to three CPUs.

1639 .sp
1640 .in +2
1641 .nf
1642 zonecfg:myzone> \fBadd capped-cpu\fR
1643 zonecfg:myzone>capped-cpu> \fBset ncpus=3\fR
1644 zonecfg:myzone>capped-cpu>capped-cpu> \fBend\fR

```

```

1645 .fi
1646 .in -2
1647 .sp

1649 .sp
1650 .LP
1651 The preceding sequence, which uses the capped-cpu property, is equivalent to
1652 the following sequence, which makes use of the \fBzone.cpu-cap\fR resource
1653 control.

1655 .sp
1656 .in +2
1657 .nf
1658 zonecfg:myzone> \fBadd rctl\fR
1659 zonecfg:myzone:rctl> \fBset name=zone.cpu-cap\fR
1660 zonecfg:myzone:rctl> \fBadd value (priv=privileged,limit=300,action=none)\fR
1661 zonecfg:myzone:rctl> \fBend\fR
1662 .fi
1663 .in -2
1664 .sp

1666 .LP
1667 \fBExample 10 \fRUsing \fBkstat\fR to Monitor CPU Caps
1668 .sp
1669 .LP
1670 The following command displays information about all CPU caps.

1672 .sp
1673 .in +2
1674 .nf
1675 # \fBkstat -n /cpucaps/\fR
1676 module: caps                               instance: 0
1677 name:    cpucaps_project_0                 class:    project_caps
1678         above_sec                           0
1679         below_sec                            2157
1680         crtime                               821.048183159
1681         maxusage                             2
1682         nwait                                0
1683         snaptime                             235885.637253027
1684         usage                                0
1685         value                               18446743151372347932
1686         zonename                            global

1688 module: caps                               instance: 0
1689 name:    cpucaps_project_1                 class:    project_caps
1690         above_sec                           0
1691         below_sec                            0
1692         crtime                               225339.192787265
1693         maxusage                             5
1694         nwait                                0
1695         snaptime                             235885.637591677
1696         usage                                5
1697         value                               18446743151372347932
1698         zonename                            global

1700 module: caps                               instance: 0
1701 name:    cpucaps_project_201                class:    project_caps
1702         above_sec                           0
1703         below_sec                            235105
1704         crtime                               780.37961782
1705         maxusage                             100
1706         nwait                                0
1707         snaptime                             235885.637789687
1708         usage                                43
1709         value                               100
1710         zonename                            global

```

```

1712 module: caps                instance: 0
1713 name:    cpucaps_project_202 class:    project_caps
1714         above_sec             0
1715         below_sec            235094
1716         crtime               791.72983782
1717         maxusage             100
1718         nwait                 0
1719         snaptime             235885.637967512
1720         usage                 48
1721         value                 100
1722         zonename             global

1724 module: caps                instance: 0
1725 name:    cpucaps_project_203 class:    project_caps
1726         above_sec             0
1727         below_sec            235034
1728         crtime               852.104401481
1729         maxusage             75
1730         nwait                 0
1731         snaptime             235885.638144304
1732         usage                 47
1733         value                 100
1734         zonename             global

1736 module: caps                instance: 0
1737 name:    cpucaps_project_86710 class:    project_caps
1738         above_sec             22
1739         below_sec            235166
1740         crtime               698.441717859
1741         maxusage             101
1742         nwait                 0
1743         snaptime             235885.638319871
1744         usage                 54
1745         value                 100
1746         zonename             global

1748 module: caps                instance: 0
1749 name:    cpucaps_zone_0      class:    zone_caps
1750         above_sec             100733
1751         below_sec            134332
1752         crtime               821.048177123
1753         maxusage             207
1754         nwait                 2
1755         snaptime             235885.638497731
1756         usage                 199
1757         value                 200
1758         zonename             global

1760 module: caps                instance: 1
1761 name:    cpucaps_project_0   class:    project_caps
1762         above_sec             0
1763         below_sec            0
1764         crtime               225360.256448422
1765         maxusage             7
1766         nwait                 0
1767         snaptime             235885.638714404
1768         usage                 7
1769         value                 18446743151372347932
1770         zonename             test_001

1772 module: caps                instance: 1
1773 name:    cpucaps_zone_1     class:    zone_caps
1774         above_sec             2
1775         below_sec            10524
1776         crtime               225360.256440278

```

```

1777         maxusage             106
1778         nwait                 0
1779         snaptime             235885.638896443
1780         usage                 7
1781         value                 100
1782         zonename             test_001
1783         .fi
1784         .in -2
1785         .sp

1787 .LP
1788 \fBExample 11 \fRDisplaying CPU Caps for a Specific Zone or Project
1789 .sp
1790 .LP
1791 Using the \fBkstat\fR \fB-c\fR and \fB-i\fR options, you can display CPU caps
1792 for a specific zone or project, as below. The first command produces a display
1793 for a specific project, the second for the same project within zone 1.

1795 .sp
1796 .in +2
1797 .nf
1798 # \fBkstat -c project_caps\fR

1800 # \fBkstat -c project_caps -i 1\fR
1801 .fi
1802 .in -2
1803 .sp

1805 .SH EXIT STATUS
1806 .sp
1807 .LP
1808 The following exit values are returned:
1809 .sp
1810 .ne 2
1811 .na
1812 \fB0\fR
1813 .ad
1814 .sp .6
1815 .RS 4n
1816 Successful completion.
1817 .RE

1819 .sp
1820 .ne 2
1821 .na
1822 \fB1\fR
1823 .ad
1824 .sp .6
1825 .RS 4n
1826 An error occurred.
1827 .RE

1829 .sp
1830 .ne 2
1831 .na
1832 \fB2\fR
1833 .ad
1834 .sp .6
1835 .RS 4n
1836 Invalid usage.
1837 .RE

1839 .SH ATTRIBUTES
1840 .sp
1841 .LP
1842 See \fBAttributes\fR(5) for descriptions of the following attributes:

```

```
1843 .sp
1845 .sp
1846 .TS
1847 box;
1848 c | c
1849 l | l .
1850 ATTRIBUTE TYPE ATTRIBUTE VALUE
1851 -
1852 Interface Stability Volatile
1853 .TE

1855 .SH SEE ALSO
1856 .sp
1857 .LP
1858 \fBppriv\fR(1), \fBprctl\fR(1), \fBzlogin\fR(1), \fBkstat\fR(1M),
1859 \fBmount\fR(1M), \fBpooladm\fR(1M), \fBpoolcfg\fR(1M), \fBpoold\fR(1M),
1860 \fBrcapd\fR(1M), \fBrcctladm\fR(1M), \fBsvcadm\fR(1M), \fBsysidtool\fR(1M),
1861 \fBzfs\fR(1M), \fBzoneadm\fR(1M), \fBpriv_str_to_set\fR(3C),
1862 \fBkstat\fR(3KSTAT), \fBvfstab\fR(4), \fBattributes\fR(5), \fBbrands\fR(5),
1863 \fBfmatch\fR(5), \fBlx\fR(5), \fBprivileges\fR(5), \fBresource_controls\fR(5),
1864 \fBzones\fR(5)
1865 .sp
1866 .LP
1867 \fISystem Administration Guide: Solaris Containers-Resource Management, and
1868 Solaris Zones\fR
1869 .SH NOTES
1870 .sp
1871 .LP
1872 All character data used by \fBzonecfg\fR must be in US-ASCII encoding.
```

\*\*\*\*\*

5997 Thu Jul 17 00:50:40 2014

new/usr/src/man/man2/pipe.2

manpage lint.

\*\*\*\*\*

```

1 \" te
2.\" Copyright (c) 2002, Sun Microsystems, Inc. All Rights Reserved.
3.\" Copyright 1989 AT&T
4.\" Portions Copyright (c) 2001, the Institute of Electrical and Electronics Eng
5.\" Portions Copyright (c) 2013, OmniTI Computer Consulting, Inc. All Rights Res
6.\" Sun Microsystems, Inc. gratefully acknowledges The Open Group for permission
7.\" http://www.opengroup.org/bookstore/.
8.\" The Institute of Electrical and Electronics Engineers and The Open Group, ha
9.\" This notice shall appear on any product containing this material.
10.\" The contents of this file are subject to the terms of the Common Development
11.\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
12.\" When distributing Covered Code, include this CDDL HEADER in each file and in
13.TH PIPE 2 "Apr 19, 2013"
14.SH NAME
15 pipe \- create an interprocess channel
16.SH SYNOPSIS
17.LP
18.nf
19#include <unistd.h>

21 \fBint\fR \fBpipe\fR(\fBint\fR \fIfildes\fR[2]);

23 \fBint\fR \fBpipe2\fR(\fBint\fR \fIfildes\fR[2], \fBint\fR \fIflags\fR);
24.fi

26.SH DESCRIPTION
27.sp
28.LP
29 The \fBpipe()\fR and \fBpipe2()\fR functions create an I/O mechanism called a
30 pipe and returns two file descriptors, \fIfildes\fR[\fB0\fR] and
31 \fIfildes\fR[\fB1\fR]. The files associated with \fIfildes\fR[\fB0\fR]
32 and \fIfildes\fR[\fB1\fR] are streams and are both opened for reading and
33 writing. The \fBpipe()\fR call will clear the \fBONDELAY\fR,
34 \fBONBLOCK\fR, and \fBFD_CLOEXEC\fR flags on both file descriptors. The
35 \fBfcntl\fR(2) function can be used to set these flags.
36.sp
37.LP
38 The \fBpipe2()\fR call will clear the \fBONDELAY\fR on both file descriptors.
39 The \fIflags\fR argument may be used to specify attributes on both file
40 descriptors. \fBpipe2()\fR called with a \fIflags\fR value of 0 will
41 behave identically to \fBpipe()\fR. Values for \fIflags\fR are constructed
42 by a bitwise-inclusive-OR of flags from the following list, defined in
43 <\fBfcntl.h\fR>.
44.RE

45.sp
46.ne 2
47.na
48 \fBONBLOCK\fR
49.ad
50.RS 12n
51 Both file descriptors will be placed in non-blocking mode. This corresponds
52 to the \fBONBLOCK\fR flag to \fBfcntl\fR(2).
53.RE

55.sp
56.ne 2
57.na
58 \fBONCLOEXEC\fR
59.ad
60.RS 12n

```

```

61 Both file descriptors will be opened with the FD_CLOEXEC flag set. Both file
62 descriptors will be closed prior to any future exec() calls.
63.RE

65.sp
66.LP
67 A read from \fIfildes\fR[\fB0\fR] accesses the data written to
68 \fIfildes\fR[\fB1\fR] on a first-in-first-out (FIFO) basis and a read from
69 \fIfildes\fR[\fB1\fR] accesses the data written to \fIfildes\fR[\fB0\fR] also
70 on a \fBFIFO\fR basis.
71.sp
72.LP
73 Upon successful completion \fBpipe()\fR marks for update the \fBst_atime\fR,
74 \fBst_ctime\fR, and \fBst_mtime\fR fields of the pipe.
75.SH RETURN VALUES
76.sp
77.LP
78 Upon successful completion, \fB0\fR is returned. Otherwise, \fB(mil\fR is
79 returned and \fBerrno\fR is set to indicate the error.
80.SH ERRORS
81.sp
82.LP
83 The \fBpipe()\fR and \fBpipe2()\fR functions will fail if:
84.sp
85.ne 2
86.na
87 \fBEMFILE\fR
88.ad
89.RS 10n
90 More than {\fBOPEN_MAX\fR} file descriptors are already in use by this process.
91.RE

93.sp
94.ne 2
95.na
96 \fBENFILE\fR
97.ad
98.RS 10n
99 The number of simultaneously open files in the system would exceed a
100 system-imposed limit.
101.RE

103.sp
104.ne 2
105.na
106 \fBEBEFILE\fR
107.ad
108.RS 10n
109 The \fIfildes[2]\fR argument points to an illegal address.
110.RE

112.sp
113.LP
114 The \fBpipe2()\fR function will also fail if:
115.sp
116.ne 2
117.na
118 \fBEBINVAL\fR
119.ad
120.RS 10n
121 The \fIflags\fR argument is illegal. Valid \fIflags\fR are zero or a
122 bitwise inclusive-OR of \fBONCLOEXEC\fR and \fBONBLOCK\fR.
123.RE

126.SH ATTRIBUTES

```

```
127 .sp
128 .LP
129 See \fBattributes\fR(5) for descriptions of the following attributes:
130 .sp

132 .sp
133 .TS
134 box;
135 c | c
136 l | l .
137 ATTRIBUTE TYPE ATTRIBUTE VALUE
138 _
139 Interface Stability Standard
140 _
141 MT-Level Async-Signal-Safe
142 .TE

144 .SH SEE ALSO
145 .sp
146 .LP
147 \fBsh\fR(1), \fBfcntl\fR(2), \fBfstat\fR(2), \fBgetmsg\fR(2), \fBpoll\fR(2),
148 \fBputmsg\fR(2), \fBread\fR(2), \fBwrite\fR(2), \fBopen\fR(2),
149 \fBattributes\fR(5), \fBstandards\fR(5), \fBstreamio\fR(7I)
150 .SH NOTES
151 .sp
152 .LP
153 Since a pipe is bi-directional, there are two separate flows of data.
154 Therefore, the size (\fBst_size\fR) returned by a call to \fBfstat\fR(2) with
155 argument \fIfildes\fR[\fB0\fR] or \fIfildes\fR[\fB1\fR] is the number of bytes
156 available for reading from \fIfildes\fR[\fB0\fR] or \fIfildes\fR[\fB1\fR]
157 respectively. Previously, the size (\fBst_size\fR) returned by a call to
158 \fBfstat()\fR with argument \fIfildes\fR[\fB1\fR] (the write-end) was the
159 number of bytes available for reading from \fIfildes\fR[\fB0\fR] (the
160 read-end).
```



new/usr/src/man/man3c/btowc.3c

1

```
*****  
3682 Thu Jul 17 00:50:40 2014  
new/usr/src/man/man3c/btowc.3c  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

\*\*\*\*\*

6351 Thu Jul 17 00:50:40 2014

new/usr/src/man/man3c/fgetwc.3c

manpage lint.

\*\*\*\*\*

```

1  \" te
2  .\" Copyright 2014 Garrett D'Amore <garrett@damore.org>
3  .\" Copyright (c) 2003, X/Open Company Limited. All Rights Reserved. Portions
4  .\" Sun Microsystems, Inc. gratefully acknowledges The Open Group for permission
5  .\" http://www.opengroup.org/bookstore/.
6  .\" The Institute of Electrical and Electronics Engineers and The Open Group, ha
7  .\" This notice shall appear on any product containing this material.
8  .\" The contents of this file are subject to the terms of the Common Development
9  .\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
10 .\" When distributing Covered Code, include this CDDL HEADER in each file and in
11 .TH FGETWC 3C "Jun 24, 2014"
12 .SH NAME
13 fgetwc, fgetwc_l \- get a wide-character code from a stream
14 .SH SYNOPSIS
15 .LP
16 .nf
17 #include <stdio.h>
18 #include <wchar.h>
19
20 \fBwint_t\fR \fBfgetwc\fR(\fBFILE *\fR\fIstream\fR);
21 .fi
22 .LP
23 .nf
24 #include <stdio.h>
25 #include <wchar.h>
26 #include <xlocale.h>
27
28 \fBwint_t\fR \fBfgetwc_l\fR(\fBFILE *\fR\fIstream\fR, \fBlocale_t\fR, \fBIloc\fR)
29 .fi
30
31 .SH DESCRIPTION
32 .sp
33 .LP
34 The \fBfgetwc()\fR and \fBfgetwc_l()\fR functions obtain the next
35 character (if present) from the input stream pointed to by \fIstream\fR,
36 convert that to the corresponding wide-character code and advance the
37 associated file position indicator for the stream (if defined).
38 Whereas \fBfgetwc()\fR uses the current locale, \fBfgetwc_l()\fR uses the
39 locale specified by \fIloc\fR.
39 locale specified by \fIloc\fR.
40 .LP
41 If an error occurs, the resulting value of the file position indicator for the
42 stream is indeterminate.
43 .LP
44 The \fBfgetwc()\fR and \fBfgetwc_l()\fR functions may mark the \fBbst_atime\fR
45 field of the file
46 associated with \fIstream\fR for update. The \fBbst_atime\fR field will be
47 marked for update by the first successful execution of \fBfgetwc()\fR,
48 \fBfgetc\fR(3C), \fBfgetc_s\fR(3C), \fBfgetws\fR(3C), \fBfread\fR(3C),
49 \fBfscanf\fR(3C), \fBfgetc_l\fR(3C), \fBfgetchar_l\fR(3C), \fBfgetc_s_l\fR(3C), or
50 \fBscanf\fR(3C) using \fIstream\fR that returns data not supplied by a prior
51 call to \fBungetc\fR(3C) or \fBungetwc\fR(3C).
52 .SH RETURN VALUES
53 .LP
54 Upon successful completion both functions return the
55 wide-character code of the character read from the input stream pointed to by
56 \fIstream\fR converted to a type \fBwint_t\fR.
57 .LP
58 For standard-conforming (see \fBstandards\fR(5)) applications, if the
59 end-of-file indicator for the stream is set, \fBfgetwc()\fR and
60 \fBfgetwc_l()\fR return \fBEOF\fR whether or not the stream is at

```

```

61 end-of-file.
62 .LP
63 If a read error occurs, the error indicator for the stream is set,
64 \fBfgetwc()\fR and \fBfgetwc_l()\fR returns \fBEOF\fR and sets
65 \fBerrno\fR to indicate the error.
66 .LP
67 If an encoding error occurs, the error indicator for the stream is set,
68 \fBfgetwc()\fR and \fBfgetwc_l()\fR return \fBEOF\fR, and \fBerrno\fR is
69 set to indicate the error.
70 .SH ERRORS
71 .LP
72 The \fBfgetwc()\fR and \fBfgetwc_l()\fR functions will fail if data needs to be
73 read and:
74 .TP
75 .B EAGAIN
76 The \fBONONBLOCK\fR flag is set for the file descriptor underlying
77 \fIstream\fR and the process would be delayed in the \fBfgetwc()\fR or
78 \fBfgetwc_l()\fR operation.
79 .TP
80 .B EBADF
81 The file descriptor underlying \fIstream\fR is not a valid file descriptor open
82 for reading.
83 .TP
84 .B EINTR
85 The read operation was terminated due to the receipt of a signal, and no data
86 was transferred.
87 .TP
88 .B EIO
89 A physical I/O error has occurred, or the process is in a background process
90 group attempting to read from its controlling terminal and either the process
91 is ignoring or blocking the \fBSIGTIN\fR signal or the process group is
92 orphaned.
93 .TP
94 .B EOVERFLOW
95 The file is a regular file and an attempt was made to read at or beyond the
96 offset maximum associated with the corresponding \fIstream\fR.
97 .LP
98 The \fBfgetwc()\fR and \fBfgetwc_l()\fR functions may fail if:
99 .TP
100 .B ENOMEM
101 Insufficient memory is available.
102 .TP
103 .B ENXIO
104 A request was made of a non-existent device, or the request was outside the
105 capabilities of the device.
106 .TP
107 .B EILSEQ
108 The data obtained from the input stream does not form a valid character.
109 .SH USAGE
110 .sp
111 .LP
112 The \fBferror\fR(3C) or \fBfeof\fR(3C) functions must be used to distinguish
113 between an error condition and an end-of-file condition.
114 .SH ATTRIBUTES
115 .sp
116 .LP
117 See \fBattributes\fR(5) for descriptions of the following attributes:
118 .TS
119 box;
120 c | c
121 l | l .
122 ATTRIBUTE TYPE ATTRIBUTE VALUE
123 -
124 CSI Enabled
125 -
126 Interface Stability See below.

```

```
127 _
128 MT-Level      MT-Safe
129 .TE

131 .LP
132 The
133 .B fgetwc()
134 function is Standard.  The
135 .B fgetwc_l()
136 function is Uncommitted.
137 .SH SEE ALSO
138 .LP
139 \fBfeof\fR(3C), \fBferror\fR(3C), \fBfgetc\fR(3C), \fBfgets\fR(3C),
140 \fBfgetws\fR(3C), \fBfopen\fR(3C), \fBfread\fR(3C), \fBfscanf\fR(3C),
141 \fBgetc\fR(3C), \fBgetchar\fR(3C), \fBgets\fR(3C), \fBscanf\fR(3C),
142 \fBnewlocale\fR(3C), \fBsetlocale\fR(3C), \fBungetc\fR(3C), \fBungetwc\fR(3C),
143 \fBuselocale\fR(3C), \fBattributes\fR(5),
144 \fBstandards\fR(5)
```

new/usr/src/man/man3c/fopen.3c

1

```
*****  
12307 Thu Jul 17 00:50:40 2014  
new/usr/src/man/man3c/fopen.3c  
manpage lint.  
*****  
_____unchanged_portion_omitted_
```

new/usr/src/man/man3c/iswctype.3c

1

```
*****  
4950 Thu Jul 17 00:50:40 2014  
new/usr/src/man/man3c/iswctype.3c  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

new/usr/src/man/man3c/mbrtowc.3c

1

```
*****  
6226 Thu Jul 17 00:50:41 2014  
new/usr/src/man/man3c/mbrtowc.3c  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

new/usr/src/man/man3c/mbsrtowcs.3c

1

```
*****  
6940 Thu Jul 17 00:50:41 2014  
new/usr/src/man/man3c/mbsrtowcs.3c  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

new/usr/src/man/man3c/newlocale.3c

1

```
*****  
4309 Thu Jul 17 00:50:41 2014  
new/usr/src/man/man3c/newlocale.3c  
manpage lint.  
*****  
_____unchanged_portion_omitted_
```



\*\*\*\*\*

4621 Thu Jul 17 00:50:41 2014

new/usr/src/man/man3c/strcoll.3c

manpage lint.

\*\*\*\*\*

```

1  \' te
2  .\" Copyright 2014 Garrett D'Amore <garrett@damore.org>
3  .\" Copyright (c) 2003, Sun Microsystems, Inc. All Rights Reserved.
4  .\" Copyright 1989 AT&T
5  .\" Portions Copyright (c) 2001, the Institute of Electrical and Electronics Eng
6  .\" Sun Microsystems, Inc. gratefully acknowledges The Open Group for permission
7  .\" http://www.opengroup.org/bookstore/.
8  .\" The Institute of Electrical and Electronics Engineers and The Open Group, ha
9  .\" This notice shall appear on any product containing this material.
10 .\" The contents of this file are subject to the terms of the Common Development
11 .\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
12 .\" When distributing Covered Code, include this CDDL HEADER in each file and in
13 .TH STRCOLL 3C "Jun 23, 2014"
14 .SH NAME
15 strcoll, strcoll_1 \- string collation
16 .SH SYNOPSIS
17 .LP
18 .nf
19 #include <string.h>

21 \fBint\fR \fBstrcoll\fR(\fBconst char *\fR\fIs1\fR, \fBconst char *\fR\fIs2\fR);
22 .fi
23 .LP
24 .nf
25 \fBint\fR \fBstrcoll_1\fR(\fBconst char *\fR\fIs1\fR, \fBconst char *\fR\fIs2\fR
26 .fi

28 .SH DESCRIPTION
29 .LP
30 Both \fBstrcoll()\fR and \fBstrxfrm\fR(3C) provide for locale-specific string
31 sorting. \fBstrcoll()\fR is intended for applications in which the number of
32 comparisons per string is small. When strings are to be compared a number of
33 times, \fBstrxfrm\fR(3C) is a more appropriate function because the
34 transformation process occurs only once.
35 .LP
36 The \fBstrcoll_1()\fR function behaves
37 identically to \fBstrcoll()\fR, except instead of operating in the current
38 locale, it operates in the locale specified by \fIloc\fR.
39 .LP
40 The \fBstrcoll()\fR function does not change the setting of \fBerrno\fR if
41 successful.
42 .LP
43 Since no return value is reserved to indicate an error, an application wishing
44 to check for error situations should set \fBerrno\fR to 0, then call
45 \fBstrcoll()\fR, then check \fBerrno\fR.
46 .SH RETURN VALUES
47 .LP
48 Upon successful completion, \fBstrcoll()\fR returns an integer greater than,
49 equal to, or less than zero in direct correlation to whether string \fIs1\fR is
50 greater than, equal to, or less than the string \fIs2\fR. The comparison is
51 based on strings interpreted as appropriate to the locale
52 category \fBLC_COLLATE\fR (see \fBsetlocale\fR(3C)).
53 .LP
54 On error, \fBstrcoll()\fR may set \fBerrno\fR, but no return value is reserved
55 to indicate an error.
56 .SH ERRORS
57 .sp
58 .LP
59 The \fBstrcoll()\fR and \fBstrcoll_1()\fR functions may fail if:
60 .sp
61 .ne 2

```

```

62 .na
63 \fB\FBEINVAL\fR\fR
64 .ad
65 .RS 10n
66 The \fIs1\fR or \fIs2\fR arguments contain characters outside the domain of the
67 collating sequence.
68 .RE
69 .SH FILES
70 .IP \fB/usr/lib/locale/\fR\fIlocale\fR/\fB/LC_COLLATE/*\fR
71 collation database for \fIlocale\fR
72 .RE

72 .SH ATTRIBUTES
73 .LP
74 See \fBattributes\fR(5) for descriptions of the following attributes:
75 .TS
76 box;
77 c | c
78 l | l .
79 ATTRIBUTE TYPE ATTRIBUTE VALUE
80 -
81 CSI Enabled
82 -
83 Interface Stability Standard
84 -
85 MT-Level MT-Safe
86 .TE

88 .SH SEE ALSO
89 .sp
90 .LP
91 \fBlocaledef\fR(1), \fBnewlocale\fR(3C), \fBsetlocale\fR(3C), \fBstring\fR(3C),
92 \fBstrxfrm\fR(3C), \fBuselocale\fR(3C),
93 \fBwtxfrm\fR(3C), \fBattributes\fR(5), \fBenviron\fR(5), \fBstandards\fR(5)

```

new/usr/src/man/man3c/string.3c

1

```
*****  
21322 Thu Jul 17 00:50:41 2014  
new/usr/src/man/man3c/string.3c  
manpage lint.  
*****  
_____unchanged_portion_omitted_
```

```

*****
1767 Thu Jul 17 00:50:41 2014
new/usr/src/man/man3c/towlower.3c
manpage lint.
*****
1 \" te
2 .\"
3 .\" This file and its contents are supplied under the terms of the
4 .\" Common Development and Distribution License ("CDDL"), version 1.0.
5 .\" You may only use this file in accordance with the terms of version
6 .\" 1.0 of the CDDL.
7 .\"
8 .\" A full copy of the text of the CDDL should have accompanied this
9 .\" source. A copy of the CDDL is also available via the Internet at
10 .\" http://www.illumos.org/license/CDDL.
11 .\"
12 .\"
13 .\" Copyright (c) 2014 Joyent, Inc. All rights reserved.
14 .\" Copyright 2014 Garrett D'Amore <garrett@damore.org>
15 .\"
16 .TH TOWLOWER 3C "Jun 21, 2014"
17 .SH NAME
18 tolower, tolower_l \- transliterate upper-case wide characters to lower-case
19 .SH SYNOPSIS
20 .LP
21 .nf
22 #include <wctype.h>

24 \fBwint_t\fR \fBtolower\fR(\fBwint_t\fR \fBIwc\fR);
25 .fi
26 .LP
27 .nf
28 \fBwint_t\fR \fBtolower_l\fR(\fBwint_t\fR \fBIwc\fR, \fBlocale_t\fR \fBiloc\fR);
29 .fi
30 .nf
31 .SH DESCRIPTION
32 The function
33 .BR tolower()
34 is the wide character equivalent of the function
35 .BR tolower(3C).
36 It converts the upper-case wide character
37 .I wc
38 to the equivalent lower-case
39 wide character, if one exists. If one does not exist, it returns
40 .I wc
41 unchanged.
42 .LP
43 The function
44 .B tolower_l()
45 is equivalent to the function
46 .BR tolower() ,
47 but instead of operating in the current locale, operates in the
48 locale specified by
49 .IR loc .
50 .SH RETURN VALUES
51 On successful completion,
52 .B tolower()
53 and
54 .B tolower_l()
55 return the lower-case character that corresponds to the argument passed.
56 Otherwise, they return the argument unchanged.
57 .SH ERRORS
58 No errors are defined.
59 .SH ATTRIBUTES
60 .TS
61 box;

```

```

61 c | c
62 l | l .
63 ATTRIBUTE TYPE ATTRIBUTE VALUE
64 _
65 Interface Stability Standard
66 _
67 MT-Level MT-Safe
68 .TE

70 .SH SEE ALSO
71 .BR newlocale(3C),
72 .BR setlocale(3C),
73 .BR towupper(3C),
74 .BR uselocale(3C),
75 .BR locale(5)

```

new/usr/src/man/man3c/uselocale.3c

1

```
*****  
2323 Thu Jul 17 00:50:41 2014  
new/usr/src/man/man3c/uselocale.3c  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

new/usr/src/man/man3c/wcscoll.3c

1

```
*****  
4735 Thu Jul 17 00:50:41 2014  
new/usr/src/man/man3c/wcscoll.3c  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

\*\*\*\*\*

6901 Thu Jul 17 00:50:41 2014

new/usr/src/man/man3c/wcsrtombs.3c

manpage lint.

\*\*\*\*\*

```

1  \" te
2  .\\ Copyright 2014 Garrett D'Amore <garrett@damore.org>
3  .\\ Copyright (c) 1992, X/Open Company Limited. All Rights Reserved. Portions C
4  .\\ Sun Microsystems, Inc. gratefully acknowledges The Open Group for permission
5  .\\ http://www.opengroup.org/bookstore/.
6  .\\ The Institute of Electrical and Electronics Engineers and The Open Group, ha
7  .\\ This notice shall appear on any product containing this material.
8  .\\ The contents of this file are subject to the terms of the Common Development
9  .\\ You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
10 .\\ When distributing Covered Code, include this CDDL HEADER in each file and in
11 .TH WCSRTOMBS 3C "Jul 13, 2014"
12 .SH NAME
13 wcsrtombs, wcsrtombs_l, wcsrtombs, wcsrtombs_l \\- convert a wide-character str
14 (restartable)
15 .SH SYNOPSIS
16 .LP
17 .nf
18 #include <wchar.h>

20 \\fBsize_t\\fR \\fBwcsrtombs\\fR(\\fBchar *restrict\\fR \\fIdst\\fR, \\fBconst wchar_t **
21   \\fBsize_t\\fR \\fIlen\\fR, \\fBmbstate_t *restrict\\fR \\fIps\\fR);
22 .fi
23 .LP
24 .nf
25 \\fBsize_t\\fR \\fBwcsrtombs\\fR(\\fBchar *restrict\\fR \\fIdst\\fR, \\fBconst wchar_t *
26   \\fBsize_t\\fR \\fInwc\\fR, \\fBsize_t\\fR \\fIlen\\fR, \\fBmbstate_t *restrict\\fR \\fI
27 .fi
28 .LP
29 .nf
30 #include <wchar.h>
31 #include <xlocale.h>

33 \\fBsize_t\\fR \\fBwcsrtombs_l\\fR(\\fBchar *restrict\\fR \\fIdst\\fR, \\fBconst wchar_t
34   \\fBsize_t\\fR \\fIlen\\fR, \\fBmbstate_t *restrict\\fR \\fIps\\fR, \\fBlocale_t\\fR
35 .fi
36 .LP
37 .nf
38 \\fBsize_t\\fR \\fBwcsrtombs_l\\fR(\\fBchar *restrict\\fR \\fIdst\\fR, \\fBconst wchar_t
39   \\fBsize_t\\fR \\fInwc\\fR, \\fBsize_t\\fR \\fIlen\\fR, \\fBmbstate_t *restrict\\fR \\fI
40 .fi
41 .SH DESCRIPTION
42 .LP
43 The \\fBwcsrtombs()\\fR function converts a sequence of wide-characters from the
44 array indirectly pointed to by \\fIsrc\\fR into a sequence of corresponding
45 characters, beginning in the conversion state described by the object pointed
46 to by \\fIps\\fR. If \\fIdst\\fR is not a null pointer, the converted characters
47 are then stored into the array pointed to by \\fIdst\\fR. Conversion continues up
48 to and including a terminating null wide-character, which is also stored.
49 Conversion stops earlier in the following cases:
50 .RS +4
51 .TP
52 .ie t \\(bu
53 .el o
54 When a code is reached that does not correspond to a valid character.
55 .RE
56 .RS +4
57 .TP
58 .ie t \\(bu
59 .el o
60 When the next character would exceed the limit of \\fIlen\\fR total bytes to be
61 stored in the array pointed to by \\fIdst\\fR (and \\fIdst\\fR is not a null

```

```

62 pointer).
63 .RE
64 .RS +4
65 .TP
66 .ie t \\(bu
67 .el o
68 In the case of \\fBwcsrtombs()\\fR and \\fBwcsrtombs_l()\\fR, when \\fInwc\\fR
69 wide characters have been completely converted.
70 .RE
71 .LP
72 Each conversion takes place as if by a call to the \\fBwcbtomb()\\fR function.
73 .LP
74 If \\fIdst\\fR is not a null pointer, the pointer object pointed to by \\fIsrc\\fR
75 is assigned either a null pointer (if conversion stopped due to reaching a
76 terminating null wide-character) or the address just past the last
77 wide-character converted (if any). If conversion stopped due to reaching a
78 terminating null wide-character, the resulting state described is the initial
79 conversion state.
80 .LP
81 If \\fIps\\fR is a null pointer, these functions uses their own
82 internal \\fBmbstate_t\\fR object, which is initialized at program startup to the
83 initial conversion state. Otherwise, the \\fBmbstate_t\\fR object pointed to by
84 \\fIps\\fR is used to completely describe the current conversion state of the
85 associated character sequence. The system will behave as if no function defined
86 in the Reference Manual calls any of these functions.
87 .LP
88 The behavior of \\fBwcsrtombs()\\fR and \\fBwcsrtombs_l()\\fR are affected by the
89 \\fBLC_CTYPE\\fR category of the current locale. See \\fBenvron()\\fR(5).
90 .LP
91 The \\fBwcsrtombs_l()\\fR and \\fBwcsrtombs_l()\\fR functions behave identically
92 to \\fBwcsrtombs()\\fR and \\fBwcsrtombs_l()\\fR respectively, except
93 that instead of operating in the current locale, they operate in the locale
94 specified by \\fIloc\\fR.
95 .SH RETURN VALUES
96 .LP
97 If conversion stops because a code is reached that does not correspond to a
98 valid character, an encoding error occurs. In this case, these
99 functions store the value of the macro \\fBEILSEQ\\fR in \\fBerrno\\fR and return
100 \\fB(size_t)\\(mil\\fR; the conversion state is undefined. Otherwise, they return
101 the number of bytes in the resulting character sequence, not including the
102 terminating null (if any).
103 .SH ERRORS
104 .LP
105 These functions may fail if:
106 .sp
107 .ne 2
108 .na
109 \\fBEBEINVAL\\fR\\fR
110 .ad
111 .RS 10n
112 The \\fIps\\fR argument points to an object that contains an invalid conversion
113 state.
114 .RE

116 .sp
117 .ne 2
118 .na
119 \\fBEBEILSEQ\\fR\\fR
120 .ad
121 .RS 10n
122 A wide-character code does not correspond to a valid character.
123 .RE
124 .SH ATTRIBUTES
125 .LP
126 See \\fBattributes\\fR(5) for descriptions of the following attributes:
127 .TS

```

```
128 box;
129 c | c
130 l | l .
131 ATTRIBUTE TYPE  ATTRIBUTE VALUE
132 _
133 Interface Stability      See below.
134 _
135 MT-Level              See below.
136 .TE

138 .LP
139 The \fBwcsrtombs()\fR and \fBwcsnrtombs()\fR functions are Standard. The
140 \fBwcsrtombs_l()\fR and \fBwcsnrtombs_l()\fR functions are Uncommitted.
141 .LP
142 If \fIps\fR is a null pointer, these functions should be considered Unsafe
143 for use in multithreaded applications. Otherwise, they are MT-Safe.
144 .SH SEE ALSO
145 .LP
146 \fBmbsinit\fR(3C), \fBnewlocale\fR(3C), \fBsetlocale\fR(3C), \fBuselocale\fR(3C)
147 \fBwcrctomb\fR(3C), \fBattributes\fR(5),
148 \fBenviron\fR(5), \fBstandards\fR(5)
```

new/usr/src/man/man3lib/libc.3lib

1

```
*****  
34015 Thu Jul 17 00:50:42 2014  
new/usr/src/man/man3lib/libc.3lib  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```



\*\*\*\*\*

2499 Thu Jul 17 00:50:42 2014

new/usr/src/man/man4/init.4

manpage lint.

\*\*\*\*\*

```

1 \" te
2.\" Copyright 2014 Garrett D'Amore
3.\" Copyright (c) 2003, Sun Microsystems, 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 INIT 4 "Mar 15, 2014"
8.TH init 4 "Mar 15, 2014"
9.SH NAME
10 init \- set default system time zone and locale
11.SH SYNOPSIS
12.LP
13.nf
14 \fB/etc/default/init\fR
15.fi

17.SH DESCRIPTION
18.sp
19.LP
20 This file sets the time zone environment variable \fBTZ\fR, and the
21 locale-related environment variables \fBLANG\fR, \fBLC_COLLATE\fR,
22 \fBLC_CTYPE\fR, \fBLC_MESSAGES\fR, \fBLC_MONETARY\fR, \fBLC_NUMERIC\fR, and
23 \fBLC_TIME\fR.
24.sp
25.LP
26 \fB/etc/TIMEZONE\fR is a symbolic link to \fB/etc/default/init\fR. This
27 link exists for compatibility with legacy software, is obsolete, and may
28 be removed in a future release.
29.sp
30.LP
31 The number of environment variables that can be set from
32 \fB/etc/default/init\fR is limited to 20.
33.sp
34.LP
35 The format of the file is:
36.sp
37.in +2
38.nf
39 \fIVAR\fR\fB=\fR\fIvalue\fR
40.fi
41.in -2
42.sp

44.sp
45.LP
46 where \fIVAR\fR is a timezone environment variable and \fIvalue\fR is the value
47 assigned to the variable. \fIvalue\fR can be enclosed in double quotes (") or
48 single quotes (\&''). The double or single quotes cannot be part of the value.
49.SH SEE ALSO
50.sp
51.LP
52 \fBinit\fR(1M), \fBrtc\fR(1M), \fBctime\fR(3C), \fBenviron\fR(5)
53.SH NOTES
54.sp
55.LP
56 When changing the \fBTZ\fR setting on x86 systems, you must make a
57 corresponding change to the \fB/etc/rtc_config\fR file to account for the new
58 timezone setting. This can be accomplished by executing the following commands,
59 followed by a reboot, to make the changes take effect:
60.sp

```

```

61.in +2
62.nf
63 # rtc \fB-z\fR \fIzone-name\fR
64 # rtc \fB-c\fR

66.fi
67.in -2
68.sp

70.sp
71.LP
72 where \fIzone-name\fR is the same name as the \fBTZ\fR variable setting.
73.sp
74.LP
75 See \fBrtc\fR(1M) for information on the \fBrtc\fR command.

```

\*\*\*\*\*

76160 Thu Jul 17 00:50:42 2014

new/usr/src/man/man5/mdoc.5

manpage lint.

\*\*\*\*\*

```

1 .\"
2 .\" Permission to use, copy, modify, and distribute this software for any
3 .\" purpose with or without fee is hereby granted, provided that the above
4 .\" copyright notice and this permission notice appear in all copies.
5 .\"
6 .\" THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES
7 .\" WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF
8 .\" MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR
9 .\" ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES
10 .\" WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN
11 .\" ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF
12 .\" OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.
13 .\"
14 .\"
15 .\" Copyright (c) 2009, 2010, 2011 Kristaps Dzonsons <kristaps@bsd.lv>
16 .\" Copyright (c) 2010, 2011 Ingo Schwarze <schwarze@openbsd.org>
17 .\" Copyright 2012 Nexenta Systems, Inc. All rights reserved.
18 .\" Copyright 2014 Garrett D'Amore <garrett@dmaore.org>
19 .\"
20 .Dd Jul 16, 2014
21 .Dt MDOC 5
22 .Os
23 .Sh NAME
24 .Nm mdoc
25 .Nd semantic markup language for formatting manual pages
26 .Sh DESCRIPTION
27 The
28 .Nm mdoc
29 language supports authoring of manual pages for the
30 .Xr man 1
31 utility by allowing semantic annotations of words, phrases,
32 page sections and complete manual pages.
33 Such annotations are used by formatting tools to achieve a uniform
34 presentation across all manuals written in
35 .Nm ,
36 and to support hyperlinking if supported by the output medium.
37 .Pp
38 This reference document describes the structure of manual pages
39 and the syntax and usage of the
40 .Nm
41 language.
42 The reference implementation of a parsing and formatting tool is
43 .Xr mandoc 1 ;
44 the
45 .Sx COMPATIBILITY
46 section describes compatibility with other implementations.
47 .Pp
48 In an
49 .Nm
50 document, lines beginning with the control character
51 .Sq \&.
52 are called
53 .Dq macro lines .
54 The first word is the macro name.
55 It consists of two or three letters.
56 Most macro names begin with a capital letter.
57 For a list of available macros, see
58 .Sx MACRO OVERVIEW .
59 The words following the macro name are arguments to the macro, optionally
60 including the names of other, callable macros; see
61 .Sx MACRO SYNTAX

```

```

62 for details.
63 .Pp
64 Lines not beginning with the control character are called
65 .Dq text lines .
66 They provide free-form text to be printed; the formatting of the text
67 depends on the respective processing context:
68 .Bd -literal -offset indent
69 \&.Sh Macro lines change control state.
70 Text lines are interpreted within the current state.
71 .Ed
72 .Pp
73 Many aspects of the basic syntax of the
74 .Nm
75 language are based on the
76 .Xr roff 5
77 language; see the
78 .Em LANGUAGE SYNTAX
79 and
80 .Em MACRO SYNTAX
81 sections in the
82 .Xr roff 5
83 manual for details, in particular regarding
84 comments, escape sequences, whitespace, and quoting.
85 However, using
86 .Xr roff 5
87 requests in
88 .Nm
89 documents is discouraged;
90 .Xr mandoc 1
91 supports some of them merely for backward compatibility.
92 .Sh MANUAL STRUCTURE
93 A well-formed
94 .Nm
95 document consists of a document prologue followed by one or more
96 sections.
97 .Pp
98 The prologue, which consists of the
99 .Sx \&Dd ,
100 .Sx \&Dt ,
101 and
102 .Sx \&Os
103 macros in that order, is required for every document.
104 .Pp
105 The first section (sections are denoted by
106 .Sx \&Sh )
107 must be the NAME section, consisting of at least one
108 .Sx \&Nm
109 followed by
110 .Sx \&Nd .
111 .Pp
112 Following that, convention dictates specifying at least the
113 .Em SYNOPSIS
114 and
115 .Em DESCRIPTION
116 sections, although this varies between manual sections.
117 .Pp
118 The following is a well-formed skeleton
119 .Nm
120 file for a utility
121 .Qq progname :
122 .Bd -literal -offset indent
123 \&.Dd Jan 1, 1970
124 \&.Dt PROGNAME section
125 \&.Os
126 \&.Sh NAME
127 \&.Nm progname

```

```

128 \&.Nd one line description
129 \&.\e\(\dq .Sh LIBRARY
130 \&.\e\(\dq For sections 2, 3, & 9 only.
131 \&.Sh SYNOPSIS
132 \&.Nm progname
133 \&.Op Fl options
134 \&.Ar
135 \&.Sh DESCRIPTION
136 The
137 \&.Nm
138 utility processes files ...
139 \&.\e\(\dq .Sh IMPLEMENTATION NOTES
140 \&.\e\(\dq .Sh RETURN VALUES
141 \&.\e\(\dq For sections 2, 3, & 9 only.
142 \&.\e\(\dq .Sh ENVIRONMENT
143 \&.\e\(\dq For sections 1, 1M, 5, & 6 only.
144 \&.\e\(\dq .Sh FILES
145 \&.\e\(\dq .Sh EXIT STATUS
146 \&.\e\(\dq For sections 1, 1M, & 6 only.
147 \&.\e\(\dq .Sh EXAMPLES
148 \&.\e\(\dq .Sh DIAGNOSTICS
149 \&.\e\(\dq For sections 1, 1M, 5, 6, & 7 only.
150 \&.\e\(\dq .Sh ERRORS
151 \&.\e\(\dq For sections 2, 3, & 9 only.
152 \&.\e\(\dq .Sh ARCHITECTURE
153 \&.\e\(\dq .Sh CODE SET INDEPENDENCE
154 \&.\e\(\dq For sections 1, 1M, & 3 only.
155 \&.\e\(\dq .Sh INTERFACE STABILITY
156 \&.\e\(\dq .Sh MT-LEVEL
157 \&.\e\(\dq For sections 2 & 3 only.
158 \&.\e\(\dq .Sh SEE ALSO
159 \&.\e\(\dq .Xr foobar 1
160 \&.\e\(\dq .Sh STANDARDS
161 \&.\e\(\dq .Sh HISTORY
162 \&.\e\(\dq .Sh AUTHORS
163 \&.\e\(\dq .Sh CAVEATS
164 \&.\e\(\dq .Sh BUGS
165 \&.\e\(\dq .Sh SECURITY CONSIDERATIONS
166 \&.\e\(\dq Not used in OpenBSD.
167 .Ed
168 .Pp
169 The sections in an
170 .Nm
171 document are conventionally ordered as they appear above.
172 Sections should be composed as follows:
173 .Bl -ohang -offset Ds
174 .It Em NAME
175 The name(s) and a one line description of the documented material.
176 The syntax for this as follows:
177 .Bd -literal -offset indent
178 \&.Nm name0 ,
179 \&.Nm name1 ,
180 \&.Nm name2
181 \&.Nd a one line description
182 .Ed
183 .Pp
184 Multiple
185 .Sq \&Nm
186 names should be separated by commas.
187 .Pp
188 The
189 .Sx \&Nm
190 macro(s) must precede the
191 .Sx \&Nd
192 macro.
193 .Pp

```

```

194 See
195 .Sx \&Nm
196 and
197 .Sx \&Nd .
198 .It Em LIBRARY
199 The name of the library containing the documented material, which is
200 assumed to be a function in a section 2, 3, or 9 manual.
201 The syntax for this is as follows:
202 .Bd -literal -offset indent
203 \&.Lb libarm
204 .Ed
205 .Pp
206 See
207 .Sx \&Lb .
208 .It Em SYNOPSIS
209 Documents the utility invocation syntax, function call syntax, or device
210 configuration.
211 .Pp
212 For the first, utilities (sections 1, 1M, and 6), this is
213 generally structured as follows:
214 .Bd -literal -offset indent
215 \&.Nm bar
216 \&.Op Fl v
217 \&.Op Fl o Ar file
218 \&.Op Ar
219 \&.Nm foo
220 \&.Op Fl v
221 \&.Op Fl o Ar file
222 \&.Op Ar
223 .Ed
224 .Pp
225 Commands should be ordered alphabetically.
226 .Pp
227 For the second, function calls (sections 2, 3, 9):
228 .Bd -literal -offset indent
229 \&.In header.h
230 \&.Vt extern const char *global;
231 \&.Ft "char *"
232 \&.Fn foo "const char *src"
233 \&.Ft "char *"
234 \&.Fn bar "const char *src"
235 .Ed
236 .Pp
237 Ordering of
238 .Sx \&In ,
239 .Sx \&Vt ,
240 .Sx \&Fn ,
241 and
242 .Sx \&Fo
243 macros should follow C header-file conventions.
244 .Pp
245 And for the third, configurations (section 7):
246 .Bd -literal -offset indent
247 \&.Cd \(\dqit* at isa? port 0x2e\(\dq
248 \&.Cd \(\dqit* at isa? port 0x4e\(\dq
249 .Ed
250 .Pp
251 Manuals not in these sections generally don't need a
252 .Em SYNOPSIS .
253 .Pp
254 Some macros are displayed differently in the
255 .Em SYNOPSIS
256 section, particularly
257 .Sx \&Nm ,
258 .Sx \&Cd ,
259 .Sx \&Fd ,

```

```

260 .Sx \&Fn ,
261 .Sx \&Fo ,
262 .Sx \&In ,
263 .Sx \&Vt ,
264 and
265 .Sx \&Ft .
266 All of these macros are output on their own line.
267 If two such dissimilar macros are pairwise invoked (except for
268 .Sx \&Ft
269 before
270 .Sx \&Fo
271 or
272 .Sx \&Fn ) ,
273 they are separated by a vertical space, unless in the case of
274 .Sx \&Fo ,
275 .Sx \&Fn ,
276 and
277 .Sx \&Ft ,
278 which are always separated by vertical space.
279 .Pp
280 When text and macros following an
281 .Sx \&Nm
282 macro starting an input line span multiple output lines,
283 all output lines but the first will be indented to align
284 with the text immediately following the
285 .Sx \&Nm
286 macro, up to the next
287 .Sx \&Nm ,
288 .Sx \&Sh ,
289 or
290 .Sx \&Ss
291 macro or the end of an enclosing block, whichever comes first.
292 .It Em DESCRIPTION
293 This begins with an expansion of the brief, one line description in
294 .Em NAME :
295 .Bd -literal -offset indent
296 The
297 \&.Nm
298 utility does this, that, and the other.
299 .Ed
300 .Pp
301 It usually follows with a breakdown of the options (if documenting a
302 command), such as:
303 .Bd -literal -offset indent
304 The arguments are as follows:
305 \&.Bl \-tag \-width Ds
306 \&.It Fl v
307 Print verbose information.
308 \&.El
309 .Ed
310 .Pp
311 Manuals not documenting a command won't include the above fragment.
312 .Pp
313 Since the
314 .Em DESCRIPTION
315 section usually contains most of the text of a manual, longer manuals
316 often use the
317 .Sx \&Ss
318 macro to form subsections.
319 In very long manuals, the
320 .Em DESCRIPTION
321 may be split into multiple sections, each started by an
322 .Sx \&Sh
323 macro followed by a non-standard section name, and each having
324 several subsections, like in the present
325 .Nm

```

```

326 manual.
327 .It Em IMPLEMENTATION NOTES
328 Implementation-specific notes should be kept here.
329 This is useful when implementing standard functions that may have side
330 effects or notable algorithmic implications.
331 .It Em RETURN VALUES
332 This section documents the
333 return values of functions in sections 2, 3, and 9.
334 .Pp
335 See
336 .Sx \&Rv .
337 .It Em ENVIRONMENT
338 Lists the environment variables used by the utility,
339 and explains the syntax and semantics of their values.
340 The
341 .Xr environ 5
342 manual provides examples of typical content and formatting.
343 .Pp
344 See
345 .Sx \&Ev .
346 .It Em FILES
347 Documents files used.
348 It's helpful to document both the file name and a short description of how
349 the file is used (created, modified, etc.).
350 .Pp
351 See
352 .Sx \&Pa .
353 .It Em EXIT STATUS
354 This section documents the
355 command exit status for section 1, 6, and 8 utilities.
356 Historically, this information was described in
357 .Em DIAGNOSTICS ,
358 a practise that is now discouraged.
359 .Pp
360 See
361 .Sx \&Ex .
362 .It Em EXAMPLES
363 Example usages.
364 This often contains snippets of well-formed, well-tested invocations.
365 Make sure that examples work properly!
366 .It Em DIAGNOSTICS
367 Documents error conditions.
368 This is most useful in section 4 manuals.
369 Historically, this section was used in place of
370 .Em EXIT STATUS
371 for manuals in sections 1, 6, and 8; however, this practise is
372 discouraged.
373 .Pp
374 See
375 .Sx \&Bl
376 .Fl diag .
377 .It Em ERRORS
378 Documents error handling in sections 2, 3, and 9.
379 .Pp
380 See
381 .Sx \&Er .
382 .It Em ARCHITECTURE
383 This section is usually absent, but will be present when the
384 interface is specific to one or more architectures.
385 .It Em CODE SET INDEPENDENCE
386 Indicates whether the interface operates correctly with various different
387 code sets. True independent code sets will support not only ASCII and
388 Extended UNIX Codesets (EUC), but also other multi-byte encodings such as
389 UTF-8 and GB2312.
390 .Pp
391 Generally there will be some limitations that are fairly standard. See

```

392 .Xr standards 5 for more information about some of these. Most interfaces  
 393 should support at least UTF-8 in addition to ASCII.  
 394 .It Em INTERFACE STABILITY  
 395 Indicates the level of commitment to the interface. Interfaces can be described  
 396 with in the following ways:  
 397 **.Bl -tag -width Ds**  
 397 **.Bl -tag**  
 398 **.It Nm Standard**  
 399 Indicates that the interface is defined by one or more standards bodies.  
 400 Generally, changes to the interface will be carefully managed to conform  
 401 to the relevant standards. These interfaces are generally the most suitable  
 402 for use in portable programs.  
 403 **.It Nm Committed**  
 404 Indicates that the interface is intended to be preserved for the long-haul, and  
 405 will rarely, if ever change, and never without notification (barring  
 406 extraordinary and extenuating circumstances). These interfaces are  
 407 preferred over other interfaces with the exception of  
 408 **.Nm Standard**  
 409 interfaces.  
 410 **.It Nm Uncommitted**  
 411 Indicates that the interface may change. Generally, changes to these interfaces  
 412 should be infrequent, and some effort will be made to address compatibility  
 413 considerations when changing or removing such interfaces. However, there is  
 414 no firm commitment to the preservation of the interface. Most often this  
 415 is applied to interfaces where operational experience with the interface  
 416 is still limited and some need to change may be anticipated.  
 417 **.Pp**  
 418 Consumers should expect to revalidate any  
 419 **.Nm Uncommitted**  
 420 interfaces when crossing release boundaries. Products intended for  
 421 use on many releases or intended to support compatibility with future  
 422 releases should avoid these interfaces.  
 423 **.It Nm Volatile**  
 424 The interface can change at any time for any reason. Often this relates to  
 425 interfaces that are part of external software components that are still evolving  
 426 rapidly. Consumers should not expect that the interface (either binary or  
 427 source level) will be unchanged from one release to the next.  
 428 **.It Nm Not-an-Interface**  
 429 Describes something that is specifically not intended for programmatic  
 430 consumption. For example, specific human-readable output, or the layout  
 431 of graphical items on a user interface, may be described this way. Generally  
 432 programmatic alternatives to these will be available, and should be used  
 433 when programmatic consumption is needed.  
 434 **.It Nm Private**  
 435 This is an internal interface. Generally these interfaces should only be  
 436 used within the project, and should not be used by other programs or modules.  
 437 The interface can and will change without notice as the project needs, at  
 438 any time.  
 439 **.Pp**  
 440 Most often, Private interfaces will lack any documentation whatsoever, and  
 441 generally any undocumented interface can be assumed to be Private.  
 442 **.It Nm Obsolete**  
 443 The interface is not intended for use in new projects or programs, and may  
 444 be removed at a future date. The  
 445 **.Nm Obsolete**  
 446 word is a modifier that can  
 447 be applied to other commitment levels. For example an  
 448 **.Nm Obsolete Committed**  
 449 interface is unlikely to be removed or changed, but nonetheless new use  
 450 is discouraged (perhaps a better newer alternative is present).  
 451 **.El**  
 452 **.It Em MT-LEVEL**  
 453 This section describes considerations for the interface when used within  
 454 programs that use multiple threads. More discussion of these considerations  
 455 is made in the MT-Level section of  
 456 **.Xr attributes 5** .

457 The interface can be described in the following ways.  
 458 **.Bl -tag -width Ds**  
 458 **.Bl -tag**  
 459 **.It Nm Safe**  
 460 Indicates the interface is safe for use within multiple threads. There  
 461 may be additional caveats that apply, in which case those will be  
 462 described. Note that some interfaces have semantics which may affect  
 463 other threads, but these should be an intrinsic part of the interface  
 464 rather than an unexpected side effect. For example, closing a file in  
 465 one thread will cause that file to be closed in all threads.  
 466 **.It Nm Unsafe**  
 467 Indicates the interface is unsuitable for concurrent use within multiple  
 468 threads. A threaded application may still make use of the interface, but  
 469 will be required to provide external synchronization means to ensure that  
 470 only a single thread calls the interface at a time.  
 471 **.It Nm MT-Safe**  
 472 Indicates that the interface is not only safe for concurrent use, but is  
 473 designed for such use. For example, a  
 474 **.Nm Safe**  
 475 interface may make use of a global lock to provide safety, but at reduced  
 476 internal concurrency, whereas an  
 477 **.Nm MT-Safe**  
 478 interface will be designed to be efficient even when used concurrently.  
 479 **.It Nm Async-Signal-Safe**  
 480 Indicates that the library is safe for use within a signal handler. An  
 481 **.Nm MT-Safe**  
 482 interface can be made  
 483 **.Nm Async-Signal-Safe**  
 484 by ensuring that it blocks signals when acquiring locks.  
 485 **.It Nm Safe with Exceptions**  
 486 As for  
 487 **.Nm Safe**  
 488 but with specific exceptions noted.  
 489 **.It Nm MT-Safe with Exceptions**  
 490 As for  
 491 **.Nm MT-Safe**  
 492 but with specific exceptions noted.  
 493 **.El**  
 494 **.It Em SEE ALSO**  
 495 References other manuals with related topics.  
 496 This section should exist for most manuals.  
 497 Cross-references should conventionally be ordered first by section, then  
 498 alphabetically.  
 499 **.Pp**  
 500 References to other documentation concerning the topic of the manual page,  
 501 for example authoritative books or journal articles, may also be  
 502 provided in this section.  
 503 **.Pp**  
 504 **See**  
 505 **.Sx \&Rs**  
 506 **and**  
 507 **.Sx \&Xr** .  
 508 **.It Em STANDARDS**  
 509 References any standards implemented or used.  
 510 If not adhering to any standards, the  
 511 **.Em HISTORY**  
 512 section should be used instead.  
 513 **.Pp**  
 514 **See**  
 515 **.Sx \&St** .  
 516 **.It Em HISTORY**  
 517 A brief history of the subject, including where it was first implemented,  
 518 and when it was ported to or reimplemented for the operating system at hand.  
 519 **.It Em AUTHORS**  
 520 Credits to the person or persons who wrote the code and/or documentation.  
 521 Authors should generally be noted by both name and email address.

```

522 .Pp
523 See
524 .Sx \&An .
525 .It Em CAVEATS
526 Common misuses and misunderstandings should be explained
527 in this section.
528 .It Em BUGS
529 Known bugs, limitations, and work-arounds should be described
530 in this section.
531 .It Em SECURITY CONSIDERATIONS
532 Documents any security precautions that operators should consider.
533 .El
534 .Sh MACRO OVERVIEW
535 This overview is sorted such that macros of similar purpose are listed
536 together, to help find the best macro for any given purpose.
537 Deprecated macros are not included in the overview, but can be found below
538 in the alphabetical
539 .Sx MACRO REFERENCE .
540 .Ss Document preamble and NAME section macros
541 .Bl -column "Brq, Bro, Brc" description
542 .It Sx \&Dd Ta document date: Ar month day , year
543 .It Sx \&Dt Ta document title: Ar TITLE SECTION Op Ar volume | arch
544 .It Sx \&Os Ta operating system version: Op Ar system Op Ar version
545 .It Sx \&Nm Ta document name (one argument)
546 .It Sx \&Nd Ta document description (one line)
547 .El
548 .Ss Sections and cross references
549 .Bl -column "Brq, Bro, Brc" description
550 .It Sx \&Sh Ta section header (one line)
551 .It Sx \&Ss Ta subsection header (one line)
552 .It Sx \&Sx Ta internal cross reference to a section or subsection
553 .It Sx \&Xr Ta cross reference to another manual page: Ar name section
554 .It Sx \&Pp , \&Lp Ta start a text paragraph (no arguments)
555 .El
556 .Ss Displays and lists
557 .Bl -column "Brq, Bro, Brc" description
558 .It Sx \&Bd , \&Ed Ta display block:
559 .Fl Ar type
560 .Op Fl offset Ar width
561 .Op Fl compact
562 .It Sx \&Dl Ta indented display (one line)
563 .It Sx \&Dl Ta indented literal display (one line)
564 .It Sx \&Bl , \&El Ta list block:
565 .Fl Ar type
566 .Op Fl width Ar val
567 .Op Fl offset Ar val
568 .Op Fl compact
569 .It Sx \&It Ta list item (syntax depends on Fl Ar type )
570 .It Sx \&Ta Ta table cell separator in Sx \&Bl Fl column No lists
571 .It Sx \&Rs , \&%* , \&Re Ta bibliographic block (references)
572 .El
573 .Ss Spacing control
574 .Bl -column "Brq, Bro, Brc" description
575 .It Sx \&PF Ta prefix, no following horizontal space (one argument)
576 .It Sx \&Ns Ta roman font, no preceding horizontal space (no arguments)
577 .It Sx \&Ap Ta apostrophe without surrounding whitespace (no arguments)
578 .It Sx \&Sm Ta switch horizontal spacing mode: Cm on | off
579 .It Sx \&Bk , \&Ek Ta keep block: Fl words
580 .It Sx \&br Ta force output line break in text mode (no arguments)
581 .It Sx \&sp Ta force vertical space: Op Ar height
582 .El
583 .Ss Semantic markup for command line utilities:
584 .Bl -column "Brq, Bro, Brc" description
585 .It Sx \&Nm Ta start a SYNOPSIS block with the name of a utility
586 .It Sx \&Fl Ta command line options (flags) (>=0 arguments)
587 .It Sx \&Cm Ta command modifier (>0 arguments)

```

```

588 .It Sx \&Ar Ta command arguments (>=0 arguments)
589 .It Sx \&Op , \&Oo , \&Oc Ta optional syntax elements (enclosure)
590 .It Sx \&Ic Ta internal or interactive command (>0 arguments)
591 .It Sx \&Ev Ta environmental variable (>0 arguments)
592 .It Sx \&Pa Ta file system path (>=0 arguments)
593 .El
594 .Ss Semantic markup for function libraries:
595 .Bl -column "Brq, Bro, Brc" description
596 .It Sx \&Lb Ta function library (one argument)
597 .It Sx \&In Ta include file (one argument)
598 .It Sx \&Ft Ta function type (>0 arguments)
599 .It Sx \&Fo , \&Fc Ta function block: Ar funcname
600 .It Sx \&Fn Ta function name:
601 .Op Ar functype
602 .Ar funcname
603 .Oo
604 .Op Ar argtype
605 .Ar argname
606 .Oc
607 .It Sx \&Fa Ta function argument (>0 arguments)
608 .It Sx \&Vt Ta variable type (>0 arguments)
609 .It Sx \&Va Ta variable name (>0 arguments)
610 .It Sx \&Dv Ta defined variable or preprocessor constant (>0 arguments)
611 .It Sx \&Er Ta error constant (>0 arguments)
612 .It Sx \&Ev Ta environmental variable (>0 arguments)
613 .El
614 .Ss Various semantic markup:
615 .Bl -column "Brq, Bro, Brc" description
616 .It Sx \&An Ta author name (>0 arguments)
617 .It Sx \&Lk Ta hyperlink: Ar uri Op Ar name
618 .It Sx \&Mt Ta Do mailto Dc hyperlink: Ar address
619 .It Sx \&Cd Ta kernel configuration declaration (>0 arguments)
620 .It Sx \&Ad Ta memory address (>0 arguments)
621 .It Sx \&Ms Ta mathematical symbol (>0 arguments)
622 .It Sx \&Tn Ta tradename (>0 arguments)
623 .El
624 .Ss Physical markup
625 .Bl -column "Brq, Bro, Brc" description
626 .It Sx \&Em Ta italic font or underline (emphasis) (>0 arguments)
627 .It Sx \&Sy Ta boldface font (symbolic) (>0 arguments)
628 .It Sx \&Li Ta typewriter font (literal) (>0 arguments)
629 .It Sx \&No Ta return to roman font (normal) (no arguments)
630 .It Sx \&Bf , \&Ef Ta font block:
631 .Op Fl Ar type | Cm \&Em | \&Li | \&Sy
632 .El
633 .Ss Physical enclosures
634 .Bl -column "Brq, Bro, Brc" description
635 .It Sx \&Dq , \&Do , \&Dc Ta enclose in typographic double quotes: Dq text
636 .It Sx \&Qq , \&Qo , \&Qc Ta enclose in typewriter double quotes: Qq text
637 .It Sx \&Sq , \&So , \&Sc Ta enclose in single quotes: Sq text
638 .It Sx \&Ql Ta single-quoted literal text: Ql text
639 .It Sx \&Pq , \&Po , \&Pc Ta enclose in parentheses: Pq text
640 .It Sx \&Bq , \&Bo , \&Bc Ta enclose in square brackets: Bq text
641 .It Sx \&Brq , \&Bro , \&Brc Ta enclose in curly braces: Brq text
642 .It Sx \&Aq , \&Ao , \&Ac Ta enclose in angle brackets: Aq text
643 .It Sx \&Eo , \&Ec Ta generic enclosure
644 .El
645 .Ss Text production
646 .Bl -column "Brq, Bro, Brc" description
647 .It Sx \&Ex Fl std Ta standard command exit values: Op Ar utility ...
648 .It Sx \&Rv Fl std Ta standard function return values: Op Ar function ...
649 .It Sx \&St Ta reference to a standards document (one argument)
650 .It Sx \&Ux Ta Ux
651 .It Sx \&At Ta At
652 .It Sx \&Bx Ta Bx
653 .It Sx \&Bsx Ta Bsx

```

654 .It Sx \&Nx Ta Nx  
 655 .It Sx \&Fx Ta Fx  
 656 .It Sx \&Ox Ta Ox  
 657 .It Sx \&Dx Ta Dx  
 658 .El  
 659 .Sh MACRO REFERENCE  
 660 This section is a canonical reference of all macros, arranged  
 661 alphabetically.  
 662 For the scoping of individual macros, see  
 663 .Sx MACRO SYNTAX .  
 664 .Ss \&%A  
 665 Author name of an  
 666 .Sx \&Rs  
 667 block.  
 668 Multiple authors should each be accorded their own  
 669 .Sx \%A  
 670 line.  
 671 Author names should be ordered with full or abbreviated forename(s)  
 672 first, then full surname.  
 673 .Ss \&%B  
 674 Book title of an  
 675 .Sx \&Rs  
 676 block.  
 677 This macro may also be used in a non-bibliographic context when  
 678 referring to book titles.  
 679 .Ss \&%C  
 680 Publication city or location of an  
 681 .Sx \&Rs  
 682 block.  
 683 .Ss \&%D  
 684 Publication date of an  
 685 .Sx \&Rs  
 686 block.  
 687 Recommended formats of arguments are  
 688 .Ar month day , year  
 689 or just  
 690 .Ar year .  
 691 .Ss \&%I  
 692 Publisher or issuer name of an  
 693 .Sx \&Rs  
 694 block.  
 695 .Ss \&%J  
 696 Journal name of an  
 697 .Sx \&Rs  
 698 block.  
 699 .Ss \&%N  
 700 Issue number (usually for journals) of an  
 701 .Sx \&Rs  
 702 block.  
 703 .Ss \&%O  
 704 Optional information of an  
 705 .Sx \&Rs  
 706 block.  
 707 .Ss \&%P  
 708 Book or journal page number of an  
 709 .Sx \&Rs  
 710 block.  
 711 .Ss \&%Q  
 712 Institutional author (school, government, etc.) of an  
 713 .Sx \&Rs  
 714 block.  
 715 Multiple institutional authors should each be accorded their own  
 716 .Sx \%Q  
 717 line.  
 718 .Ss \&%R  
 719 Technical report name of an

720 .Sx \&Rs  
 721 block.  
 722 .Ss \&%T  
 723 Article title of an  
 724 .Sx \&Rs  
 725 block.  
 726 This macro may also be used in a non-bibliographical context when  
 727 referring to article titles.  
 728 .Ss \&%U  
 729 URI of reference document.  
 730 .Ss \&%V  
 731 Volume number of an  
 732 .Sx \&Rs  
 733 block.  
 734 .Ss \&Ac  
 735 Close an  
 736 .Sx \&Ao  
 737 block.  
 738 Does not have any tail arguments.  
 739 .Ss \&Ad  
 740 Memory address.  
 741 Do not use this for postal addresses.  
 742 .Pp  
 743 Examples:  
 744 .Dl \&.Ad [0,\$]  
 745 .Dl \&.Ad 0x00000000  
 746 .Ss \&An  
 747 Author name.  
 748 Can be used both for the authors of the program, function, or driver  
 749 documented in the manual, or for the authors of the manual itself.  
 750 Requires either the name of an author or one of the following arguments:  
 751 .Pp  
 752 .Bl -tag -width "-nosplitX" -offset indent -compact  
 753 .It Fl split  
 754 Start a new output line before each subsequent invocation of  
 755 .Sx \&An .  
 756 .It Fl nosplit  
 757 The opposite of  
 758 .Fl split .  
 759 .El  
 760 .Pp  
 761 The default is  
 762 .Fl nosplit .  
 763 The effect of selecting either of the  
 764 .Fl split  
 765 modes ends at the beginning of the  
 766 .Em AUTHORS  
 767 section.  
 768 In the  
 769 .Em AUTHORS  
 770 section, the default is  
 771 .Fl nosplit  
 772 for the first author listing and  
 773 .Fl split  
 774 for all other author listings.  
 775 .Pp  
 776 Examples:  
 777 .Dl \&.An -nosplit  
 778 .Dl \&.An Kristaps Dzonsons \&Aq kristaps@bsd.lv  
 779 .Ss \&Ao  
 780 Begin a block enclosed by angle brackets.  
 781 Does not have any head arguments.  
 782 .Pp  
 783 Examples:  
 784 .Dl \&.Fl -key= \&Ns \&Ao \&Ar val \&Ac  
 785 .Pp

786 See also  
787 .Sx \&Aq .  
788 .Ss \&Ap  
789 Inserts an apostrophe without any surrounding whitespace.  
790 This is generally used as a grammatical device when referring to the verb  
791 form of a function.  
792 .Pp  
793 Examples:  
794 .Dl \&.Fn execve \&Ap d  
795 .Ss \&Aq  
796 Encloses its arguments in angle brackets.  
797 .Pp  
798 Examples:  
799 .Dl \&.Fl -key= \&Ns \&Aq \&Ar val  
800 .Pp  
801 .Em Remarks :  
802 this macro is often abused for rendering URIs, which should instead use  
803 .Sx \&Lk  
804 or  
805 .Sx \&Mt ,  
806 or to note pre-processor  
807 .Dq Li #include  
808 statements, which should use  
809 .Sx \&In .  
810 .Pp  
811 See also  
812 .Sx \&Ao .  
813 .Ss \&Ar  
814 Command arguments.  
815 If an argument is not provided, the string  
816 .Dq file ...\  
817 is used as a default.  
818 .Pp  
819 Examples:  
820 .Dl ".Fl o Ar file"  
821 .Dl ".Ar"  
822 .Dl ".Ar arg1 , arg2 ."  
823 .Pp  
824 The arguments to the  
825 .Sx \&Ar  
826 macro are names and placeholders for command arguments;  
827 for fixed strings to be passed verbatim as arguments, use  
828 .Sx \&Fl  
829 or  
830 .Sx \&Cm .  
831 .Ss \&At  
832 Formats an AT&T version.  
833 Accepts one optional argument:  
834 .Pp  
835 .Bl -tag -width "v[1-7] | 32vX" -offset indent -compact  
836 .It Cm v[1-7] | 32v  
837 A version of  
838 .At .  
839 .It Cm III  
840 .At III .  
841 .It Cm V[. [1-4]]?  
842 A version of  
843 .At V .  
844 .El  
845 .Pp  
846 Note that these arguments do not begin with a hyphen.  
847 .Pp  
848 Examples:  
849 .Dl \&.At  
850 .Dl \&.At III  
851 .Dl \&.At V.1

852 .Pp  
853 See also  
854 .Sx \&Bsx ,  
855 .Sx \&Bx ,  
856 .Sx \&Dx ,  
857 .Sx \&Fx ,  
858 .Sx \&Nx ,  
859 .Sx \&Ox ,  
860 and  
861 .Sx \&Ux .  
862 .Ss \&Bc  
863 Close a  
864 .Sx \&Bo  
865 block.  
866 Does not have any tail arguments.  
867 .Ss \&Bd  
868 Begin a display block.  
869 Its syntax is as follows:  
870 .Bd -ragged -offset indent  
871 .Pf \. Sx \&Bd  
872 .Fl Ns Ar type  
873 .Op Fl offset Ar width  
874 .Op Fl compact  
875 .Ed  
876 .Pp  
877 Display blocks are used to select a different indentation and  
878 justification than the one used by the surrounding text.  
879 They may contain both macro lines and text lines.  
880 By default, a display block is preceded by a vertical space.  
881 .Pp  
882 The  
883 .Ar type  
884 must be one of the following:  
885 .Bl -tag -width l3n -offset indent  
886 .It Fl centered  
887 Produce one output line from each input line, and centre-justify each line.  
888 Using this display type is not recommended; many  
889 .Nm  
890 implementations render it poorly.  
891 .It Fl filled  
892 Change the positions of line breaks to fill each line, and left- and  
893 right-justify the resulting block.  
894 .It Fl literal  
895 Produce one output line from each input line,  
896 and do not justify the block at all.  
897 Preserve white space as it appears in the input.  
898 Always use a constant-width font.  
899 Use this for displaying source code.  
900 .It Fl ragged  
901 Change the positions of line breaks to fill each line, and left-justify  
902 the resulting block.  
903 .It Fl unfilled  
904 The same as  
905 .Fl literal ,  
906 but using the same font as for normal text, which is a variable width font  
907 if supported by the output device.  
908 .El  
909 .Pp  
910 The  
911 .Ar type  
912 must be provided first.  
913 Additional arguments may follow:  
914 .Bl -tag -width l3n -offset indent  
915 .It Fl offset Ar width  
916 Indent the display by the  
917 .Ar width ,



```

918 which may be one of the following:
919 .Bl -item
920 .It
921 One of the pre-defined strings
922 .Cm indent ,
923 the width of a standard indentation (six constant width characters);
924 .Cm indent-two ,
925 twice
926 .Cm indent ;
927 .Cm left ,
928 which has no effect;
929 .Cm right ,
930 which justifies to the right margin; or
931 .Cm center ,
932 which aligns around an imagined centre axis.
933 .It
934 A macro invocation, which selects a predefined width
935 associated with that macro.
936 The most popular is the imaginary macro
937 .Ar \&Ds ,
938 which resolves to
939 .Sy 6n .
940 .It
941 A width using the syntax described in
942 .Sx Scaling Widths .
943 .It
944 An arbitrary string, which indents by the length of this string.
945 .El
946 .Pp
947 When the argument is missing,
948 .Fl offset
949 is ignored.
950 .It Fl compact
951 Do not assert vertical space before the display.
952 .El
953 .Pp
954 Examples:
955 .Bd -literal -offset indent
956 \&.Bd \-literal \-offset indent \-compact
957 Hello world.
958 \&.Ed
959 .Ed
960 .Pp
961 See also
962 .Sx \&Dl
963 and
964 .Sx \&Dl .
965 .Ss \&Bf
966 Change the font mode for a scoped block of text.
967 Its syntax is as follows:
968 .Bd -ragged -offset indent
969 .Pf \. Sx \&Bf
970 .Oo
971 .Fl emphasis | literal | symbolic |
972 .Cm \&Em | \&Li | \&Sy
973 .Oc
974 .Ed
975 .Pp
976 The
977 .Fl emphasis
978 and
979 .Cm \&Em
980 argument are equivalent, as are
981 .Fl symbolic
982 and
983 .Cm \&Sy ,

```

```

984 and
985 .Fl literal
986 and
987 .Cm \&Li .
988 Without an argument, this macro does nothing.
989 The font mode continues until broken by a new font mode in a nested
990 scope or
991 .Sx \&Ef
992 is encountered.
993 .Pp
994 See also
995 .Sx \&Li ,
996 .Sx \&Ef ;
997 .Sx \&Em ;
998 and
999 .Sx \&Sy .
1000 .Ss \&Bk
1001 For each macro, keep its output together on the same output line,
1002 until the end of the macro or the end of the input line is reached,
1003 whichever comes first.
1004 Line breaks in text lines are unaffected.
1005 The syntax is as follows:
1006 .Pp
1007 .Dl Pf \. Sx \&Bk Fl words
1008 .Pp
1009 The
1010 .Fl words
1011 argument is required; additional arguments are ignored.
1012 .Pp
1013 The following example will not break within each
1014 .Sx \&Op
1015 macro line:
1016 .Bd -literal -offset indent
1017 \&.Bk \-words
1018 \&.Op Fl f Ar flags
1019 \&.Op Fl o Ar output
1020 \&.Ek
1021 .Ed
1022 .Pp
1023 Be careful in using over-long lines within a keep block!
1024 Doing so will clobber the right margin.
1025 .Ss \&Bl
1026 Begin a list.
1027 Lists consist of items specified using the
1028 .Sx \&It
1029 macro, containing a head or a body or both.
1030 The list syntax is as follows:
1031 .Bd -ragged -offset indent
1032 .Pf \. Sx \&Bl
1033 .Fl Ns Ar type
1034 .Op Fl width Ar val
1035 .Op Fl offset Ar val
1036 .Op Fl compact
1037 .Op HEAD ...
1038 .Ed
1039 .Pp
1040 The list
1041 .Ar type
1042 is mandatory and must be specified first.
1043 The
1044 .Fl width
1045 and
1046 .Fl offset
1047 arguments accept
1048 .Sx Scaling Widths
1049 or use the length of the given string.

```

1050 The  
 1051 .Fl offset  
 1052 is a global indentation for the whole list, affecting both item heads  
 1053 and bodies.  
 1054 For those list types supporting it, the  
 1055 .Fl width  
 1056 argument requests an additional indentation of item bodies,  
 1057 to be added to the  
 1058 .Fl offset .  
 1059 Unless the  
 1060 .Fl compact  
 1061 argument is specified, list entries are separated by vertical space.  
 1062 .Pp  
 1063 A list must specify one of the following list types:  
 1064 .Bl -tag -width l2n -offset indent  
 1065 .It Fl bullet  
 1066 No item heads can be specified, but a bullet will be printed at the head  
 1067 of each item.  
 1068 Item bodies start on the same output line as the bullet  
 1069 and are indented according to the  
 1070 .Fl width  
 1071 argument.  
 1072 .It Fl column  
 1073 A columnated list.  
 1074 The  
 1075 .Fl width  
 1076 argument has no effect; instead, each argument specifies the width  
 1077 of one column, using either the  
 1078 .Sx Scaling Widths  
 1079 syntax or the string length of the argument.  
 1080 If the first line of the body of a  
 1081 .Fl column  
 1082 list is not an  
 1083 .Sx \&It  
 1084 macro line,  
 1085 .Sx \&It  
 1086 contexts spanning one input line each are implied until an  
 1087 .Sx \&It  
 1088 macro line is encountered, at which point items start being interpreted as  
 1089 described in the  
 1090 .Sx \&It  
 1091 documentation.  
 1092 .It Fl dash  
 1093 Like  
 1094 .Fl bullet ,  
 1095 except that dashes are used in place of bullets.  
 1096 .It Fl diag  
 1097 Like  
 1098 .Fl inset ,  
 1099 except that item heads are not parsed for macro invocations.  
 1100 Most often used in the  
 1101 .Em DIAGNOSTICS  
 1102 section with error constants in the item heads.  
 1103 .It Fl enum  
 1104 A numbered list.  
 1105 No item heads can be specified.  
 1106 Formatted like  
 1107 .Fl bullet ,  
 1108 except that cardinal numbers are used in place of bullets,  
 1109 starting at 1.  
 1110 .It Fl hang  
 1111 Like  
 1112 .Fl tag ,  
 1113 except that the first lines of item bodies are not indented, but follow  
 1114 the item heads like in  
 1115 .Fl inset

1116 lists.  
 1117 .It Fl hyphen  
 1118 Synonym for  
 1119 .Fl dash .  
 1120 .It Fl inset  
 1121 Item bodies follow items heads on the same line, using normal inter-word  
 1122 spacing.  
 1123 Bodies are not indented, and the  
 1124 .Fl width  
 1125 argument is ignored.  
 1126 .It Fl item  
 1127 No item heads can be specified, and none are printed.  
 1128 Bodies are not indented, and the  
 1129 .Fl width  
 1130 argument is ignored.  
 1131 .It Fl ohang  
 1132 Item bodies start on the line following item heads and are not indented.  
 1133 The  
 1134 .Fl width  
 1135 argument is ignored.  
 1136 .It Fl tag  
 1137 Item bodies are indented according to the  
 1138 .Fl width  
 1139 argument.  
 1140 When an item head fits inside the indentation, the item body follows  
 1141 this head on the same output line.  
 1142 Otherwise, the body starts on the output line following the head.  
 1143 .El  
 1144 .Pp  
 1145 Lists may be nested within lists and displays.  
 1146 Nesting of  
 1147 .Fl column  
 1148 and  
 1149 .Fl enum  
 1150 lists may not be portable.  
 1151 .Pp  
 1152 See also  
 1153 .Sx \&El  
 1154 and  
 1155 .Sx \&It .  
 1156 .Ss \&Bo  
 1157 Begin a block enclosed by square brackets.  
 1158 Does not have any head arguments.  
 1159 .Pp  
 1160 Examples:  
 1161 .Bd -literal -offset indent -compact  
 1162 \&.Bo l ,  
 1163 \&.Dv BUFSIZ \&Bc  
 1164 .Ed  
 1165 .Pp  
 1166 See also  
 1167 .Sx \&Bq .  
 1168 .Ss \&Bq  
 1169 Encloses its arguments in square brackets.  
 1170 .Pp  
 1171 Examples:  
 1172 .Dl \&.Bq l , \&Dv BUFSIZ  
 1173 .Pp  
 1174 .Em Remarks :  
 1175 this macro is sometimes abused to emulate optional arguments for  
 1176 commands; the correct macros to use for this purpose are  
 1177 .Sx \&Op ,  
 1178 .Sx \&Oo ,  
 1179 and  
 1180 .Sx \&Oc .  
 1181 .Pp

1182 See also  
 1183 .Sx \&Bo .  
 1184 .Ss \&Brc  
 1185 Close a  
 1186 .Sx \&Bro  
 1187 block.  
 1188 Does not have any tail arguments.  
 1189 .Ss \&Bro  
 1190 Begin a block enclosed by curly braces.  
 1191 Does not have any head arguments.  
 1192 .Pp  
 1193 Examples:  
 1194 .Bd -literal -offset indent -compact  
 1195 \&.Bro l , ... ,  
 1196 \&.Va n \&Brc  
 1197 .Ed  
 1198 .Pp  
 1199 See also  
 1200 .Sx \&Brq .  
 1201 .Ss \&Brq  
 1202 Encloses its arguments in curly braces.  
 1203 .Pp  
 1204 Examples:  
 1205 .Dl \&.Brq l , ... , \&Va n  
 1206 .Pp  
 1207 See also  
 1208 .Sx \&Bro .  
 1209 .Ss \&Bsx  
 1210 Format the BSD/OS version provided as an argument, or a default value if  
 1211 no argument is provided.  
 1212 .Pp  
 1213 Examples:  
 1214 .Dl \&.Bsx 1.0  
 1215 .Dl \&.Bsx  
 1216 .Pp  
 1217 See also  
 1218 .Sx \&At ,  
 1219 .Sx \&Bx ,  
 1220 .Sx \&Dx ,  
 1221 .Sx \&Fx ,  
 1222 .Sx \&Nx ,  
 1223 .Sx \&Ox ,  
 1224 and  
 1225 .Sx \&Ux .  
 1226 .Ss \&Bt  
 1227 Prints  
 1228 .Dq is currently in beta test.  
 1229 .Ss \&Bx  
 1230 Format the BSD version provided as an argument, or a default value if no  
 1231 argument is provided.  
 1232 .Pp  
 1233 Examples:  
 1234 .Dl \&.Bx 4.3 Tahoe  
 1235 .Dl \&.Bx 4.4  
 1236 .Dl \&.Bx  
 1237 .Pp  
 1238 See also  
 1239 .Sx \&At ,  
 1240 .Sx \&Bsx ,  
 1241 .Sx \&Dx ,  
 1242 .Sx \&Fx ,  
 1243 .Sx \&Nx ,  
 1244 .Sx \&Ox ,  
 1245 and  
 1246 .Sx \&Ux .  
 1247 .Ss \&Cd

1248 Kernel configuration declaration.  
 1249 This denotes strings accepted by  
 1250 .Xr config 8 .  
 1251 It is most often used in section 4 manual pages.  
 1252 .Pp  
 1253 Examples:  
 1254 .Dl \&.Cd device le0 at scode?  
 1255 .Pp  
 1256 .Em Remarks :  
 1257 this macro is commonly abused by using quoted literals to retain  
 1258 whitespace and align consecutive  
 1259 .Sx \&Cd  
 1260 declarations.  
 1261 This practise is discouraged.  
 1262 .Ss \&Cm  
 1263 Command modifiers.  
 1264 Typically used for fixed strings passed as arguments, unless  
 1265 .Sx \&Fl  
 1266 is more appropriate.  
 1267 Also useful when specifying configuration options or keys.  
 1268 .Pp  
 1269 Examples:  
 1270 .Dl ".Nm mt Fl f Ar device Cm rewind"  
 1271 .Dl ".Nm ps Fl o Cm pid , Ns Cm command"  
 1272 .Dl ".Nm dd Cm if= Ns Ar file1 Cm of= Ns Ar file2"  
 1273 .Dl ".Cm IdentityFile Pa ~/.ssh/id\_rsa"  
 1274 .Dl ".Cm LogLevel Dv DEBUG"  
 1275 .Ss \&Dl  
 1276 One-line indented display.  
 1277 This is formatted by the default rules and is useful for simple indented  
 1278 statements.  
 1279 It is followed by a newline.  
 1280 .Pp  
 1281 Examples:  
 1282 .Dl \&.Dl \&Fl abcdefgh  
 1283 .Pp  
 1284 See also  
 1285 .Sx \&Bd  
 1286 and  
 1287 .Sx \&Dl .  
 1288 .Ss \&Db  
 1289 Switch debugging mode.  
 1290 Its syntax is as follows:  
 1291 .Pp  
 1292 .Dl Pf \. Sx \&Db Cm on | off  
 1293 .Pp  
 1294 This macro is ignored by  
 1295 .Xr mandoc 1 .  
 1296 .Ss \&Dc  
 1297 Close a  
 1298 .Sx \&Do  
 1299 block.  
 1300 Does not have any tail arguments.  
 1301 .Ss \&Dd  
 1302 Document date.  
 1303 This is the mandatory first macro of any  
 1304 .Nm  
 1305 manual.  
 1306 Its syntax is as follows:  
 1307 .Pp  
 1308 .Dl Pf \. Sx \&Dd Ar month day , year  
 1309 .Pp  
 1310 The  
 1311 .Ar month  
 1312 is the full English month name, the  
 1313 .Ar day

1314 is an optionally zero-padded numeral, and the  
 1315 .Ar year  
 1316 is the full four-digit year.  
 1317 .Pp  
 1318 Other arguments are not portable; the  
 1319 .Xr mandoc 1  
 1320 utility handles them as follows:  
 1321 .Bl -dash -offset 3n -compact  
 1322 .It  
 1323 To have the date automatically filled in by the  
 1324 .Ox  
 1325 version of  
 1326 .Xr cvs 1 ,  
 1327 the special string  
 1328 .Dq \$&Mdocdate\$  
 1329 can be given as an argument.  
 1330 .It  
 1331 A few alternative date formats are accepted as well  
 1332 and converted to the standard form.  
 1333 .It  
 1334 If a date string cannot be parsed, it is used verbatim.  
 1335 .It  
 1336 If no date string is given, the current date is used.  
 1337 .El  
 1338 .Pp  
 1339 Examples:  
 1340 .Dl \&.Dd \$&Mdocdate\$  
 1341 .Dl \&.Dd \$&Mdocdate: July 21 2007\$  
 1342 .Dl \&.Dd July 21, 2007  
 1343 .Pp  
 1344 See also  
 1345 .Sx \&Dt  
 1346 and  
 1347 .Sx \&Os .  
 1348 .Ss \&Dl  
 1349 One-line intended display.  
 1350 This is formatted as literal text and is useful for commands and  
 1351 invocations.  
 1352 It is followed by a newline.  
 1353 .Pp  
 1354 Examples:  
 1355 .Dl \&.Dl % mandoc mdoc.5 \e(ba less  
 1356 .Pp  
 1357 See also  
 1358 .Sx \&Bd  
 1359 and  
 1360 .Sx \&Dl .  
 1361 .Ss \&Do  
 1362 Begin a block enclosed by double quotes.  
 1363 Does not have any head arguments.  
 1364 .Pp  
 1365 Examples:  
 1366 .Bd -literal -offset indent -compact  
 1367 \&.Do  
 1368 April is the cruellest month  
 1369 \&.Dc  
 1370 \e(em T.S. Eliot  
 1371 .Ed  
 1372 .Pp  
 1373 See also  
 1374 .Sx \&Dq .  
 1375 .Ss \&Dq  
 1376 Encloses its arguments in  
 1377 .Dq typographic  
 1378 double-quotes.  
 1379 .Pp

1380 Examples:  
 1381 .Bd -literal -offset indent -compact  
 1382 \&.Dq April is the cruellest month  
 1383 \e(em T.S. Eliot  
 1384 .Ed  
 1385 .Pp  
 1386 See also  
 1387 .Sx \&Qq ,  
 1388 .Sx \&Sq ,  
 1389 and  
 1390 .Sx \&Do .  
 1391 .Ss \&Dt  
 1392 Document title.  
 1393 This is the mandatory second macro of any  
 1394 .Nm  
 1395 file.  
 1396 Its syntax is as follows:  
 1397 .Bd -ragged -offset indent  
 1398 .Pf \. Sx \&Dt  
 1399 .Oo  
 1400 .Ar title  
 1401 .Oo  
 1402 .Ar section  
 1403 .Op Ar volume  
 1404 .Op Ar arch  
 1405 .Oc  
 1406 .Oc  
 1407 .Ed  
 1408 .Pp  
 1409 Its arguments are as follows:  
 1410 .Bl -tag -width Ds -offset Ds  
 1411 .It Ar title  
 1412 The document's title (name), defaulting to  
 1413 .Dq UNKNOWN  
 1414 if unspecified.  
 1415 It should be capitalised.  
 1416 .It Ar section  
 1417 The manual section. It should correspond to the manual's filename suffix  
 1418 and defaults to  
 1419 .Dq 1  
 1420 if unspecified.  
 1421 .It Ar volume  
 1422 This overrides the volume inferred from  
 1423 .Ar section .  
 1424 This field is optional.  
 1425 .It Ar arch  
 1426 This specifies the machine architecture a manual page applies to,  
 1427 where relevant.  
 1428 .El  
 1429 .Ss \&Dv  
 1430 Defined variables such as preprocessor constants, constant symbols,  
 1431 enumeration values, and so on.  
 1432 .Pp  
 1433 Examples:  
 1434 .Dl \&.Dv NULL  
 1435 .Dl \&.Dv BUFSIZ  
 1436 .Dl \&.Dv STDOUT\_FILENO  
 1437 .Pp  
 1438 See also  
 1439 .Sx \&Er  
 1440 and  
 1441 .Sx \&Ev  
 1442 for special-purpose constants and  
 1443 .Sx \&Va  
 1444 for variable symbols.  
 1445 .Ss \&Dx

1446 Format the DragonFly BSD version provided as an argument, or a default  
 1447 value if no argument is provided.  
 1448 .Pp  
 1449 Examples:  
 1450 .Dl \&.Dx 2.4.1  
 1451 .Dl \&.Dx  
 1452 .Pp  
 1453 See also  
 1454 .Sx \&At ,  
 1455 .Sx \&Bsx ,  
 1456 .Sx \&Bx ,  
 1457 .Sx \&Fx ,  
 1458 .Sx \&Nx ,  
 1459 .Sx \&Ox ,  
 1460 and  
 1461 .Sx \&Ux .  
 1462 .Ss \&Ec  
 1463 Close a scope started by  
 1464 .Sx \&Eo .  
 1465 Its syntax is as follows:  
 1466 .Pp  
 1467 .Dl Pf \. Sx \&Ec Op Ar TERM  
 1468 .Pp  
 1469 The  
 1470 .Ar TERM  
 1471 argument is used as the enclosure tail, for example, specifying \e(rq  
 1472 will emulate  
 1473 .Sx \&Dc .  
 1474 .Ss \&Ed  
 1475 End a display context started by  
 1476 .Sx \&Bd .  
 1477 .Ss \&Ef  
 1478 End a font mode context started by  
 1479 .Sx \&Bf .  
 1480 .Ss \&Ek  
 1481 End a keep context started by  
 1482 .Sx \&Bk .  
 1483 .Ss \&El  
 1484 End a list context started by  
 1485 .Sx \&Bl .  
 1486 .Pp  
 1487 See also  
 1488 .Sx \&Bl  
 1489 and  
 1490 .Sx \&It .  
 1491 .Ss \&Em  
 1492 Denotes text that should be  
 1493 .Em emphasised .  
 1494 Note that this is a presentation term and should not be used for  
 1495 stylistically decorating technical terms.  
 1496 Depending on the output device, this is usually represented  
 1497 using an italic font or underlined characters.  
 1498 .Pp  
 1499 Examples:  
 1500 .Dl \&.Em Warnings!  
 1501 .Dl \&.Em Remarks :  
 1502 .Pp  
 1503 See also  
 1504 .Sx \&Bf ,  
 1505 .Sx \&Li ,  
 1506 .Sx \&No ,  
 1507 and  
 1508 .Sx \&Sy .  
 1509 .Ss \&En  
 1510 This macro is obsolete and not implemented in  
 1511 .Xr mandoc 1 .

1512 .Ss \&Eo  
 1513 An arbitrary enclosure.  
 1514 Its syntax is as follows:  
 1515 .Pp  
 1516 .Dl Pf \. Sx \&Eo Op Ar TERM  
 1517 .Pp  
 1518 The  
 1519 .Ar TERM  
 1520 argument is used as the enclosure head, for example, specifying \e(lq  
 1521 will emulate  
 1522 .Sx \&Do .  
 1523 .Ss \&Er  
 1524 Error constants for definitions of the  
 1525 .Va errno  
 1526 libc global variable.  
 1527 This is most often used in section 2 and 3 manual pages.  
 1528 .Pp  
 1529 Examples:  
 1530 .Dl \&.Er EPERM  
 1531 .Dl \&.Er ENOENT  
 1532 .Pp  
 1533 See also  
 1534 .Sx \&Dv  
 1535 for general constants.  
 1536 .Ss \&Es  
 1537 This macro is obsolete and not implemented.  
 1538 .Ss \&Ev  
 1539 Environmental variables such as those specified in  
 1540 .Xr environ 5 .  
 1541 .Pp  
 1542 Examples:  
 1543 .Dl \&.Ev DISPLAY  
 1544 .Dl \&.Ev PATH  
 1545 .Pp  
 1546 See also  
 1547 .Sx \&Dv  
 1548 for general constants.  
 1549 .Ss \&Ex  
 1550 Insert a standard sentence regarding command exit values of 0 on success  
 1551 and >0 on failure.  
 1552 This is most often used in section 1, 6, and 8 manual pages.  
 1553 Its syntax is as follows:  
 1554 .Pp  
 1555 .Dl Pf \. Sx \&Ex Fl std Op Ar utility ...  
 1556 .Pp  
 1557 If  
 1558 .Ar utility  
 1559 is not specified, the document's name set by  
 1560 .Sx \&Nm  
 1561 is used.  
 1562 Multiple  
 1563 .Ar utility  
 1564 arguments are treated as separate utilities.  
 1565 .Pp  
 1566 See also  
 1567 .Sx \&Rv .  
 1568 .Ss \&Fa  
 1569 Function argument.  
 1570 Its syntax is as follows:  
 1571 .Bd -ragged -offset indent  
 1572 .Pf \. Sx \&Fa  
 1573 .Op Cm argtype  
 1574 .Cm argname  
 1575 .Ed  
 1576 .Pp  
 1577 This may be invoked for names with or without the corresponding type.

1578 It is also used to specify the field name of a structure.  
 1579 Most often, the  
 1580 .Sx \&Fa  
 1581 macro is used in the  
 1582 .Em SYNOPSIS  
 1583 within  
 1584 .Sx \&Fo  
 1585 section when documenting multi-line function prototypes.  
 1586 If invoked with multiple arguments, the arguments are separated by a  
 1587 comma.  
 1588 Furthermore, if the following macro is another  
 1589 .Sx \&Fa ,  
 1590 the last argument will also have a trailing comma.  
 1591 .Pp  
 1592 Examples:  
 1593 .Dl \&Fa \((dqconst char \*p\((dq  
 1594 .Dl \&Fa \((dqint a\((dq \((dqint b\((dq \((dqint c\((dq  
 1595 .Dl \&Fa foo  
 1596 .Pp  
 1597 See also  
 1598 .Sx \&Fo .  
 1599 .Ss \&Fc  
 1600 End a function context started by  
 1601 .Sx \&Fo .  
 1602 .Ss \&Fd  
 1603 Historically used to document include files.  
 1604 This usage has been deprecated in favour of  
 1605 .Sx \&In .  
 1606 Do not use this macro.  
 1607 .Pp  
 1608 See also  
 1609 .Sx MANUAL STRUCTURE  
 1610 and  
 1611 .Sx \&In .  
 1612 .Ss \&Fl  
 1613 Command-line flag or option.  
 1614 Used when listing arguments to command-line utilities.  
 1615 Prints a fixed-width hyphen  
 1616 .Sq \-  
 1617 directly followed by each argument.  
 1618 If no arguments are provided, a hyphen is printed followed by a space.  
 1619 If the argument is a macro, a hyphen is prefixed to the subsequent macro  
 1620 output.  
 1621 .Pp  
 1622 Examples:  
 1623 .Dl ".Fl R Op Fl H | L | P"  
 1624 .Dl ".Op Fl lAaCcdFfgHhikLllmnopqRrSsTtux"  
 1625 .Dl ".Fl type Cm d Fl name Pa CVS"  
 1626 .Dl ".Fl Ar signal\_number"  
 1627 .Dl ".Fl o Fl"  
 1628 .Pp  
 1629 See also  
 1630 .Sx \&Cm .  
 1631 .Ss \&Fn  
 1632 A function name.  
 1633 Its syntax is as follows:  
 1634 .Bd -ragged -offset indent  
 1635 .Pf \. Ns Sx \&Fn  
 1636 .Op Ar functype  
 1637 .Ar funcname  
 1638 .Op Oo Ar argtype Oc Ar argname  
 1639 .Ed  
 1640 .Pp  
 1641 Function arguments are surrounded in parenthesis and  
 1642 are delimited by commas.  
 1643 If no arguments are specified, blank parenthesis are output.

1644 In the  
 1645 .Em SYNOPSIS  
 1646 section, this macro starts a new output line,  
 1647 and a blank line is automatically inserted between function definitions.  
 1648 .Pp  
 1649 Examples:  
 1650 .Dl \&Fn \((dqint funcname\((dq \((dqint arg0\((dq \((dqint arg1\((dq  
 1651 .Dl \&Fn funcname \((dqint arg0\((dq  
 1652 .Dl \&Fn funcname arg0  
 1653 .Pp  
 1654 .Bd -literal -offset indent -compact  
 1655 \&Ft functype  
 1656 \&Fn funcname  
 1657 .Ed  
 1658 .Pp  
 1659 When referring to a function documented in another manual page, use  
 1660 .Sx \&Xr  
 1661 instead.  
 1662 See also  
 1663 .Sx MANUAL STRUCTURE ,  
 1664 .Sx \&Fo ,  
 1665 and  
 1666 .Sx \&Ft .  
 1667 .Ss \&Fo  
 1668 Begin a function block.  
 1669 This is a multi-line version of  
 1670 .Sx \&Fn .  
 1671 Its syntax is as follows:  
 1672 .Pp  
 1673 .Dl Pf \. Sx \&Fo Ar funcname  
 1674 .Pp  
 1675 Invocations usually occur in the following context:  
 1676 .Bd -ragged -offset indent  
 1677 .Pf \. Sx \&Ft Ar functype  
 1678 .br  
 1679 .Pf \. Sx \&Fo Ar funcname  
 1680 .br  
 1681 .Pf \. Sx \&Fa Oo Ar argtype Oc Ar argname  
 1682 .br  
 1683 \&.\.\  
 1684 .br  
 1685 .Pf \. Sx \&Fc  
 1686 .Ed  
 1687 .Pp  
 1688 A  
 1689 .Sx \&Fo  
 1690 scope is closed by  
 1691 .Sx \&Fc .  
 1692 .Pp  
 1693 See also  
 1694 .Sx MANUAL STRUCTURE ,  
 1695 .Sx \&Fa ,  
 1696 .Sx \&Fc ,  
 1697 and  
 1698 .Sx \&Ft .  
 1699 .Ss \&Fr  
 1700 This macro is obsolete and not implemented in  
 1701 .Xr mandoc 1 .  
 1702 .Pp  
 1703 It was used to show function return values.  
 1704 The syntax was:  
 1705 .Pp  
 1706 .Dl Pf . Sx \&Fr Ar value  
 1707 .Ss \&Ft  
 1708 A function type.  
 1709 Its syntax is as follows:

1710 .Pp  
 1711 .Dl Pf \. Sx \&Ft Ar functype  
 1712 .Pp  
 1713 In the  
 1714 .Em SYNOPSIS  
 1715 section, a new output line is started after this macro.  
 1716 .Pp  
 1717 Examples:  
 1718 .Dl \&.Ft int  
 1719 .Bd -literal -offset indent -compact  
 1720 \&.Ft functype  
 1721 \&.Fn funcname  
 1722 .Ed  
 1723 .Pp  
 1724 See also  
 1725 .Sx MANUAL STRUCTURE ,  
 1726 .Sx \&Fn ,  
 1727 and  
 1728 .Sx \&Fo .  
 1729 .Ss \&Fx  
 1730 Format the  
 1731 .Fx  
 1732 version provided as an argument, or a default value  
 1733 if no argument is provided.  
 1734 .Pp  
 1735 Examples:  
 1736 .Dl \&.Fx 7.1  
 1737 .Dl \&.Fx  
 1738 .Pp  
 1739 See also  
 1740 .Sx \&At ,  
 1741 .Sx \&Bsx ,  
 1742 .Sx \&Bx ,  
 1743 .Sx \&Dx ,  
 1744 .Sx \&Nx ,  
 1745 .Sx \&Ox ,  
 1746 and  
 1747 .Sx \&Ux .  
 1748 .Ss \&Hf  
 1749 This macro is not implemented in  
 1750 .Xr mandoc 1 .  
 1751 .Pp  
 1752 It was used to include the contents of a (header) file literally.  
 1753 The syntax was:  
 1754 .Pp  
 1755 .Dl Pf . Sx \&Hf Ar filename  
 1756 .Ss \&Ic  
 1757 Designate an internal or interactive command.  
 1758 This is similar to  
 1759 .Sx \&Cm  
 1760 but used for instructions rather than values.  
 1761 .Pp  
 1762 Examples:  
 1763 .Dl \&.Ic :wq  
 1764 .Dl \&.Ic hash  
 1765 .Dl \&.Ic alias  
 1766 .Pp  
 1767 Note that using  
 1768 .Sx \&Bd Fl literal  
 1769 or  
 1770 .Sx \&Dl  
 1771 is preferred for displaying code; the  
 1772 .Sx \&Ic  
 1773 macro is used when referring to specific instructions.  
 1774 .Ss \&In  
 1775 An

1776 .Dq include  
 1777 file.  
 1778 When invoked as the first macro on an input line in the  
 1779 .Em SYNOPSIS  
 1780 section, the argument is displayed in angle brackets  
 1781 and preceded by  
 1782 .Dq #include ,  
 1783 and a blank line is inserted in front if there is a preceding  
 1784 function declaration.  
 1785 This is most often used in section 2, 3, and 9 manual pages.  
 1786 .Pp  
 1787 Examples:  
 1788 .Dl \&.In sys/types.h  
 1789 .Pp  
 1790 See also  
 1791 .Sx MANUAL STRUCTURE .  
 1792 .Ss \&It  
 1793 A list item.  
 1794 The syntax of this macro depends on the list type.  
 1795 .Pp  
 1796 Lists  
 1797 of type  
 1798 .Fl hang ,  
 1799 .Fl ohang ,  
 1800 .Fl inset ,  
 1801 and  
 1802 .Fl diag  
 1803 have the following syntax:  
 1804 .Pp  
 1805 .Dl Pf \. Sx \&It Ar args  
 1806 .Pp  
 1807 Lists of type  
 1808 .Fl bullet ,  
 1809 .Fl dash ,  
 1810 .Fl enum ,  
 1811 .Fl hyphen  
 1812 and  
 1813 .Fl item  
 1814 have the following syntax:  
 1815 .Pp  
 1816 .Dl Pf \. Sx \&It  
 1817 .Pp  
 1818 with subsequent lines interpreted within the scope of the  
 1819 .Sx \&It  
 1820 until either a closing  
 1821 .Sx \&El  
 1822 or another  
 1823 .Sx \&It .  
 1824 .Pp  
 1825 The  
 1826 .Fl tag  
 1827 list has the following syntax:  
 1828 .Pp  
 1829 .Dl Pf \. Sx \&It Op Cm args  
 1830 .Pp  
 1831 Subsequent lines are interpreted as with  
 1832 .Fl bullet  
 1833 and family.  
 1834 The line arguments correspond to the list's left-hand side; body  
 1835 arguments correspond to the list's contents.  
 1836 .Pp  
 1837 The  
 1838 .Fl column  
 1839 list is the most complicated.  
 1840 Its syntax is as follows:  
 1841 .Pp

1842 .Dl Pf \. Sx \&It Ar cell Op <TAB> Ar cell ...  
 1843 .Dl Pf \. Sx \&It Ar cell Op Sx \&Ta Ar cell ...  
 1844 .Pp  
 1845 The arguments consist of one or more lines of text and macros  
 1846 representing a complete table line.  
 1847 Cells within the line are delimited by tabs or by the special  
 1848 .Sx \&Ta  
 1849 block macro.  
 1850 The tab cell delimiter may only be used within the  
 1851 .Sx \&It  
 1852 line itself; on following lines, only the  
 1853 .Sx \&Ta  
 1854 macro can be used to delimit cells, and  
 1855 .Sx \&Ta  
 1856 is only recognised as a macro when called by other macros,  
 1857 not as the first macro on a line.  
 1858 .Pp  
 1859 Note that quoted strings may span tab-delimited cells on an  
 1860 .Sx \&It  
 1861 line.  
 1862 For example,  
 1863 .Pp  
 1864 .Dl .It \((dqcoll ; <TAB> col2 ; \((dq \&;  
 1865 .Pp  
 1866 will preserve the semicolon whitespace except for the last.  
 1867 .Pp  
 1868 See also  
 1869 .Sx \&Bl .  
 1870 .Ss \&Lb  
 1871 Specify a library.  
 1872 The syntax is as follows:  
 1873 .Pp  
 1874 .Dl Pf \. Sx \&Lb Ar library  
 1875 .Pp  
 1876 The  
 1877 .Ar library  
 1878 parameter may be a system library, such as  
 1879 .Cm libz  
 1880 or  
 1881 .Cm libpam ,  
 1882 in which case a small library description is printed next to the linker  
 1883 invocation; or a custom library, in which case the library name is  
 1884 printed in quotes.  
 1885 This is most commonly used in the  
 1886 .Em SYNOPSIS  
 1887 section as described in  
 1888 .Sx MANUAL STRUCTURE .  
 1889 .Pp  
 1890 Examples:  
 1891 .Dl \&.Lb libz  
 1892 .Dl \&.Lb mdoc  
 1893 .Ss \&Li  
 1894 Denotes text that should be in a  
 1895 .Li literal  
 1896 font mode.  
 1897 Note that this is a presentation term and should not be used for  
 1898 stylistically decorating technical terms.  
 1899 .Pp  
 1900 On terminal output devices, this is often indistinguishable from  
 1901 normal text.  
 1902 .Pp  
 1903 See also  
 1904 .Sx \&Bf ,  
 1905 .Sx \&Em ,  
 1906 .Sx \&No ,  
 1907 and

1908 .Sx \&Sy .  
 1909 .Ss \&Lk  
 1910 Format a hyperlink.  
 1911 Its syntax is as follows:  
 1912 .Pp  
 1913 .Dl Pf \. Sx \&Lk Ar uri Op Ar name  
 1914 .Pp  
 1915 Examples:  
 1916 .Dl \&.Lk http://bsd.lv \((dqThe BSD.lv Project\((dq  
 1917 .Dl \&.Lk http://bsd.lv  
 1918 .Pp  
 1919 See also  
 1920 .Sx \&Mt .  
 1921 .Ss \&Lp  
 1922 Synonym for  
 1923 .Sx \&Pp .  
 1924 .Ss \&Ms  
 1925 Display a mathematical symbol.  
 1926 Its syntax is as follows:  
 1927 .Pp  
 1928 .Dl Pf \. Sx \&Ms Ar symbol  
 1929 .Pp  
 1930 Examples:  
 1931 .Dl \&.Ms sigma  
 1932 .Dl \&.Ms aleph  
 1933 .Ss \&Mt  
 1934 Format a  
 1935 .Dg mailto:  
 1936 hyperlink.  
 1937 Its syntax is as follows:  
 1938 .Pp  
 1939 .Dl Pf \. Sx \&Mt Ar address  
 1940 .Pp  
 1941 Examples:  
 1942 .Dl \&.Mt discuss@manpages.bsd.lv  
 1943 .Ss \&Nd  
 1944 A one line description of the manual's content.  
 1945 This may only be invoked in the  
 1946 .Em SYNOPSIS  
 1947 section subsequent the  
 1948 .Sx \&Nm  
 1949 macro.  
 1950 .Pp  
 1951 Examples:  
 1952 .Dl Pf . Sx \&Nd mdoc language reference  
 1953 .Dl Pf . Sx \&Nd format and display UNIX manuals  
 1954 .Pp  
 1955 The  
 1956 .Sx \&Nd  
 1957 macro technically accepts child macros and terminates with a subsequent  
 1958 .Sx \&Sh  
 1959 invocation.  
 1960 Do not assume this behaviour: some  
 1961 .Xr whatis 1  
 1962 database generators are not smart enough to parse more than the line  
 1963 arguments and will display macros verbatim.  
 1964 .Pp  
 1965 See also  
 1966 .Sx \&Nm .  
 1967 .Ss \&Nm  
 1968 The name of the manual page, or \((em in particular in section 1, 6,  
 1969 and 8 pages \((em of an additional command or feature documented in  
 1970 the manual page.  
 1971 When first invoked, the  
 1972 .Sx \&Nm  
 1973 macro expects a single argument, the name of the manual page.



1974 Usually, the first invocation happens in the  
 1975 .Em NAME  
 1976 section of the page.  
 1977 The specified name will be remembered and used whenever the macro is  
 1978 called again without arguments later in the page.  
 1979 The  
 1980 .Sx \&Nm  
 1981 macro uses  
 1982 .Sx Block full-implicit  
 1983 semantics when invoked as the first macro on an input line in the  
 1984 .Em SYNOPSIS  
 1985 section; otherwise, it uses ordinary  
 1986 .Sx In-line  
 1987 semantics.  
 1988 .Pp  
 1989 Examples:  
 1990 .Bd -literal -offset indent  
 1991 \&.Sh SYNOPSIS  
 1992 \&.Nm cat  
 1993 \&.Op Fl benstuv  
 1994 \&.Op Ar  
 1995 .Ed  
 1996 .Pp  
 1997 In the  
 1998 .Em SYNOPSIS  
 1999 of section 2, 3 and 9 manual pages, use the  
 2000 .Sx \&Fn  
 2001 macro rather than  
 2002 .Sx \&Nm  
 2003 to mark up the name of the manual page.  
 2004 .Ss \&No  
 2005 Normal text.  
 2006 Closes the scope of any preceding in-line macro.  
 2007 When used after physical formatting macros like  
 2008 .Sx \&Em  
 2009 or  
 2010 .Sx \&Sy ,  
 2011 switches back to the standard font face and weight.  
 2012 Can also be used to embed plain text strings in macro lines  
 2013 using semantic annotation macros.  
 2014 .Pp  
 2015 Examples:  
 2016 .Dl ".Em italic , Sy bold , No and roman"  
 2017 .Pp  
 2018 .Bd -literal -offset indent -compact  
 2019 \&.Sm off  
 2020 \&.Cm :C No / Ar pattern No / Ar replacement No /  
 2021 \&.Sm on  
 2022 .Ed  
 2023 .Pp  
 2024 See also  
 2025 .Sx \&Em ,  
 2026 .Sx \&Li ,  
 2027 and  
 2028 .Sx \&Sy .  
 2029 .Ss \&Ns  
 2030 Suppress a space between the output of the preceding macro  
 2031 and the following text or macro.  
 2032 Following invocation, input is interpreted as normal text  
 2033 just like after an  
 2034 .Sx \&No  
 2035 macro.  
 2036 .Pp  
 2037 This has no effect when invoked at the start of a macro line.  
 2038 .Pp  
 2039 Examples:

2040 .Dl ".Ar name Ns = Ns Ar value"  
 2041 .Dl ".Cm :M Ns Ar pattern"  
 2042 .Dl ".Fl o Ns Ar output"  
 2043 .Pp  
 2044 See also  
 2045 .Sx \&No  
 2046 and  
 2047 .Sx \&Sm .  
 2048 .Ss \&Nx  
 2049 Format the  
 2050 .Nx  
 2051 version provided as an argument, or a default value if  
 2052 no argument is provided.  
 2053 .Pp  
 2054 Examples:  
 2055 .Dl \&.Nx 5.01  
 2056 .Dl \&.Nx  
 2057 .Pp  
 2058 See also  
 2059 .Sx \&At ,  
 2060 .Sx \&Bsx ,  
 2061 .Sx \&Bx ,  
 2062 .Sx \&Dx ,  
 2063 .Sx \&Fx ,  
 2064 .Sx \&Ox ,  
 2065 and  
 2066 .Sx \&Ux .  
 2067 .Ss \&Oc  
 2068 Close multi-line  
 2069 .Sx \&Oo  
 2070 context.  
 2071 .Ss \&Oo  
 2072 Multi-line version of  
 2073 .Sx \&Op .  
 2074 .Pp  
 2075 Examples:  
 2076 .Bd -literal -offset indent -compact  
 2077 \&.Oo  
 2078 \&.Op Fl flag Ns Ar value  
 2079 \&.Oc  
 2080 .Ed  
 2081 .Ss \&Op  
 2082 Optional part of a command line.  
 2083 Prints the argument(s) in brackets.  
 2084 This is most often used in the  
 2085 .Em SYNOPSIS  
 2086 section of section 1 and 8 manual pages.  
 2087 .Pp  
 2088 Examples:  
 2089 .Dl \&.Op \&Fl a \&Ar b  
 2090 .Dl \&.Op \&Ar a | b  
 2091 .Pp  
 2092 See also  
 2093 .Sx \&Oo .  
 2094 .Ss \&Os  
 2095 Document operating system version.  
 2096 This is the mandatory third macro of  
 2097 any  
 2098 .Nm  
 2099 file.  
 2100 Its syntax is as follows:  
 2101 .Pp  
 2102 .Dl Pf \. Sx \&Os Op Ar system Op Ar version  
 2103 .Pp  
 2104 The optional  
 2105 .Ar system

2106 parameter specifies the relevant operating system or environment.  
 2107 Left unspecified, it defaults to the local operating system version.  
 2108 This is the suggested form.  
 2109 .Pp  
 2110 Examples:  
 2111 .Dl \&.Os  
 2112 .Dl \&.Os KTH/CSC/TCS  
 2113 .Dl \&.Os BSD 4.3  
 2114 .Pp  
 2115 See also  
 2116 .Sx \&Dd  
 2117 and  
 2118 .Sx \&Dt .  
 2119 .Ss \&Ot  
 2120 This macro is obsolete and not implemented in  
 2121 .Xr mandoc 1 .  
 2122 .Pp  
 2123 Historical  
 2124 .Xr mdoc 5  
 2125 packages described it as  
 2126 .Dq "old function type (FORTRAN)" .  
 2127 .Ss \&Ox  
 2128 Format the  
 2129 .Ox  
 2130 version provided as an argument, or a default value  
 2131 if no argument is provided.  
 2132 .Pp  
 2133 Examples:  
 2134 .Dl \&.Ox 4.5  
 2135 .Dl \&.Ox  
 2136 .Pp  
 2137 See also  
 2138 .Sx \&At ,  
 2139 .Sx \&BSx ,  
 2140 .Sx \&Bx ,  
 2141 .Sx \&Dx ,  
 2142 .Sx \&Fx ,  
 2143 .Sx \&Nx ,  
 2144 and  
 2145 .Sx \&Ux .  
 2146 .Ss \&Pa  
 2147 An absolute or relative file system path, or a file or directory name.  
 2148 If an argument is not provided, the character  
 2149 .Sq \ (ti  
 2150 is used as a default.  
 2151 .Pp  
 2152 Examples:  
 2153 .Dl \&.Pa /usr/bin/mandoc  
 2154 .Dl \&.Pa /usr/share/man/man5/mdoc.5  
 2155 .Pp  
 2156 See also  
 2157 .Sx \&Lk .  
 2158 .Ss \&Pc  
 2159 Close parenthesised context opened by  
 2160 .Sx \&Po .  
 2161 .Ss \&Pf  
 2162 Removes the space between its argument  
 2163 .Pq Dq prefix  
 2164 and the following macro.  
 2165 Its syntax is as follows:  
 2166 .Pp  
 2167 .Dl .Pf Ar prefix macro arguments ...  
 2168 .Pp  
 2169 This is equivalent to:  
 2170 .Pp  
 2171 .Dl .No Ar prefix No \&Ns Ar macro arguments ...

2172 .Pp  
 2173 Examples:  
 2174 .Dl ".Pf \$ Ar variable\_name"  
 2175 .Dl ".Pf 0x Ar hex\_digits"  
 2176 .Pp  
 2177 See also  
 2178 .Sx \&Ns  
 2179 and  
 2180 .Sx \&Sm .  
 2181 .Ss \&Po  
 2182 Multi-line version of  
 2183 .Sx \&Pq .  
 2184 .Ss \&Pp  
 2185 Break a paragraph.  
 2186 This will assert vertical space between prior and subsequent macros  
 2187 and/or text.  
 2188 .Pp  
 2189 Paragraph breaks are not needed before or after  
 2190 .Sx \&Sh  
 2191 or  
 2192 .Sx \&Ss  
 2193 macros or before displays  
 2194 .Pq Sx \&Bd  
 2195 or lists  
 2196 .Pq Sx \&Bl  
 2197 unless the  
 2198 .Fl compact  
 2199 flag is given.  
 2200 .Ss \&Pq  
 2201 Parenthesised enclosure.  
 2202 .Pp  
 2203 See also  
 2204 .Sx \&Po .  
 2205 .Ss \&Qc  
 2206 Close quoted context opened by  
 2207 .Sx \&Qo .  
 2208 .Ss \&Ql  
 2209 Format a single-quoted literal.  
 2210 See also  
 2211 .Sx \&Qq  
 2212 and  
 2213 .Sx \&Sq .  
 2214 .Ss \&Qo  
 2215 Multi-line version of  
 2216 .Sx \&Qq .  
 2217 .Ss \&Qq  
 2218 Encloses its arguments in  
 2219 .Qq typewriter  
 2220 double-quotes.  
 2221 Consider using  
 2222 .Sx \&Dq .  
 2223 .Pp  
 2224 See also  
 2225 .Sx \&Dq ,  
 2226 .Sx \&Sq ,  
 2227 and  
 2228 .Sx \&Qo .  
 2229 .Ss \&Re  
 2230 Close an  
 2231 .Sx \&Rs  
 2232 block.  
 2233 Does not have any tail arguments.  
 2234 .Ss \&Rs  
 2235 Begin a bibliographic  
 2236 .Pq Dq reference  
 2237 block.

2238 Does not have any head arguments.  
 2239 The block macro may only contain  
 2240 .Sx \&%A ,  
 2241 .Sx \&%B ,  
 2242 .Sx \&%C ,  
 2243 .Sx \&%D ,  
 2244 .Sx \&%I ,  
 2245 .Sx \&%J ,  
 2246 .Sx \&%N ,  
 2247 .Sx \&%O ,  
 2248 .Sx \&%P ,  
 2249 .Sx \&%Q ,  
 2250 .Sx \&%R ,  
 2251 .Sx \&%T ,  
 2252 .Sx \&%U ,  
 2253 and  
 2254 .Sx \&%V  
 2255 child macros (at least one must be specified).  
 2256 .Pp  
 2257 Examples:  
 2258 .Bd -literal -offset indent -compact  
 2259 \&.Rs  
 2260 \&.%A J. E. Hopcroft  
 2261 \&.%A J. D. Ullman  
 2262 \&.%B Introduction to Automata Theory, Languages, and Computation  
 2263 \&.%I Addison-Wesley  
 2264 \&.%C Reading, Massachusetts  
 2265 \&.%D 1979  
 2266 \&.Re  
 2267 .Ed  
 2268 .Pp  
 2269 If an  
 2270 .Sx \&Rs  
 2271 block is used within a SEE ALSO section, a vertical space is asserted  
 2272 before the rendered output, else the block continues on the current  
 2273 line.  
 2274 .Ss \&Rv  
 2275 Insert a standard sentence regarding a function call's return value of 0  
 2276 on success and \-1 on error, with the  
 2277 .Va errno  
 2278 libc global variable set on error.  
 2279 Its syntax is as follows:  
 2280 .Pp  
 2281 .Dl Pf \. Sx \&Rv Fl std Op Ar function ...  
 2282 .Pp  
 2283 If  
 2284 .Ar function  
 2285 is not specified, the document's name set by  
 2286 .Sx \&Nm  
 2287 is used.  
 2288 Multiple  
 2289 .Ar function  
 2290 arguments are treated as separate functions.  
 2291 .Pp  
 2292 See also  
 2293 .Sx \&Ex .  
 2294 .Ss \&Sc  
 2295 Close single-quoted context opened by  
 2296 .Sx \&So .  
 2297 .Ss \&Sh  
 2298 Begin a new section.  
 2299 For a list of conventional manual sections, see  
 2300 .Sx MANUAL STRUCTURE .  
 2301 These sections should be used unless it's absolutely necessary that  
 2302 custom sections be used.  
 2303 .Pp

2304 Section names should be unique so that they may be keyed by  
 2305 .Sx \&Sx .  
 2306 Although this macro is parsed, it should not consist of child node or it  
 2307 may not be linked with  
 2308 .Sx \&Sx .  
 2309 .Pp  
 2310 See also  
 2311 .Sx \&Pp ,  
 2312 .Sx \&Ss ,  
 2313 and  
 2314 .Sx \&Sx .  
 2315 .Ss \&Sm  
 2316 Switches the spacing mode for output generated from macros.  
 2317 Its syntax is as follows:  
 2318 .Pp  
 2319 .Dl Pf \. Sx \&Sm Cm on | off  
 2320 .Pp  
 2321 By default, spacing is  
 2322 .Cm on .  
 2323 When switched  
 2324 .Cm off ,  
 2325 no white space is inserted between macro arguments and between the  
 2326 output generated from adjacent macros, but text lines  
 2327 still get normal spacing between words and sentences.  
 2328 .Ss \&So  
 2329 Multi-line version of  
 2330 .Sx \&Sq .  
 2331 .Ss \&Sq  
 2332 Encloses its arguments in  
 2333 .Sq typewriter  
 2334 single-quotes.  
 2335 .Pp  
 2336 See also  
 2337 .Sx \&Dq ,  
 2338 .Sx \&Qq ,  
 2339 and  
 2340 .Sx \&So .  
 2341 .Ss \&Ss  
 2342 Begin a new subsection.  
 2343 Unlike with  
 2344 .Sx \&Sh ,  
 2345 there is no convention for the naming of subsections.  
 2346 Except  
 2347 .Em DESCRIPTION ,  
 2348 the conventional sections described in  
 2349 .Sx MANUAL STRUCTURE  
 2350 rarely have subsections.  
 2351 .Pp  
 2352 Sub-section names should be unique so that they may be keyed by  
 2353 .Sx \&Sx .  
 2354 Although this macro is parsed, it should not consist of child node or it  
 2355 may not be linked with  
 2356 .Sx \&Sx .  
 2357 .Pp  
 2358 See also  
 2359 .Sx \&Pp ,  
 2360 .Sx \&Sh ,  
 2361 and  
 2362 .Sx \&Sx .  
 2363 .Ss \&St  
 2364 Replace an abbreviation for a standard with the full form.  
 2365 The following standards are recognised:  
 2366 .Pp  
 2367 .Bl -tag -width "-p1003.1g-2000X" -compact  
 2368 .It \-p1003.1-88  
 2369 .St -p1003.1-88

```

2370 .It \-p1003.1-90
2371 .St -p1003.1-90
2372 .It \-p1003.1-96
2373 .St -p1003.1-96
2374 .It \-p1003.1-2001
2375 .St -p1003.1-2001
2376 .It \-p1003.1-2004
2377 .St -p1003.1-2004
2378 .It \-p1003.1-2008
2379 .St -p1003.1-2008
2380 .It \-p1003.1
2381 .St -p1003.1
2382 .It \-p1003.1b
2383 .St -p1003.1b
2384 .It \-p1003.1b-93
2385 .St -p1003.1b-93
2386 .It \-p1003.1c-95
2387 .St -p1003.1c-95
2388 .It \-p1003.1g-2000
2389 .St -p1003.1g-2000
2390 .It \-p1003.1i-95
2391 .St -p1003.1i-95
2392 .It \-p1003.2-92
2393 .St -p1003.2-92
2394 .It \-p1003.2a-92
2395 .St -p1003.2a-92
2396 .It \-p1387.2-95
2397 .St -p1387.2-95
2398 .It \-p1003.2
2399 .St -p1003.2
2400 .It \-p1387.2
2401 .St -p1387.2
2402 .It \-isoC
2403 .St -isoC
2404 .It \-isoC-90
2405 .St -isoC-90
2406 .It \-isoC-amd1
2407 .St -isoC-amd1
2408 .It \-isoC-tcor1
2409 .St -isoC-tcor1
2410 .It \-isoC-tcor2
2411 .St -isoC-tcor2
2412 .It \-isoC-99
2413 .St -isoC-99
2414 .It \-isoC-2011
2415 .St -isoC-2011
2416 .It \-iso9945-1-90
2417 .St -iso9945-1-90
2418 .It \-iso9945-1-96
2419 .St -iso9945-1-96
2420 .It \-iso9945-2-93
2421 .St -iso9945-2-93
2422 .It \-ansiC
2423 .St -ansiC
2424 .It \-ansiC-89
2425 .St -ansiC-89
2426 .It \-ansiC-99
2427 .St -ansiC-99
2428 .It \-ieee754
2429 .St -ieee754
2430 .It \-iso8802-3
2431 .St -iso8802-3
2432 .It \-iso8601
2433 .St -iso8601
2434 .It \-ieee1275-94
2435 .St -ieee1275-94

```

```

2436 .It \-xpg3
2437 .St -xpg3
2438 .It \-xpg4
2439 .St -xpg4
2440 .It \-xpg4.2
2441 .St -xpg4.2
2442 .It \-xpg4.3
2443 .St -xpg4.3
2444 .It \-xbd5
2445 .St -xbd5
2446 .It \-xcu5
2447 .St -xcu5
2448 .It \-xsh5
2449 .St -xsh5
2450 .It \-xns5
2451 .St -xns5
2452 .It \-xns5.2
2453 .St -xns5.2
2454 .It \-xns5.2d2.0
2455 .St -xns5.2d2.0
2456 .It \-xcurses4.2
2457 .St -xcurses4.2
2458 .It \-susv2
2459 .St -susv2
2460 .It \-susv3
2461 .St -susv3
2462 .It \-svid4
2463 .St -svid4
2464 .El
2465 .Ss \&Sx
2466 Reference a section or subsection in the same manual page.
2467 The referenced section or subsection name must be identical to the
2468 enclosed argument, including whitespace.
2469 .Pp
2470 Examples:
2471 .Dl \&Sx MANUAL STRUCTURE
2472 .Pp
2473 See also
2474 .Sx \&Sh
2475 and
2476 .Sx \&Ss .
2477 .Ss \&Sy
2478 Format enclosed arguments in symbolic
2479 .Pq Dq boldface .
2480 Note that this is a presentation term and should not be used for
2481 stylistically decorating technical terms.
2482 .Pp
2483 See also
2484 .Sx \&Bf ,
2485 .Sx \&Em ,
2486 .Sx \&Li ,
2487 and
2488 .Sx \&No .
2489 .Ss \&Ta
2490 Table cell separator in
2491 .Sx \&Bl Fl column
2492 lists; can only be used below
2493 .Sx \&It .
2494 .Ss \&Tn
2495 Format a tradename.
2496 .Pp
2497 Since this macro is often implemented to use a small caps font,
2498 it has historically been used for acronyms (like ASCII) as well.
2499 Such usage is not recommended because it would use the same macro
2500 sometimes for semantical annotation, sometimes for physical formatting.
2501 .Pp

```

2502 Examples:  
 2503 .Dl \&.Tn IBM  
 2504 .Ss \&Ud  
 2505 Prints out  
 2506 .Dq currently under development.  
 2507 .Ss \&Ux  
 2508 Format the UNIX name.  
 2509 Accepts no argument.  
 2510 .Pp  
 2511 Examples:  
 2512 .Dl \&.Ux  
 2513 .Pp  
 2514 See also  
 2515 .Sx \&At ,  
 2516 .Sx \&Bsx ,  
 2517 .Sx \&Bx ,  
 2518 .Sx \&Dx ,  
 2519 .Sx \&Fx ,  
 2520 .Sx \&Nx ,  
 2521 and  
 2522 .Sx \&Ox .  
 2523 .Ss \&Va  
 2524 A variable name.  
 2525 .Pp  
 2526 Examples:  
 2527 .Dl \&.Va foo  
 2528 .Dl \&.Va const char \*bar ;  
 2529 .Ss \&Vt  
 2530 A variable type.  
 2531 This is also used for indicating global variables in the  
 2532 .Em SYNOPSIS  
 2533 section, in which case a variable name is also specified.  
 2534 Note that it accepts  
 2535 .Sx Block partial-implicit  
 2536 syntax when invoked as the first macro on an input line in the  
 2537 .Em SYNOPSIS  
 2538 section, else it accepts ordinary  
 2539 .Sx In-line  
 2540 syntax.  
 2541 In the former case, this macro starts a new output line,  
 2542 and a blank line is inserted in front if there is a preceding  
 2543 function definition or include directive.  
 2544 .Pp  
 2545 Note that this should not be confused with  
 2546 .Sx \&Ft ,  
 2547 which is used for function return types.  
 2548 .Pp  
 2549 Examples:  
 2550 .Dl \&.Vt unsigned char  
 2551 .Dl \&.Vt extern const char \* const sys\_signame[] \&;  
 2552 .Pp  
 2553 See also  
 2554 .Sx MANUAL STRUCTURE  
 2555 and  
 2556 .Sx \&Va .  
 2557 .Ss \&Xc  
 2558 Close a scope opened by  
 2559 .Sx \&Xo .  
 2560 .Ss \&Xo .  
 2561 Extend the header of an  
 2562 .Sx \&It  
 2563 macro or the body of a partial-implicit block macro  
 2564 beyond the end of the input line.  
 2565 This macro originally existed to work around the 9-argument limit  
 2566 of historic  
 2567 .Xr roff 5 .

2568 .Ss \&Xr  
 2569 Link to another manual  
 2570 .Pg Qq cross-reference .  
 2571 Its syntax is as follows:  
 2572 .Pp  
 2573 .Dl Pf \. Sx \&Xr Ar name section  
 2574 .Pp  
 2575 The  
 2576 .Ar name  
 2577 and  
 2578 .Ar section  
 2579 are the name and section of the linked manual.  
 2580 If  
 2581 .Ar section  
 2582 is followed by non-punctuation, an  
 2583 .Sx \&Ns  
 2584 is inserted into the token stream.  
 2585 This behaviour is for compatibility with  
 2586 GNU troff.  
 2587 .Pp  
 2588 Examples:  
 2589 .Dl \&.Xr mandoc 1  
 2590 .Dl \&.Xr mandoc 1 \&;  
 2591 .Dl \&.Xr mandoc 1 \&Ns s behaviour  
 2592 .Ss \&br  
 2593 Emits a line-break.  
 2594 This macro should not be used; it is implemented for compatibility with  
 2595 historical manuals.  
 2596 .Pp  
 2597 Consider using  
 2598 .Sx \&Pp  
 2599 in the event of natural paragraph breaks.  
 2600 .Ss \&sp  
 2601 Emits vertical space.  
 2602 This macro should not be used; it is implemented for compatibility with  
 2603 historical manuals.  
 2604 Its syntax is as follows:  
 2605 .Pp  
 2606 .Dl Pf \. Sx \&sp Op Ar height  
 2607 .Pp  
 2608 The  
 2609 .Ar height  
 2610 argument must be formatted as described in  
 2611 .Sx Scaling Widths .  
 2612 If unspecified,  
 2613 .Sx \&sp  
 2614 asserts a single vertical space.  
 2615 .Sh MACRO SYNTAX  
 2616 The syntax of a macro depends on its classification.  
 2617 In this section,  
 2618 .Sq \-arg  
 2619 refers to macro arguments, which may be followed by zero or more  
 2620 .Sq parm  
 2621 parameters;  
 2622 .Sq \&Yo  
 2623 opens the scope of a macro; and if specified,  
 2624 .Sq \&Yc  
 2625 closes it out.  
 2626 .Pp  
 2627 The  
 2628 .Em Callable  
 2629 column indicates that the macro may also be called by passing its name  
 2630 as an argument to another macro.  
 2631 For example,  
 2632 .Sq \&.Op \&F1 O \&Ar file  
 2633 produces

```

2634 .Sq Op Fl O Ar file .
2635 To prevent a macro call and render the macro name literally,
2636 escape it by prepending a zero-width space,
2637 .Sq \e& .
2638 For example,
2639 .Sq \&Op \e&Fl O
2640 produces
2641 .Sq Op \&Fl O .
2642 If a macro is not callable but its name appears as an argument
2643 to another macro, it is interpreted as opaque text.
2644 For example,
2645 .Sq \&.Fl \&Sh
2646 produces
2647 .Sq Fl \&Sh .
2648 .Pp
2649 The
2650 .Em Parsed
2651 column indicates whether the macro may call other macros by receiving
2652 their names as arguments.
2653 If a macro is not parsed but the name of another macro appears
2654 as an argument, it is interpreted as opaque text.
2655 .Pp
2656 The
2657 .Em Scope
2658 column, if applicable, describes closure rules.
2659 .Ss Block full-explicit
2660 Multi-line scope closed by an explicit closing macro.
2661 All macros contains bodies; only
2662 .Sx \&Bf
2663 and
2664 .Pq optionally
2665 .Sx \&Bl
2666 contain a head.
2667 .Bd -literal -offset indent
2668 \&.Yo \(\LB\arg \(\LBparm...\(rB\(\rB \(\LBhead...\(rB
2669 \(\LBbody...\(rB
2670 \&.Yc
2671 .Ed
2672 .Bl -column "MacroX" "CallableX" "ParsedX" "closed by XXX" -offset indent
2673 .It Em Macro Ta Em Callable Ta Em Parsed Ta Em Scope
2674 .It Sx \&Bd Ta \&No Ta \&No Ta closed by Sx \&Ed
2675 .It Sx \&Bf Ta \&No Ta \&No Ta closed by Sx \&Ef
2676 .It Sx \&Bk Ta \&No Ta \&No Ta closed by Sx \&Ek
2677 .It Sx \&Bl Ta \&No Ta \&No Ta closed by Sx \&El
2678 .It Sx \&Ed Ta \&No Ta \&No Ta opened by Sx \&Bd
2679 .It Sx \&Ef Ta \&No Ta \&No Ta opened by Sx \&Bf
2680 .It Sx \&Ek Ta \&No Ta \&No Ta opened by Sx \&Bk
2681 .It Sx \&El Ta \&No Ta \&No Ta opened by Sx \&Bl
2682 .El
2683 .Ss Block full-implicit
2684 Multi-line scope closed by end-of-file or implicitly by another macro.
2685 All macros have bodies; some
2686 .Po
2687 .Sx \&It Fl bullet ,
2688 .Fl hyphen ,
2689 .Fl dash ,
2690 .Fl enum ,
2691 .Fl item
2692 .Pc
2693 don't have heads; only one
2694 .Po
2695 .Sx \&It
2696 in
2697 .Sx \&Bl Fl column
2698 .Pc
2699 has multiple heads.

```

```

2700 .Bd -literal -offset indent
2701 \&.Yo \(\LB\arg \(\LBparm...\(rB\(\rB \(\LBhead...\(rB\(\rB
2702 \(\LBbody...\(rB
2703 .Ed
2704 .Bl -column "MacroX" "CallableX" "ParsedX" "closed by XXXXXXXXXXXX" -offset inden
2705 .It Em Macro Ta Em Callable Ta Em Parsed Ta Em Scope
2706 .It Sx \&It Ta \&No Ta Yes Ta closed by Sx \&It , Sx \&El
2707 .It Sx \&Nd Ta \&No Ta \&No Ta closed by Sx \&Sh
2708 .It Sx \&Nm Ta \&No Ta Yes Ta closed by Sx \&Nm , Sx \&Sh , Sx \&Ss
2709 .It Sx \&Sh Ta \&No Ta Yes Ta closed by Sx \&Sh
2710 .It Sx \&Ss Ta \&No Ta Yes Ta closed by Sx \&Sh , Sx \&Ss
2711 .El
2712 .Pp
2713 Note that the
2714 .Sx \&Nm
2715 macro is a
2716 .Sx Block full-implicit
2717 macro only when invoked as the first macro
2718 in a
2719 .Em SYNOPSIS
2720 section line, else it is
2721 .Sx In-line .
2722 .Ss Block partial-explicit
2723 Like block full-explicit, but also with single-line scope.
2724 Each has at least a body and, in limited circumstances, a head
2725 .Po
2726 .Sx \&Fo ,
2727 .Sx \&Eo
2728 .Pc
2729 and/or tail
2730 .Pq Sx \&Ec .
2731 .Bd -literal -offset indent
2732 \&.Yo \(\LB\arg \(\LBparm...\(rB\(\rB \(\LBhead...\(rB
2733 \(\LBbody...\(rB
2734 \&.Yc \(\LBtail...\(rB
2736 \&.Yo \(\LB\arg \(\LBparm...\(rB\(\rB \(\LBhead...\(rB
2737 \(\LBbody...\(rB \&Yc \(\LBtail...\(rB
2738 .Ed
2739 .Bl -column "MacroX" "CallableX" "ParsedX" "closed by XXXX" -offset indent
2740 .It Em Macro Ta Em Callable Ta Em Parsed Ta Em Scope
2741 .It Sx \&Ac Ta Yes Ta Yes Ta opened by Sx \&Ao
2742 .It Sx \&Ao Ta Yes Ta Yes Ta closed by Sx \&Ac
2743 .It Sx \&Bc Ta Yes Ta Yes Ta closed by Sx \&Bo
2744 .It Sx \&Bo Ta Yes Ta Yes Ta opened by Sx \&Bc
2745 .It Sx \&Brc Ta Yes Ta Yes Ta opened by Sx \&Bro
2746 .It Sx \&Bro Ta Yes Ta Yes Ta closed by Sx \&Brc
2747 .It Sx \&Dc Ta Yes Ta Yes Ta opened by Sx \&Do
2748 .It Sx \&Do Ta Yes Ta Yes Ta closed by Sx \&Dc
2749 .It Sx \&Ec Ta Yes Ta Yes Ta opened by Sx \&Eo
2750 .It Sx \&Eo Ta Yes Ta Yes Ta closed by Sx \&Ec
2751 .It Sx \&Fc Ta Yes Ta Yes Ta opened by Sx \&Fo
2752 .It Sx \&Fo Ta \&No Ta \&No Ta closed by Sx \&Fc
2753 .It Sx \&Oc Ta Yes Ta Yes Ta closed by Sx \&Oo
2754 .It Sx \&Oo Ta Yes Ta Yes Ta opened by Sx \&Oc
2755 .It Sx \&Pc Ta Yes Ta Yes Ta closed by Sx \&Po
2756 .It Sx \&Po Ta Yes Ta Yes Ta opened by Sx \&Pc
2757 .It Sx \&Qc Ta Yes Ta Yes Ta opened by Sx \&Qo
2758 .It Sx \&Qo Ta Yes Ta Yes Ta closed by Sx \&Qc
2759 .It Sx \&Rc Ta \&No Ta \&No Ta opened by Sx \&Rs
2760 .It Sx \&Rc Ta \&No Ta \&No Ta closed by Sx \&Re
2761 .It Sx \&Sc Ta Yes Ta Yes Ta opened by Sx \&So
2762 .It Sx \&So Ta Yes Ta Yes Ta closed by Sx \&Sc
2763 .It Sx \&Xc Ta Yes Ta Yes Ta opened by Sx \&Xo
2764 .It Sx \&Xo Ta Yes Ta Yes Ta closed by Sx \&Xc
2765 .El

```

```

2766 .Ss Block partial-implicit
2767 Like block full-implicit, but with single-line scope closed by the
2768 end of the line.
2769 .Bd -literal -offset indent
2770 &\.Yo \(\B-arg \(\Bval...\(rB\(\Bbody...\(rB \(\Bres...\(rB
2771 .Ed
2772 .Bl -column "MacroX" "CallableX" "ParsedX" -offset indent
2773 .It Em Macro Ta Em Callable Ta Em Parsed
2774 .It Sx \&Aq Ta Yes Ta Yes
2775 .It Sx \&Bq Ta Yes Ta Yes
2776 .It Sx \&Brq Ta Yes Ta Yes
2777 .It Sx \&Dl Ta \&No Ta \&Yes
2778 .It Sx \&Dl Ta \&No Ta Yes
2779 .It Sx \&Dq Ta Yes Ta Yes
2780 .It Sx \&Op Ta Yes Ta Yes
2781 .It Sx \&Pq Ta Yes Ta Yes
2782 .It Sx \&Ql Ta Yes Ta Yes
2783 .It Sx \&Qq Ta Yes Ta Yes
2784 .It Sx \&Sq Ta Yes Ta Yes
2785 .It Sx \&Vt Ta Yes Ta Yes
2786 .El
2787 .Pp
2788 Note that the
2789 .Sx \&Vt
2790 macro is a
2791 .Sx Block partial-implicit
2792 only when invoked as the first macro
2793 in a
2794 .Em SYNOPSIS
2795 section line, else it is
2796 .Sx In-line .
2797 .Ss Special block macro
2798 The
2799 .Sx \&Ta
2800 macro can only be used below
2801 .Sx \&It
2802 in
2803 .Sx \&Bl Fl column
2804 lists.
2805 It delimits blocks representing table cells;
2806 these blocks have bodies, but no heads.
2807 .Bl -column "MacroX" "CallableX" "ParsedX" "closed by XXXX" -offset indent
2808 .It Em Macro Ta Em Callable Ta Em Parsed Ta Em Scope
2809 .It Sx \&Ta Ta Yes Ta Yes Ta closed by Sx \&Ta , Sx \&It
2810 .El
2811 .Ss In-line
2812 Closed by the end of the line, fixed argument lengths,
2813 and/or subsequent macros.
2814 In-line macros have only text children.
2815 If a number (or inequality) of arguments is
2816 .Pq n ,
2817 then the macro accepts an arbitrary number of arguments.
2818 .Bd -literal -offset indent
2819 &\.Yo \(\B-arg \(\Bval...\(rB\(\Bargs...\(rB \(\Bres...\(rB
2821 &\.Yo \(\B-arg \(\Bval...\(rB\(\Bargs...\(rB Yc...
2823 &\.Yo \(\B-arg \(\Bval...\(rB\(\B arg0 arg1 argN
2824 .Ed
2825 .Bl -column "MacroX" "CallableX" "ParsedX" "Arguments" -offset indent
2826 .It Em Macro Ta Em Callable Ta Em Parsed Ta Em Arguments
2827 .It Sx \&A Ta \&No Ta \&No Ta >0
2828 .It Sx \&B Ta \&No Ta \&No Ta >0
2829 .It Sx \&C Ta \&No Ta \&No Ta >0
2830 .It Sx \&D Ta \&No Ta \&No Ta >0
2831 .It Sx \&I Ta \&No Ta \&No Ta >0

```

```

2832 .It Sx \&J Ta \&No Ta \&No Ta >0
2833 .It Sx \&N Ta \&No Ta \&No Ta >0
2834 .It Sx \&O Ta \&No Ta \&No Ta >0
2835 .It Sx \&P Ta \&No Ta \&No Ta >0
2836 .It Sx \&Q Ta \&No Ta \&No Ta >0
2837 .It Sx \&R Ta \&No Ta \&No Ta >0
2838 .It Sx \&T Ta \&No Ta \&No Ta >0
2839 .It Sx \&U Ta \&No Ta \&No Ta >0
2840 .It Sx \&V Ta \&No Ta \&No Ta >0
2841 .It Sx \&Ad Ta Yes Ta Yes Ta >0
2842 .It Sx \&An Ta Yes Ta Yes Ta >0
2843 .It Sx \&Ap Ta Yes Ta Yes Ta 0
2844 .It Sx \&Ar Ta Yes Ta Yes Ta n
2845 .It Sx \&At Ta Yes Ta Yes Ta 1
2846 .It Sx \&Bsx Ta Yes Ta Yes Ta n
2847 .It Sx \&Bt Ta \&No Ta \&No Ta 0
2848 .It Sx \&Bx Ta Yes Ta Yes Ta n
2849 .It Sx \&Cd Ta Yes Ta Yes Ta >0
2850 .It Sx \&Cm Ta Yes Ta Yes Ta >0
2851 .It Sx \&Db Ta \&No Ta \&No Ta 1
2852 .It Sx \&Dd Ta \&No Ta \&No Ta n
2853 .It Sx \&Dt Ta \&No Ta \&No Ta n
2854 .It Sx \&Dv Ta Yes Ta Yes Ta >0
2855 .It Sx \&Dx Ta Yes Ta Yes Ta n
2856 .It Sx \&Em Ta Yes Ta Yes Ta >0
2857 .It Sx \&En Ta \&No Ta \&No Ta 0
2858 .It Sx \&Er Ta Yes Ta Yes Ta >0
2859 .It Sx \&Es Ta \&No Ta \&No Ta 0
2860 .It Sx \&Ev Ta Yes Ta Yes Ta >0
2861 .It Sx \&Ex Ta \&No Ta \&No Ta n
2862 .It Sx \&Fa Ta Yes Ta Yes Ta >0
2863 .It Sx \&Fd Ta \&No Ta \&No Ta >0
2864 .It Sx \&Fl Ta Yes Ta Yes Ta n
2865 .It Sx \&Fn Ta Yes Ta Yes Ta >0
2866 .It Sx \&Fr Ta \&No Ta \&No Ta n
2867 .It Sx \&Ft Ta Yes Ta Yes Ta >0
2868 .It Sx \&Fx Ta Yes Ta Yes Ta n
2869 .It Sx \&Hf Ta \&No Ta \&No Ta n
2870 .It Sx \&Ic Ta Yes Ta Yes Ta >0
2871 .It Sx \&In Ta \&No Ta \&No Ta 1
2872 .It Sx \&Lb Ta \&No Ta \&No Ta 1
2873 .It Sx \&Li Ta Yes Ta Yes Ta >0
2874 .It Sx \&Lk Ta Yes Ta Yes Ta >0
2875 .It Sx \&Lp Ta \&No Ta \&No Ta 0
2876 .It Sx \&Ms Ta Yes Ta Yes Ta >0
2877 .It Sx \&Mt Ta Yes Ta Yes Ta >0
2878 .It Sx \&Nm Ta Yes Ta Yes Ta n
2879 .It Sx \&No Ta Yes Ta Yes Ta 0
2880 .It Sx \&Ns Ta Yes Ta Yes Ta 0
2881 .It Sx \&Nx Ta Yes Ta Yes Ta n
2882 .It Sx \&Os Ta \&No Ta \&No Ta n
2883 .It Sx \&Ot Ta \&No Ta \&No Ta n
2884 .It Sx \&Ox Ta Yes Ta Yes Ta n
2885 .It Sx \&Pa Ta Yes Ta Yes Ta n
2886 .It Sx \&Pf Ta Yes Ta Yes Ta 1
2887 .It Sx \&Pp Ta \&No Ta \&No Ta 0
2888 .It Sx \&Rv Ta \&No Ta \&No Ta n
2889 .It Sx \&Sm Ta \&No Ta \&No Ta 1
2890 .It Sx \&St Ta \&No Ta Yes Ta 1
2891 .It Sx \&Sx Ta Yes Ta Yes Ta >0
2892 .It Sx \&Sy Ta Yes Ta Yes Ta >0
2893 .It Sx \&Tn Ta Yes Ta Yes Ta >0
2894 .It Sx \&Ud Ta \&No Ta \&No Ta 0
2895 .It Sx \&Ux Ta Yes Ta Yes Ta n
2896 .It Sx \&Va Ta Yes Ta Yes Ta n
2897 .It Sx \&Vt Ta Yes Ta Yes Ta >0

```

2898 .It Sx \&Xr Ta Yes Ta Yes Ta >0  
 2899 .It Sx \&br Ta \&No Ta \&No Ta 0  
 2900 .It Sx \&sp Ta \&No Ta \&No Ta 1  
 2901 .El  
 2902 .Ss Delimiters  
 2903 When a macro argument consists of one single input character  
 2904 considered as a delimiter, the argument gets special handling.  
 2905 This does not apply when delimiters appear in arguments containing  
 2906 more than one character.  
 2907 Consequently, to prevent special handling and just handle it  
 2908 like any other argument, a delimiter can be escaped by prepending  
 2909 a zero-width space  
 2910 .Pq Sq \e& .  
 2911 In text lines, delimiters never need escaping, but may be used  
 2912 as normal punctuation.  
 2913 .Pp  
 2914 For many macros, when the leading arguments are opening delimiters,  
 2915 these delimiters are put before the macro scope,  
 2916 and when the trailing arguments are closing delimiters,  
 2917 these delimiters are put after the macro scope.  
 2918 For example,  
 2919 .Pp  
 2920 .Dl Pf \. \&Aq "( [ word ] ) ."  
 2921 .Pp  
 2922 renders as:  
 2923 .Pp  
 2924 .Dl Aq ( [ word ] ) .  
 2925 .Pp  
 2926 Opening delimiters are:  
 2927 .Pp  
 2928 .Bl -tag -width Ds -offset indent -compact  
 2929 .It \&(  
 2930 left parenthesis  
 2931 .It \&[  
 2932 left bracket  
 2933 .El  
 2934 .Pp  
 2935 Closing delimiters are:  
 2936 .Pp  
 2937 .Bl -tag -width Ds -offset indent -compact  
 2938 .It \&.  
 2939 period  
 2940 .It \&,  
 2941 comma  
 2942 .It \&:  
 2943 colon  
 2944 .It \& ;  
 2945 semicolon  
 2946 .It \& )  
 2947 right parenthesis  
 2948 .It \& ]  
 2949 right bracket  
 2950 .It \& ?  
 2951 question mark  
 2952 .It \& !  
 2953 exclamation mark  
 2954 .El  
 2955 .Pp  
 2956 Note that even a period preceded by a backslash  
 2957 .Pq Sq \e.\&  
 2958 gets this special handling; use  
 2959 .Sq \e& .  
 2960 to prevent that.  
 2961 .Pp  
 2962 Many in-line macros interrupt their scope when they encounter  
 2963 delimiters, and resume their scope when more arguments follow that

2964 are not delimiters.  
 2965 For example,  
 2966 .Pp  
 2967 .Dl Pf \. \&Fl "a ( b | c \e\*(Ba d ) e"  
 2968 .Pp  
 2969 renders as:  
 2970 .Pp  
 2971 .Dl Fl a ( b | c \\*(Ba d ) e  
 2972 .Pp  
 2973 This applies to both opening and closing delimiters,  
 2974 and also to the middle delimiter:  
 2975 .Pp  
 2976 .Bl -tag -width Ds -offset indent -compact  
 2977 .It \&|  
 2978 vertical bar  
 2979 .El  
 2980 .Pp  
 2981 As a special case, the predefined string \e\*(Ba is handled and rendered  
 2982 in the same way as a plain  
 2983 .Sq \&|  
 2984 character.  
 2985 Using this predefined string is not recommended in new manuals.  
 2986 .Ss Font handling  
 2987 .In  
 2988 .Nm  
 2989 documents, usage of semantic markup is recommended in order to have  
 2990 proper fonts automatically selected; only when no fitting semantic markup  
 2991 is available, consider falling back to  
 2992 .Sx Physical markup  
 2993 macros.  
 2994 Whenever any  
 2995 .Nm  
 2996 macro switches the  
 2997 .Xr roff 5  
 2998 font mode, it will automatically restore the previous font when exiting  
 2999 its scope.  
 3000 Manually switching the font using the  
 3001 .Xr roff 5  
 3002 .Ql \ef  
 3003 font escape sequences is never required.  
 3004 .Sh COMPATIBILITY  
 3005 This section documents compatibility between mandoc and other other  
 3006 troff implementations, at this time limited to GNU troff  
 3007 .Pq Qq groff .  
 3008 The term  
 3009 .Qq historic groff  
 3010 refers to groff versions before 1.17,  
 3011 which featured a significant update of the  
 3012 .Pa doc.tmac  
 3013 file.  
 3014 .Pp  
 3015 Heirloom troff, the other significant troff implementation accepting  
 3016 \-mdoc, is similar to historic groff.  
 3017 .Pp  
 3018 The following problematic behaviour is found in groff:  
 3019 .ds hist (Historic groff only.)  
 3020 .Pp  
 3021 .Bl -dash -compact  
 3022 .It  
 3023 Display macros  
 3024 .Po  
 3025 .Sx \&Bd ,  
 3026 .Sx \&Dl ,  
 3027 and  
 3028 .Sx \&Dl  
 3029 .Pc



3030 may not be nested.  
 3031 `\*[hist]`  
 3032 `.It`  
 3033 `.Sx \&At`  
 3034 with unknown arguments produces no output at all.  
 3035 `\*[hist]`  
 3036 Newer groff and mandoc print  
 3037 `.Qq AT&T UNIX`  
 3038 and the arguments.  
 3039 `.It`  
 3040 `.Sx \&Bl Fl column`  
 3041 does not recognise trailing punctuation characters when they immediately  
 3042 precede tabulator characters, but treats them as normal text and  
 3043 outputs a space before them.  
 3044 `.It`  
 3045 `.Sx \&Bd Fl ragged compact`  
 3046 does not start a new line.  
 3047 `\*[hist]`  
 3048 `.It`  
 3049 `.Sx \&Dd`  
 3050 with non-standard arguments behaves very strangely.  
 3051 When there are three arguments, they are printed verbatim.  
 3052 Any other number of arguments is replaced by the current date,  
 3053 but without any arguments the string  
 3054 `.Dq Epoch`  
 3055 is printed.  
 3056 `.It`  
 3057 `.Sx \&Fl`  
 3058 does not print a dash for an empty argument.  
 3059 `\*[hist]`  
 3060 `.It`  
 3061 `.Sx \&Fn`  
 3062 does not start a new line unless invoked as the line macro in the  
 3063 `.Em SYNOPSIS`  
 3064 section.  
 3065 `\*[hist]`  
 3066 `.It`  
 3067 `.Sx \&Fo`  
 3068 with  
 3069 `.Pf non- Sx \&Fa`  
 3070 children causes inconsistent spacing between arguments.  
 3071 In mandoc, a single space is always inserted between arguments.  
 3072 `.It`  
 3073 `.Sx \&Ft`  
 3074 in the  
 3075 `.Em SYNOPSIS`  
 3076 causes inconsistent vertical spacing, depending on whether a prior  
 3077 `.Sx \&Fn`  
 3078 has been invoked.  
 3079 See  
 3080 `.Sx \&Ft`  
 3081 and  
 3082 `.Sx \&Fn`  
 3083 for the normalised behaviour in mandoc.  
 3084 `.It`  
 3085 `.Sx \&In`  
 3086 ignores additional arguments and is not treated specially in the  
 3087 `.Em SYNOPSIS`  
 3088 `\*[hist]`  
 3089 `.It`  
 3090 `.Sx \&It`  
 3091 sometimes requires a  
 3092 `.Fl nested`  
 3093 flag.  
 3094 `\*[hist]`  
 3095 In new groff and mandoc, any list may be nested by default and

3096 `.Fl enum`  
 3097 lists will restart the sequence only for the sub-list.  
 3098 `.It`  
 3099 `.Sx \&Li`  
 3100 followed by a delimiter is incorrectly used in some manuals  
 3101 instead of properly quoting that character, which sometimes works with  
 3102 historic groff.  
 3103 `.It`  
 3104 `.Sx \&Lk`  
 3105 only accepts a single link-name argument; the remainder is misformatted.  
 3106 `.It`  
 3107 `.Sx \&Pa`  
 3108 does not format its arguments when used in the FILES section under  
 3109 certain list types.  
 3110 `.It`  
 3111 `.Sx \&Ta`  
 3112 can only be called by other macros, but not at the beginning of a line.  
 3113 `.It`  
 3114 `.Sx \&%C`  
 3115 is not implemented.  
 3116 `.It`  
 3117 Historic groff only allows up to eight or nine arguments per macro input  
 3118 line, depending on the exact situation.  
 3119 Providing more arguments causes garbled output.  
 3120 The number of arguments on one input line is not limited with mandoc.  
 3121 `.It`  
 3122 Historic groff has many un-callable macros.  
 3123 Most of these (excluding some block-level macros) are callable  
 3124 in new groff and mandoc.  
 3125 `.It`  
 3126 `.Sq \(\ba`  
 3127 (vertical bar) is not fully supported as a delimiter.  
 3128 `\*[hist]`  
 3129 `.It`  
 3130 `.Sq \ef`  
 3131 `.Pq font face`  
 3132 and  
 3133 `.Sq \ef`  
 3134 `.Pq font family face`  
 3135 `.Sx Text Decoration`  
 3136 escapes behave irregularly when specified within line-macro scopes.  
 3137 `.It`  
 3138 Negative scaling units return to prior lines.  
 3139 Instead, mandoc truncates them to zero.  
 3140 `.El`  
 3141 `.Pp`  
 3142 The following features are unimplemented in mandoc:  
 3143 `.Pp`  
 3144 `.Bl -dash -compact`  
 3145 `.It`  
 3146 `.Sx \&Bd`  
 3147 `.Fl file Ar file .`  
 3148 `.It`  
 3149 `.Sx \&Bd`  
 3150 `.Fl offset Ar center`  
 3151 and  
 3152 `.Fl offset Ar right .`  
 3153 Groff does not implement centred and flush-right rendering either,  
 3154 but produces large indentations.  
 3155 `.It`  
 3156 The  
 3157 `.Sq \eh`  
 3158 `.Pq horizontal position ,`  
 3159 `.Sq \ev`  
 3160 `.Pq vertical position ,`  
 3161 `.Sq \em`

3162 .Pq text colour ,  
3163 .Sq \eM  
3164 .Pq text filling colour ,  
3165 .Sq \ez  
3166 .Pq zero-length character ,  
3167 .Sq \ew  
3168 .Pq string length ,  
3169 .Sq \ek  
3170 .Pq horizontal position marker ,  
3171 .Sq \eo  
3172 .Pq text overstrike ,  
3173 and  
3174 .Sq \es  
3175 .Pq text size  
3176 escape sequences are all discarded in mandoc.  
3177 .It  
3178 The  
3179 .Sq \ef  
3180 scaling unit is accepted by mandoc, but rendered as the default unit.  
3181 .It  
3182 In quoted literals, groff allows pairwise double-quotes to produce a  
3183 standalone double-quote in formatted output.  
3184 This is not supported by mandoc.  
3185 .El  
3186 .Sh SEE ALSO  
3187 .Xr man 1 ,  
3188 .Xr mandoc 1 ,  
3189 .Xr eqn 5 ,  
3190 .Xr man 5 ,  
3191 .Xr mandoc\_char 5 ,  
3192 .Xr roff 5 ,  
3193 .Xr tbl 5  
3194 .Sh HISTORY  
3195 The  
3196 .Nm  
3197 language first appeared as a troff macro package in  
3198 .Bx 4.4 .  
3199 It was later significantly updated by Werner Lemberg and Ruslan Ermilov  
3200 in groff-1.17.  
3201 The standalone implementation that is part of the  
3202 .Xr mandoc 1  
3203 utility written by Kristaps Dzonsons appeared in  
3204 .Ox 4.6 .  
3205 in July, 2014.  
3206 .Sh AUTHORS  
3207 The  
3208 .Nm  
3209 reference was written by  
3210 .An Kristaps Dzonsons ,  
3211 .Mt kristaps@bsd.lv .

\*\*\*\*\*

4186 Thu Jul 17 00:50:42 2014

new/usr/src/man/man7d/cpqary3.7d

manpage lint.

\*\*\*\*\*

```

1  \."
2  \." This file and its contents are supplied under the terms of the
3  \." Common Development and Distribution License ("CDDL"), version 1.0.
4  \." You may only use this file in accordance with the terms of version
5  \." 1.0 of the CDDL.
6  \."
7  \." A full copy of the text of the CDDL should have accompanied this
8  \." source. A copy of the CDDL is also available via the Internet at
9  \." http://www.illumos.org/license/CDDL.
10 \."
11 \."
12 \." Copyright (C) 2013 Hewlett-Packard Development Company, L.P.
13 \."
14 .TH CPQARY3 7D Aug 26, 2013"
15 .SH NAME
16 .LP
16 cpqary3 - provides disk and SCSI tape support for HP Smart Array controllers
18 .LP
17 .SH DESCRIPTION
18 .LP
19 The cpqary3 module provides low-level interface routines between the common
20 disk I/O subsystem and the HP SMART Array controllers. The cpqary3 driver
21 provides disk and SCSI tape support for the HP Smart Array controllers.
22 .LP
23 Please refer to the cpqary3 release notes, for the supported HP Smart Array
24 Controllers and Storage boxes.
25 .LP
26 Each controller should be the sole initiator on a SCSI bus. Auto
27 configuration code determines if the adapter is present at the Configured
28 address and what types of devices are attached to it.
29 .SH CONFIGURATION
30 Use the Array Configuration Utility to configure the controllers. Each
31 controller can support up to 32 logical volumes. In addition, each controller
32 supports up to a maximum of 28 connected SCSI tape drives.
33 With 1.90 and later versions of cpqary3 driver, HP Smart Array SAS controllers,
34 having Firmware Revision 5.10 or later, will support 64 logical drives. This
35 firmware also supports Dual Domain Multipath configurations.
36 .LP
37 The driver attempts to initialize itself in accordance with the information
38 found in the configuration file, /kernel/drv/cpqary3.conf.
39 .LP
40 New component - hmaeventd which logs the storage events onto console and to the
41 Integrated Management Log is made a part of HPQhma 5.0.0 package, which is not
42 part of the operating system. Therefore, by default, notify on event
43 functionality is disabled in the driver from 2.1.0 onwards. Storage event
44 logging may be enabled in the driver by modifying cpqary3.conf to set the
45 cpqary3.noevent property to "on". Modification of driver properties requires
46 reloading the driver, which in most cases will occur only following a reboot of
47 the system.
48 .LP
49 The target driver's configuration file shall need entries if support is needed
50 for targets numbering greater than the default number of targets supported by
51 the corresponding target driver.
52 .LP
53 By default, entries for SCSI target numbers 0 to 15 are present in sd.conf.
54 Entries for target numbers 16 and above must be added to the '&'scsi' class in
55 sd.conf to support additional corresponding logical volumes.
56 .LP
57 If SCSI tape drives are connected to the supported controllers, entries for
58 target IDs from 33 to 33+N must be added in the /kernel/drv/st.conf file under
59 '&'scsi' class, where N is the total number of SCSI tape drives connected to the

```

```

60 controller with largest number of tape drives connected to it, in the existing
61 configuration. For example, two supported controller, c1 and c2 are present in
62 the system. If controller c1 has two (2) tape drives and controller c2 has five
63 (5) tape drives connected, then entries for target IDs 33 thru 38 are required
64 under 'scsi' class in /kernel/drv/st.conf file. The maximum number of tape
65 drives that can be connected to a controller is 28. With 1.90 and later versions
66 of cpqary3 driver, if tape drives are connected to Smart Array SAS controllers,
67 then target ID entries for tape drives from 65 to 65+N must be added in
68 /kernel/drv/st.conf file under the '&'scsi' class.
69 .SH FILES
70 .PD 0
71 .TP 25
72 .B /kernel/drv/cpqary3.conf
73 - configuration file for CPQary3
74 .PD
75 .SH "SEE ALSO"
76 .BR driver.conf (4),
77 .BR sd (7D),
78 .BR st (7D)
81 .LP
79 .SH NOTES
80 .LP
81 The Smart Array controllers supported by the current version of the
82 cpqary3 driver do not support 'format unit' SCSI command. Hence, selecting
83 '&'format' option under 'format' utility main menu is not supported. In addition
84 the 'repair' option under 'format' utility main menu is not supported as this
85 operation is not applicable to Logical volumes connected to the supported Smart
86 Array controllers.

```

new/usr/src/man/man7d/nv\_sata.7d

1

\*\*\*\*\*

2008 Thu Jul 17 00:50:42 2014

new/usr/src/man/man7d/nv\_sata.7d

manpage lint.

\*\*\*\*\*

```
1 \" te
2.\" Copyright (c) 2008, Sun Microsystems, Inc. All Rights Reserved
3.\" Copyright 2011 Nexenta Systems, Inc. All rights reserved.
4.\" The contents of this file are subject to the terms of the Common Development
5.\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
6.\" When distributing Covered Code, include this CDDL HEADER in each file and in
7 .TH NV_SATA 7D "Sep 25, 2011"
8 .TH nv_sata 7D "25 Sep 2011"
9 .SH NAME
10 nv_sata \- NVIDIA CK804/MCP04/MCP51/MCP55/MCP61 SATA controller driver
11 .SH SYNOPSIS
12 .LP
13 \fBsata@unit-address\fR
14 .fi
16 .SH DESCRIPTION
17 .sp
18 .LP
19 The \fBnv_sata\fR driver is a SATA HBA driver that supports NVIDIA CK804/MCP04
20 and MCP51/MCP55/MCP61 SATA HBA controllers. Note that while these controllers
21 support standard SATA features including SATA-II drives, NCQ, hotplug and ATAPI
22 drives, the driver currently does not support NCQ features.
23 .SH CONFIGURATION
24 .sp
25 .LP
26 The \fBnv_sata\fR module contains no user configurable parameters.
27 .SH FILES
28 .sp
29 .ne 2
30 .na
31 \fB\fB/kernel/drv/nv_sata\fR\fR
32 .ad
33 .sp .6
34 .RS 4n
35 32-bit \fBEBELF\fR kernel module (x86).
36 .RE
38 .sp
39 .ne 2
40 .na
41 \fB\fB/kernel/drv/amd64/nv_sata\fR\fR
42 .ad
43 .sp .6
44 .RS 4n
45 64-bit \fBEBELF\fR kernel module (x86).
46 .RE
48 .SH ATTRIBUTES
49 .sp
50 .LP
51 See \fBattributes\fR(5) for descriptions of the following attribute:
52 .sp
54 .sp
55 .TS
56 box;
57 c | c
58 l | l .
59 ATTRIBUTE TYPE ATTRIBUTE VALUE
60 _
```

new/usr/src/man/man7d/nv\_sata.7d

2

```
61 Architecture x86
62 .TE
64 .SH SEE ALSO
65 .sp
66 .LP
67 \fBbcfgadm\fR(1M), \fBbcfgadm_sata\fR(1M), \fBprtconf\fR(1M), \fBsata\fR(7D),
68 \fBsd\fR(7D)
69 .sp
70 .LP
71 \fBIWriting Device Drivers\fR
```

\*\*\*\*\*

4981 Thu Jul 17 00:50:42 2014

new/usr/src/man/man7d/pcn.7d

manpage lint.

\*\*\*\*\*

```

1 \" te
2.\" Copyright 2011 Jason King. All rights reserved.
3.\" Copyright (c) 2001-2007 by Garrett D'Amore.
4.\" Redistribution and use in source and binary forms, with or without
5.\" modification, are permitted provided that the following conditions are met:
6.\" 1. Redistributions of source code must retain the above copyright notice,
7.\" this list of conditions and the following disclaimer.
8.\" 2. Redistributions in binary form must reproduce the above copyright notice,
9.\" this list of conditions and the following disclaimer in the documentation
10.\" and/or other materials provided with the distribution.
11.\" 3. Neither the name of the author nor the names of any co-contributors may
12.\" be used to endorse or promote products derived from this software
13.\" without specific prior written permission.
14.\" THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER AND CONTRIBUTORS
15.\" 'AS IS' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED
16.\" TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
17.\" PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR
18.\" CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,
19.\" EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
20.\" PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS;
21.\" OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
22.\" WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR
23.\" OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF
24.\" ADVISED OF THE POSSIBILITY OF SUCH DAMAGE
25.\" Portions Copyright (c) 2007 by Sun Microsystems, Inc. All Rights Reserved.

```

27.TH "PCN" "7D" "Sep 16, 2011"

28.

29.SH "NAME"

30.\fBpcn\fR \- PCnet Ethernet device driver

31.SH "SYNOPSIS"

32.LP

33.nf

34.\fB/dev/pcn\fR

35.fi

37.SH "DESCRIPTION"

38.sp

39.LP

40.The \fBpcn\fR driver is a multi-threaded, loadable, clonable GLDv3-based

41.STREAMS driver supporting the Data Link Provider Interface \fBdmpi\fR(7P) for

42.the AMD PCnet family of Ethernet controllers\.

43.SH "APPLICATION PROGRAMMING INTERFACE"

44.The \fBpcn\fR driver can be used as either a "style 1" or a "style 2" Data Link

45.Service Provider\.

46. **instance number of the \fBpcn\fR controller as assigned by the**46. *instance number of the \fBpcn\fR controller as assigned by the *Illumos**

47. operating environment\.

48.sp

49.LP

50.The values returned by the driver in the \fBBDL\_INFO\_ACK\fR response are:

51.RS +4

52.TP

53.ie t \(\bu

54.el o

55.Maximum SDU is 1500\.

56.RE

57.RS +4

58.TP

59.ie t \(\bu

60.el o

61 Minimum SDU is 0\.

62.RE

63.RS +4

64.TP

65.ie t \(\bu

66.el o

67.The dlsap address length is 8\.

68.RE

69.RS +4

70.TP

71.ie t \(\bu

72.el o

73.MAC type is \fBBDL\_ETHER\fR\.

74.RE

75.RS +4

76.TP

77.ie t \(\bu

78.el o

79.SAP length is \-2\.

80. The 6-byte physical address is immediately followed by a

80. 2-byte SAP\.

81.RE

82.RS +4

83.TP

84.ie t \(\bu

85.el o

86.Service mode is \fBBDL\_CLDLS\fR\.

87.RE

88.RS +4

89.TP

90.ie t \(\bu

91.el o

92.The broadcast address is the 6-byte Ethernet broadcast address

93. (\fBff:ff:ff:ff:ff:ff\fR)\.

94.SH "CONFIGURATION"

95.sp

96.LP

97.The \fBpcn\fR driver performs auto-negotiation to select the link speed and

98.mode\.

98. Link speed may be 100Mbps full-duplex, 100Mbps half-duplex,

99. 10Mbps full-duplex, or 10Mbps half-duplex, depending on the hardware

100.adapter type\.

100. See the \fBIEEE802.3\fR standard for more information\.

101.sp

102.LP

103.The capabilities advertised by the \fBpcn\fR device can be set using

104.\fBdldm\fR(lm)\.

104. The driver supports a number of parameters whose names

105.being with \fBbn\_\fR (see below)\.

105. Each of these parameters contains a

106.boolean value that determines if the devices advertises that mode of

107.operations\.

107. The \fBadv\_autoneg\_cap\fR parameter controls whether

108.auto-negotiation is performed\.

108. If \fBadv\_autoneg\_cap\fR is set to 0, the

109.driver forces the mode of operation selected by the first non-zero

110.parameter in priority order as shown below:

111.sp

112.in +2

113.nf

114

115.en\_100fdx\_cap

116.en\_10fdx\_cap

117

118.fi

119.in -2

(highest priority/greatest throughput)

100Mbps full duplex

10Mbps full duplex

(lowest priority/least throughput)

121.sp

122.LP

123.All capabilities default to enabled\.

123. Note that changing any capability

124.parameter causes te link to go down while the link partners renegotiate

125.the link speed/duplex using the newly changed capabilities\.

126.SH "ATTRIBUTES"

```
127 .sp
128 .LP
129 See \fBattributes\fR(5) for a description of the following attributes:
130 .sp

132 .sp
133 .TS
134 box;
135 c | c
136 l | l .
137 ATTRIBUTE TYPE ATTRIBUTE VALUE
138 _
139 Architecture x86
140 _
141 Interface Stability Committed
142 .TE

144 .SH "FILES"
145 .sp
146 .ne 2
147 .na
148 \fB\dev\pcn\fR\fR
149 .ad
150 .sp .6
151 .RS 4n
152 Special character device\
153 .RE

155 .sp
156 .ne 2
157 .na
158 \fB\kernel\drv\pcn\fR\fR
159 .ad
160 .sp 6
161 .RS 4n
162 32\bit driver binary\
163 .RE

165 .sp
166 .ne 2
167 .na
168 \fB\kernel\drv\amd64\pcn\fR\fR
169 .ad
170 .sp .6
171 .RS 4n
172 64\bit driver binary (x86)\
173 .RE

175 .SH "SEE ALSO"
176 .sp
177 .LP
178 \fB\dldpi\fR(1m), \fBattributes\fR(5), \fBstreamio\fR(7I), \fBdldpi\fR(7p)
179 .sp
180 .LP
181 \fIIEEE 802.3\fR \f(em Institute of Electrical and Electronics Engineers, 2002
```

new/usr/src/man/man9f/ddi\_dmae.9f

1

```
*****  
11098 Thu Jul 17 00:50:42 2014  
new/usr/src/man/man9f/ddi_dmae.9f  
manpage lint.  
*****  
_____unchanged_portion_omitted_____
```

\*\*\*\*\*

7939 Thu Jul 17 00:50:42 2014

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

manpage lint.

\*\*\*\*\*

```

1 \" te
2.\" Copyright (c) 2006 Sun Microsystems, Inc. All Rights Reserved
3.\" Copyright (c) 2013, Joyent, Inc. All rights reserved.
4.\" The contents of this file are subject to the terms of the Common Development
5.\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
6.\" When distributing Covered Code, include this CDDL HEADER in each file and in
7.TH RWLOCK 9F "Sep 19, 2013"
8.SH NAME
9 rwlock, rw_init, rw_destroy, rw_enter, rw_exit, rw_tryenter, rw_downgrade,
10 rw_tryupgrade, rw_read_locked \- readers/writer lock functions
11.SH SYNOPSIS
12.LP
13.nf
14 #include <sys/ksynch.h>

```

```

18 \fBvoid\fR \fBBrw_init\fR(\fBkrwlock_t *\fR\fRwlp\fR, \fBchar *\fR\fRname\fR, \fR
19 .fi

```

```

21 .LP
22 .nf
23 \fBvoid\fR \fBBrw_destroy\fR(\fBkrwlock_t *\fR\fRwlp\fR);
24 .fi

```

```

26 .LP
27 .nf
28 \fBvoid\fR \fBBrw_enter\fR(\fBkrwlock_t *\fR\fRwlp\fR, \fBkrw_t\fR \fRenter_type
29 .fi

```

```

31 .LP
32 .nf
33 \fBvoid\fR \fBBrw_exit\fR(\fBkrwlock_t *\fR\fRwlp\fR);
34 .fi

```

```

36 .LP
37 .nf
38 \fBint\fR \fBBrw_tryenter\fR(\fBkrwlock_t *\fR\fRwlp\fR, \fBkrw_t\fR \fRenter_ty
39 .fi

```

```

41 .LP
42 .nf
43 \fBvoid\fR \fBBrw_downgrade\fR(\fBkrwlock_t *\fR\fRwlp\fR);
44 .fi

```

```

46 .LP
47 .nf
48 \fBint\fR \fBBrw_tryupgrade\fR(\fBkrwlock_t *\fR\fRwlp\fR);
49 .fi

```

```

51 .LP
52 .nf
53 \fBint\fR \fBBrw_read_locked\fR(\fBkrwlock_t *\fR\fRwlp\fR);
54 .fi

```

```

56 .SH INTERFACE LEVEL
57 .sp
58 .LP
59 Solaris DDI specific (Solaris DDI).
60 .SH PARAMETERS
61 .sp

```

```

62 .ne 2
63 .na
64 \fB\fRwlp\fR\fR
65 .ad
66 .RS 14n
67 Pointer to a \fBkrwlock_t\fR readers/writer lock.
68 .RE

```

```

70 .sp
71 .ne 2
72 .na
73 \fB\fRname\fR\fR
74 .ad
75 .RS 14n
76 Descriptive string. This is obsolete and should be \fBNNULL\fR. (Non-null
77 strings are legal, but they're a waste of kernel memory.)
78 .RE

```

```

80 .sp
81 .ne 2
82 .na
83 \fB\fRtype\fR\fR
84 .ad
85 .RS 14n
86 Type of readers/writer lock.
87 .RE

```

```

89 .sp
90 .ne 2
91 .na
92 \fB\fRarg\fR\fR
93 .ad
94 .RS 14n
95 Type-specific argument for initialization function.
96 .RE

```

```

98 .sp
99 .ne 2
100 .na
101 \fB\fRenter_type\fR\fR
102 .ad
103 .RS 14n
104 One of the values \fBWRITER\fR, \fBREADER\fR or
105 \fBREADER_STARVEWRITER\fR, indicating whether the
106 lock is to be acquired exclusively (\fBWRITER\fR), non-exclusively
107 (\fBREADER\fR) or non-exclusively without regard to any threads
108 that may be blocked on exclusive access (\fBREADER_STARVEWRITER\fR).
109 .RE

```

```

111 .SH DESCRIPTION
112 .sp
113 .LP
114 A multiple-readers, single-writer lock is represented by the \fBkrwlock_t\fR
115 data type. This type of lock will allow many threads to have simultaneous
116 read-only access to an object. Only one thread may have write access at any one
117 time. An object that is searched more frequently than it is changed is a good
118 candidate for a readers/writer lock.
119 .sp
120 .LP
121 Readers/writer locks are slightly more expensive than mutex locks, and the
122 advantage of multiple read access may not occur if the lock will only be held
123 for a short time.
124 .sp
125 .LP
126 The \fBBrw_init()\fR function initializes a readers/writer lock. It is an error
127 to initialize a lock more than once. The \fRtype\fR argument should be set to

```



```

128 \fBRW_DRIVER\fR. If the lock is used by the interrupt handler, the
129 type-specific argument, \fiarg\fR, should be the interrupt priority returned
130 from \fBddi_intr_get_pri\fR(9F) or \fBddi_intr_get_softint_pri\fR(9F). Note
131 that arg should be the value of the interrupt priority cast by calling the
132 \fBBDDI_INTR_PRI\fR macro. If the lock is not used by any interrupt handler, the
133 argument should be \fINULL.\fR
134 .sp
135 .LP
136 The \fBrw_destroy()\fR function releases any resources that might have been
137 allocated by \fBrw_init()\fR. It should be called before freeing the memory
138 containing the lock. The lock must not be held by any thread when it is
139 destroyed.
140 .sp
141 .LP
142 The \fBrw_enter()\fR function acquires the lock, and blocks if necessary. If
143 \fiEnter_type\fR is \fBRW_WRITER\fR, the caller blocks if any thread holds
144 the lock. If \fiEnter_type\fR is \fBRW_READER\fR, the caller blocks if there
145 is a writer or a thread attempting to enter for writing. If \fiEnter_type\fR
146 is \fBRW_READER_STARVEMWRITER\fR, the caller blocks only if there is a writer;
147 if the lock is held for reading and a thread is blocked attempting to enter
148 for writing, the caller will acquire the lock as a reader instead of
149 blocking on the pending writer.

151 .sp
152 .LP
153 NOTE: It is a programming error for any thread to acquire an rwlock as
154 \fBRW_READER\fR that it already holds. Doing so can deadlock the system: if
155 thread R acquires the lock as \fBRW_READER\fR, then thread W tries to acquire
156 the lock as a writer, W will set write-wanted and block. When R tries to get
157 its second read hold on the lock, it will honor the write-wanted bit and block
158 waiting for W; but W cannot run until R drops the lock. Thus threads R and W
159 deadlock. To opt out of this behavior -- that is, to safely allow a lock to
160 be grabbed recursively as a reader -- the lock should be acquired as
161 \fBRW_READER_STARVEMWRITER\fR, which will allow R to get its second read hold
162 without regard for the write-wanted bit set by W. Note that the
163 \fBRW_READER_STARVEMWRITER\fR behavior will starve writers in the presence of
164 infinite readers; it should be used with care, and only where the default
165 \fBRW_READER\fR behavior is unacceptable.
166 .sp
167 .LP
168 The \fBrw_exit()\fR function releases the lock and may wake up one or more
169 threads waiting on the lock.
170 .sp
171 .LP
172 The \fBrw_tryenter()\fR function attempts to enter the lock, like
173 \fBrw_enter()\fR, but never blocks. It returns a non-zero value if the lock was
174 successfully entered, and zero otherwise.
175 .sp
176 .LP
177 A thread that holds the lock exclusively (entered with \fBRW_WRITER\fR), may
178 call \fBrw_downgrade()\fR to convert to holding the lock non-exclusively (as if
179 entered with \fBRW_READER\fR). One or more waiting readers may be unblocked.
180 .sp
181 .LP
182 The \fBrw_tryupgrade()\fR function can be called by a thread that holds the
183 lock for reading to attempt to convert to holding it for writing. This upgrade
184 can only succeed if no other thread is holding the lock and no other thread is
185 blocked waiting to acquire the lock for writing.
186 .sp
187 .LP
188 The \fBrw_read_locked()\fR function returns non-zero if the calling thread
189 holds the lock for read, and zero if the caller holds the lock for write. The
190 caller must hold the lock. The system may panic if \fBrw_read_locked()\fR is
191 called for a lock that isn't held by the caller.
192 .SH RETURN VALUES
193 .sp

```

```

194 .ne 2
195 .na
196 \fB\b0\fR\fR
197 .ad
198 .RS 12n
199 \fBrw_tryenter()\fR could not obtain the lock without blocking.
200 .RE

202 .sp
203 .ne 2
204 .na
205 \fB\b0\fR\fR
206 .ad
207 .RS 12n
208 \fBrw_tryupgrade()\fR was unable to perform the upgrade because of other
209 threads holding or waiting to hold the lock.
210 .RE

212 .sp
213 .ne 2
214 .na
215 \fB\b0\fR\fR
216 .ad
217 .RS 12n
218 \fBrw_read_locked()\fR returns \fB\b0\fR if the lock is held by the caller for
219 write.
220 .RE

222 .sp
223 .ne 2
224 .na
225 \fB\bnon-zero\fR\fR
226 .ad
227 .RS 12n
228 from \fBrw_read_locked()\fR if the lock is held by the caller for read.
229 .RE

231 .sp
232 .ne 2
233 .na
234 \fB\bnon-zero\fR\fR
235 .ad
236 .RS 12n
237 successful return from \fBrw_tryenter()\fR or \fBrw_tryupgrade()\fR.
238 .RE

240 .SH CONTEXT
241 .sp
242 .LP
243 These functions can be called from user, interrupt, or kernel context, except
244 for \fBrw_init()\fR and \fBrw_destroy()\fR, which can be called from user
245 context only.
246 .SH SEE ALSO
247 .sp
248 .LP
249 \fBbcondvar\fR(9F), \fBddi_intr_alloc\fR(9F), \fBddi_intr_add_handler\fR(9F),
250 \fBddi_intr_get_pri\fR(9F), \fBddi_intr_get_softint_pri\fR(9F),
251 \fBbmutex\fR(9F), \fBbsemaphore\fR(9F)
252 .sp
253 .LP
254 \fIWriting Device Drivers\fR
255 .SH NOTES
256 .sp
257 .LP
258 Compiling with \fB_LOCKTEST\fR or \fB_MPSTATS\fR defined no longer has any
259 effect. To gather lock statistics, see \fBblockstat\fR(1M).

```

new/usr/src/man/man9f/scsi\_hba\_attach\_setup.9f

1

\*\*\*\*\*

7145 Thu Jul 17 00:50:43 2014

new/usr/src/man/man9f/scsi\_hba\_attach\_setup.9f

manpage lint.

\*\*\*\*\*

\_\_\_\_\_unchanged\_portion\_omitted\_\_\_\_\_

new/usr/src/tools/onbld/Checks/Makefile

1

\*\*\*\*\*

1425 Thu Jul 17 00:50:43 2014

new/usr/src/tools/onbld/Checks/Makefile

manpage lint.

\*\*\*\*\*

```
1 #
2 # CDDL HEADER START
3 #
4 # The contents of this file are subject to the terms of the
5 # Common Development and Distribution License (the "License").
6 # You may not use this file except in compliance with the License.
7 #
8 # You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9 # or http://www.opensolaris.org/os/licensing.
10 # See the License for the specific language governing permissions
11 # and limitations under the License.
12 #
13 # When distributing Covered Code, include this CDDL HEADER in each
14 # file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 # If applicable, add the following below this CDDL HEADER, with the
16 # fields enclosed by brackets "[]" replaced with your own identifying
17 # information: Portions Copyright [yyyy] [name of copyright owner]
18 #
19 # CDDL HEADER END
20 #
21 #
22 #
23 # Copyright (c) 2008, 2010, Oracle and/or its affiliates. All rights reserved.
24 #
25 #
26 # Copyright 2010, Richard Lowe
27 # Copyright 2014 Garrett D'Amore <garrett@damore.org>
28 #
29 include $(SRC)/Makefile.master
30 include ../../Makefile.tools
31 #
32 PYSRCS = \
33     CStyle.py      \
34     Cddl.py        \
35     CmtBlk.py     \
36     Comments.py   \
37     Copyright.py  \
38     DbLookups.py  \
39     HdrChk.py     \
40     JStyle.py     \
41     Keywords.py   \
42     ManLint.py    \
43     Mapfile.py    \
44     ProcessCheck.py \
45     __init__.py
46 #
47 PYOBS = $(PYSRCS:%.py=%.pyc)
48 PYTOPDIR = $(ROOTONBLDLIB)
49 PYMODDIR = onbld/Checks
50 #
51 include ../../Makefile.python
52 #
53 all: $(PYVERSOJBS)
54 #
55 install: all $(ROOTPYFILES)
56 #
57 clean:
58 #
59 clobber: clean pyclobber
```

```
*****  
1631 Thu Jul 17 00:50:43 2014  
new/usr/src/tools/onbld/Checks/ManLint.py  
manpage lint.  
*****
```

```
1 #  
2 # CDDL HEADER START  
3 #  
4 # The contents of this file are subject to the terms of the  
5 # Common Development and Distribution License (the "License").  
6 # You may not use this file except in compliance with the License.  
7 #  
8 # You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE  
9 # or http://www.opensolaris.org/os/licensing.  
10 # See the License for the specific language governing permissions  
11 # and limitations under the License.  
12 #  
13 # When distributing Covered Code, include this CDDL HEADER in each  
14 # file and include the License file at usr/src/OPENSOLARIS.LICENSE.  
15 # If applicable, add the following below this CDDL HEADER, with the  
16 # fields enclosed by brackets "[]" replaced with your own identifying  
17 # information: Portions Copyright [yyyy] [name of copyright owner]  
18 #  
19 # CDDL HEADER END  
20 #  
  
22 #  
23 # Copyright 2008 Sun Microsystems, Inc. All rights reserved.  
24 # Use is subject to license terms.  
25 #  
26 # Copyright 2014 Garrett D'Amore <garrett@damore.org>  
27 #  
  
29 #  
30 # ManLint, wrap the mandoc lint tool in a pythonic API  
31 #  
  
33 import sys  
34 from onbld.Checks.ProcessCheck import processcheck  
  
36 def manlint(fh, filename=None, output=sys.stderr, **opts):  
37     opttrans = { 'picky': None }  
  
39     for x in filter(lambda x: x not in opttrans, opts):  
40         raise TypeError('mandoc() got an unexpected keyword '  
41             'argument %s' % x)  
  
43     options = [opttrans[x] for x in opts if opts[x] and opttrans[x]]  
44     options.append('-Tlint')  
  
46     if not filename:  
47         filename = fh.name  
  
49     ret, tmpfile = processcheck('mandoc', options, fh, output)  
  
51     if tmpfile:  
52         for line in tmpfile:  
53             line = line.replace('<stdin>', filename)  
54             output.write(line)  
  
56     tmpfile.close()  
57     return ret
```

new/usr/src/tools/onbld/Checks/\_\_init\_\_.py

1

\*\*\*\*\*

1193 Thu Jul 17 00:50:43 2014

new/usr/src/tools/onbld/Checks/\_\_init\_\_.py

manpage lint.

\*\*\*\*\*

```
1 #! /usr/bin/python
2 #
3 # CDDL HEADER START
4 #
5 # The contents of this file are subject to the terms of the
6 # Common Development and Distribution License (the "License").
7 # You may not use this file except in compliance with the License.
8 #
9 # You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
10 # or http://www.opensolaris.org/os/licensing.
11 # See the License for the specific language governing permissions
12 # and limitations under the License.
13 #
14 # When distributing Covered Code, include this CDDL HEADER in each
15 # file and include the License file at usr/src/OPENSOLARIS.LICENSE.
16 # If applicable, add the following below this CDDL HEADER, with the
17 # fields enclosed by brackets "[]" replaced with your own identifying
18 # information: Portions Copyright [yyyy] [name of copyright owner]
19 #
20 # CDDL HEADER END
21 #
22 #
23 #
24 # Copyright 2009 Sun Microsystems, Inc. All rights reserved.
25 # Use is subject to license terms.
26 #
27 #
28 # Copyright 2010, Richard Lowe
29 # Copyright 2014 Garrett D'Amore <garrett@damore.org>
30 #
31 #
32 # The 'checks' package contains various checks that may be run
33 #
34 #
35 __all__ = [
36     'Cddl',
37     'Comments',
38     'Copyright',
39     'CStyle',
40     'HdrChk',
41     'JStyle',
42     'Keywords',
43     'ManLint',
44     'Mapfile']
```

```

*****
51625 Thu Jul 17 00:50:43 2014
new/usr/src/tools/onbld/hgext/cdm.py
manpage lint.
*****
1 #
2 # This program is free software; you can redistribute it and/or modify
3 # it under the terms of the GNU General Public License version 2
4 # as published by the Free Software Foundation.
5 #
6 # This program is distributed in the hope that it will be useful,
7 # but WITHOUT ANY WARRANTY; without even the implied warranty of
8 # MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
9 # GNU General Public License for more details.
10 #
11 # You should have received a copy of the GNU General Public License
12 # along with this program; if not, write to the Free Software
13 # Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
14 #
16 #
17 # Copyright (c) 2008, 2010, Oracle and/or its affiliates. All rights reserved.
18 # Copyright 2008, 2011 Richard Lowe
19 # Copyright 2014 Garrett D'Amore <garrett@damore.org>
20 #
22 '''OpenSolaris extensions to Mercurial
24 This extension contains a number of commands to help you work with
25 the OpenSolaris consolidations. It provides commands to check your
26 changes against the various style rules used for OpenSolaris, to
27 backup and restore your changes, to generate code reviews, and to
28 prepare your changes for integration.
31 The Parent
33 To provide a uniform notion of parent workspace regardless of
34 filesystem-based access, Cadmium uses the highest numbered changeset
35 on the current branch that is also in the parent workspace to
36 represent the parent workspace.
39 The Active List
41 Many Cadmium commands operate on the active list, the set of
42 files ('active files') you have changed in this workspace in relation
43 to its parent workspace, and the metadata (commentary, primarily)
44 associated with those changes.
47 NOT Files
49 Many of Cadmium's commands to check that your work obeys the
50 various stylistic rules of the OpenSolaris consolidations (such as
51 those run by 'hg nits') allow files to be excluded from this checking
52 by means of NOT files kept in the .hg/cdm/ directory of the Mercurial
53 repository for one-time exceptions, and in the exception_lists
54 directory at the repository root for permanent exceptions. (For ON,
55 these would mean one in $CODEMGR_WS and one in
56 $CODEMGR_WS/usr/closed).
58 These files are in the same format as the Mercurial hgignore
59 file, a description of which is available in the hgignore(5) manual
60 page.

```

```

63 Common Tasks
65 - Show diffs relative to parent workspace - pdiffs
66 - Check source style rules - nits
67 - Run pre-integration checks - pbchk
68 - Collapse all your changes into a single changeset - recommit
69 '''
71 import atexit, os, re, sys, stat, termios
74 #
75 # Adjust the load path based on the location of cdm.py and the version
76 # of python into which it is being loaded. This assumes the normal
77 # onbld directory structure, where cdm.py is in
78 # lib/python(version)?/onbld/hgext/. If that changes so too must
79 # this.
80 #
81 # This and the case below are not equivalent. In this case we may be
82 # loading a cdm.py in python2.X/ via the lib/python/ symlink but need
83 # python2.Y in sys.path.
84 #
85 sys.path.insert(1, os.path.join(os.path.dirname(__file__), "..", "..", "..",
86 "python%d.%d" % sys.version_info[:2]))
88 #
89 # Add the relative path from cdm.py to usr/src/tools to the load path,
90 # such that a cdm.py loaded from the source tree uses the modules also
91 # within the source tree.
92 #
93 sys.path.insert(2, os.path.join(os.path.dirname(__file__), "..", ".."))
95 from onbld.Scm import Version
96 from mercurial import util
98 try:
99 Version.check_version()
100 except Version.VersionMismatch, badversion:
101 raise util.Abort("Version Mismatch:\n %s\n" % badversion)
103 from mercurial import cmdutil, ignore, node, patch
105 from onbld.Scm.WorkSpace import Workspace, WorkList
106 from onbld.Scm.Backup import CdmBackup
107 from onbld.Checks import Cddl, Comments, Copyright, CStyle, HdrChk
108 from onbld.Checks import JStyle, Keywords, ManLint, Mapfile
109 from onbld.Checks import JStyle, Keywords, Mapfile
111 def yes_no(ui, msg, default):
112 if default:
113 prompt = '[Y/n]:'
114 defanswer = 'y'
115 else:
116 prompt = '[y/N]:'
117 defanswer = 'n'
119 if Version.at_least("1.4"):
120 index = ui.promptchoice(msg + prompt, ['&yes', '&no'],
121 default=['y', 'n'].index(defanswer))
122 resp = ('y', 'n')[index]
123 else:
124 resp = ui.prompt(msg + prompt, ['&yes', '&no'], default=defanswer)
126 return resp[0] in ('Y', 'y')

```

```

129 def buildfilelist(ws, parent, files):
130     '''Build a list of files in which we're interested.

132     If no files are specified take files from the active list relative
133     to 'parent'.

135     Return a list of 2-tuples the first element being a path relative
136     to the current directory and the second an entry from the active
137     list, or None if an explicit file list was given.'''

139     if files:
140         return [(path, None) for path in sorted(files)]
141     else:
142         active = ws.active(parent=parent)
143         return [(ws.filepath(e.name), e) for e in sorted(active)]
144 buildfilelist = util.cachefunc(buildfilelist)

147 def not_check(repo, cmd):
148     '''return a function which returns boolean indicating whether a file
149     should be skipped for CMD.'''

151     #
152     # The ignore routines need a canonical path to the file (relative to the
153     # repo root), whereas the check commands get paths relative to the cwd.
154     #
155     # Wrap our argument such that the path is canonified before it is checked.
156     #
157     def canonified_check(ignfunc):
158         def f(path):
159             cpath = util.canonpath(repo.root, repo.getcwd(), path)
160             return ignfunc(cpath)
161         return f

163     ignorefiles = []

165     for f in [repo.join('cdm/%s.NOT' % cmd),
166              repo.wjoin('exception_lists/%s' % cmd)]:
167         if os.path.exists(f):
168             ignorefiles.append(f)

170     if ignorefiles:
171         ign = ignore.ignore(repo.root, ignorefiles, repo.ui.warn)
172         return canonified_check(ign)
173     else:
174         return util.never

177 def abort_if_dirty(ws):
178     '''Abort if the workspace has uncommitted changes, merges,
179     branches, or has Mq patches applied'''

181     if ws.modified():
182         raise util.Abort('workspace has uncommitted changes')
183     if ws.merged():
184         raise util.Abort('workspace contains uncommitted merge')
185     if ws.branched():
186         raise util.Abort('workspace contains uncommitted branch')
187     if ws.mq_applied():
188         raise util.Abort('workspace has Mq patches applied')

191 #
192 # Adding a reference to Workspace from a repo causes a circular reference

```

```

193 # repo <-> Workspace.
194 #
195 # This prevents repo, Workspace and members thereof from being garbage
196 # collected. Since transactions are aborted when the transaction object
197 # is collected, and localrepo holds a reference to the most recently created
198 # transaction, this prevents transactions from cleanly aborting.
199 #
200 # Instead, we hold the repo->Workspace association in a dictionary, breaking
201 # that dependence.
202 #
203 wslst = {}

206 def reposetup(ui, repo):
207     if repo.local() and repo not in wslst:
208         wslst[repo] = Workspace(repo)

210     if ui.interactive() and sys.stdin.isatty():
211         ui.setconfig('hooks', 'preoutgoing.cdm_pbconfirm',
212                    'python:hgext_cdm_pbconfirm')

215 def pbconfirm(ui, repo, hooktype, source):
216     def wrapper(settings=None):
217         termios.tcsetattr(sys.stdin_FILENO(), termios.TCSANOW, settings)

219     if source == 'push':
220         if not yes_no(ui, "Are you sure you wish to push?", False):
221             return 1
222     else:
223         settings = termios.tcgetattr(sys.stdin_FILENO())
224         orig = list(settings)
225         atexit.register(wrapper, orig)
226         settings[3] = settings[3] & (~termios.ISIG) # c_lflag
227         termios.tcsetattr(sys.stdin_FILENO(), termios.TCSANOW, settings)

230 def cdm_pdiffs(ui, repo, *pats, **opts):
231     '''diff workspace against its parent

233     Show differences between this workspace and its parent workspace
234     in the same manner as 'hg diff'.

236     For a description of the changeset used to represent the parent
237     workspace, see The Parent in the extension documentation ('hg help
238     cdm').
239     '''

241     act = wslst[repo].active(opts.get('parent'))
242     if not act.revs:
243         return

245     #
246     # If no patterns were specified, either explicitly or via -I or -X
247     # use the active list files to avoid a workspace walk.
248     #
249     if pats or opts.get('include') or opts.get('exclude'):
250         matchfunc = wslst[repo].matcher(pats=pats, opts=opts)
251     else:
252         matchfunc = wslst[repo].matcher(files=act.files())

254     opts = patch.diffopts(ui, opts)
255     diffs = wslst[repo].diff(act.parenttip.node(), act.localtip.node(),
256                             match=matchfunc, opts=opts)
257     if diffs:
258         ui.write(diffs)

```

```

261 def cdm_list(ui, repo, **opts):
262     '''list active files (those changed in this workspace)

264     Display a list of files changed in this workspace as compared to
265     its parent workspace.

267     File names are displayed one-per line, grouped by manner in which
268     they changed (added, modified, removed). Information about
269     renames or copies is output in parentheses following the file
270     name.

272     For a description of the changeset used to represent the parent
273     workspace, see The Parent in the extension documentation ('hg help
274     cdm').

276     Output can be filtered by change type with --added, --modified,
277     and --removed. By default, all files are shown.
278     '''

280     act = wslist[repo].active(opts['parent'])
281     wanted = set(x for x in ('added', 'modified', 'removed') if opts[x])
282     changes = {}

284     for entry in act:
285         if wanted and (entry.change not in wanted):
286             continue

288         if entry.change not in changes:
289             changes[entry.change] = []
290             changes[entry.change].append(entry)

292     for change in sorted(changes.keys()):
293         ui.write(change + ':\n')

295         for entry in sorted(changes[change]):
296             if entry.is_renamed():
297                 ui.write('\t%s (renamed from %s)\n' % (entry.name,
298                                                         entry.parentname))
299             elif entry.is_copied():
300                 ui.write('\t%s (copied from %s)\n' % (entry.name,
301                                                         entry.parentname))
302             else:
303                 ui.write('\t%s\n' % entry.name)

306 def cdm_bugs(ui, repo, parent=None):
307     '''show all bug IDs referenced in changeset comments'''

309     act = wslist[repo].active(parent)

311     for elt in set(filter(Comments.isBug, act.comments())):
312         ui.write(elt + '\n')

315 def cdm_comments(ui, repo, parent=None):
316     '''show changeset commentary for all active changesets'''
317     act = wslist[repo].active(parent)

319     for elt in act.comments():
320         ui.write(elt + '\n')

323 def cdm_renamed(ui, repo, parent=None):
324     '''show renamed active files

```

```

326     Renamed files are shown in the format::

328         new-name old-name

330     One pair per-line.
331     '''

333     act = wslist[repo].active(parent)

335     for entry in sorted(filter(lambda x: x.is_renamed(), act)):
336         ui.write('%s %s\n' % (entry.name, entry.parentname))

339 def cdm_comchk(ui, repo, **opts):
340     '''check active changeset comment formatting

342     Check that active changeset comments conform to O/N rules.

344     Each comment line must contain either one bug or ARC case ID
345     followed by its synopsis, or credit an external contributor.
346     '''

348     active = wslist[repo].active(opts.get('parent'))

350     ui.write('Comments check:\n')

352     check_db = not opts.get('nocheck')
353     return Comments.comchk(active.comments(), check_db=check_db, output=ui)

356 def cdm_cddlchk(ui, repo, *args, **opts):
357     '''check for a valid CDDL header comment in all active files.

359     Check active files for a valid Common Development and Distribution
360     License (CDDL) block comment.

362     Newly added files are checked for a copy of the CDDL header
363     comment. Modified files are only checked if they contain what
364     appears to be an existing CDDL header comment.

366     Files can be excluded from this check using the cddlchk.NOT file.
367     See NOT Files in the extension documentation ('hg help cdm').
368     '''

370     filelist = buildfilelist(wslist[repo], opts.get('parent'), args)
371     exclude = not_check(repo, 'cddlchk')
372     lenient = True
373     ret = 0

375     ui.write('CDDL block check:\n')

377     for f, e in filelist:
378         if e and e.is_removed():
379             continue
380         elif (e or opts.get('honour_not')) and exclude(f):
381             ui.status('Skipping %s...\n' % f)
382             continue
383         elif e and e.is_added():
384             lenient = False
385         else:
386             lenient = True

388         fh = open(f, 'r')
389         ret |= Cddl.cddlchk(fh, lenient=lenient, output=ui)
390         fh.close()

```



```

391     return ret

394 def cdm_manlintchk(ui, repo, *args, **opts):
395     '''check for mandoc lint

397     Check for man page formatting errors.

399     Files can be excluded from this check using the manlint.NOT
400     file. See NOT Files in the extension documentation ('hg help
401     cdm').
402     '''

404     filelist = buildfilelist(wslst[repo], opts.get('parent'), args)
405     exclude = not_check(repo, 'manlint')
406     ret = 0

408     # Man pages are identified as having a suffix starting with a digit.
409     ManfileRE = re.compile(r'.*\.[0-9][a-z]*$', re.IGNORECASE)

411     ui.write('Man format check:\n')

413     for f, e in filelist:
414         if e and e.is_removed():
415             continue
416         elif (not ManfileRE.match(f)):
417             continue
418         elif (e or opts.get('honour_notls')) and exclude(f):
419             ui.status('Skipping %s...\n' % f)
420             continue

422         fh = open(f, 'r')
423         ret |= ManLint.manlint(fh, output=ui, picky=True)
424         fh.close()
425     return ret

428 def cdm_mapfilechk(ui, repo, *args, **opts):
429     '''check for a valid mapfile header block in active files

431     Check that all link-editor mapfiles contain the standard mapfile
432     header comment directing the reader to the document containing
433     Solaris object versioning rules (README.mapfile).

435     Files can be excluded from this check using the mapfilechk.NOT
436     file. See NOT Files in the extension documentation ('hg help
437     cdm').
438     '''

440     filelist = buildfilelist(wslst[repo], opts.get('parent'), args)
441     exclude = not_check(repo, 'mapfilechk')
442     ret = 0

444     # We are interested in examining any file that has the following
445     # in its final path segment:
446     # - Contains the word 'mapfile'
447     # - Begins with 'map.'
448     # - Ends with '.map'
449     # We don't want to match unless these things occur in final path segment
450     # because directory names with these strings don't indicate a mapfile.
451     # We also ignore files with suffixes that tell us that the files
452     # are not mapfiles.
453     MapfileRE = re.compile(r'.*((mapfile[^/]*)|(/map\.[^/]*)|(\.map))$',
454         re.IGNORECASE)
455     NotMapSuffixRE = re.compile(r'.*\.[ch]$', re.IGNORECASE)

```

```

457     ui.write('Mapfile comment check:\n')

459     for f, e in filelist:
460         if e and e.is_removed():
461             continue
462         elif (not MapfileRE.match(f)) or NotMapSuffixRE.match(f):
463             continue
464         elif (e or opts.get('honour_notls')) and exclude(f):
465             ui.status('Skipping %s...\n' % f)
466             continue

468         fh = open(f, 'r')
469         ret |= Mapfile.mapfilechk(fh, output=ui)
470         fh.close()
471     return ret

474 def cdm_copyright(ui, repo, *args, **opts):
475     '''check each active file for a current and correct copyright notice

477     Check that all active files have a correctly formed copyright
478     notice containing the current year.

480     See the Non-Formatting Considerations section of the OpenSolaris
481     Developer's Reference for more info on the correct form of
482     copyright notice.
483     (http://hub.opensolaris.org/bin/view/Community+Group+on/devref\_7#H723NonForm)

485     Files can be excluded from this check using the copyright.NOT file.
486     See NOT Files in the extension documentation ('hg help cdm').
487     '''

489     filelist = buildfilelist(wslst[repo], opts.get('parent'), args)
490     exclude = not_check(repo, 'copyright')
491     ret = 0

493     ui.write('Copyright check:\n')

495     for f, e in filelist:
496         if e and e.is_removed():
497             continue
498         elif (e or opts.get('honour_notls')) and exclude(f):
499             ui.status('Skipping %s...\n' % f)
500             continue

502         fh = open(f, 'r')
503         ret |= Copyright.copyright(fh, output=ui)
504         fh.close()
505     return ret

508 def cdm_hdrchk(ui, repo, *args, **opts):
509     '''check active C header files conform to the O/N header rules

511     Check that any added or modified C header files conform to the O/N
512     header rules.

514     See the section 'HEADER STANDARDS' in the hdrchk(1) manual page
515     for more information on the rules for O/N header file formatting.

517     Files can be excluded from this check using the hdrchk.NOT file.
518     See NOT Files in the extension documentation ('hg help cdm').
519     '''

521     filelist = buildfilelist(wslst[repo], opts.get('parent'), args)
522     exclude = not_check(repo, 'hdrchk')

```

```

523     ret = 0
525     ui.write('Header format check:\n')
527     for f, e in filelist:
528         if e and e.is_removed():
529             continue
530         elif not f.endswith('.h'):
531             continue
532         elif (e or opts.get('honour_not')) and exclude(f):
533             ui.status('Skipping %s...\n' % f)
534             continue
536         fh = open(f, 'r')
537         ret |= HdrChk.hdrchk(fh, lenient=True, output=ui)
538         fh.close()
539     return ret
542 def cdm_cstyle(ui, repo, *args, **opts):
543     '''check active C source files conform to the C Style Guide
545     Check that any added or modified C source file conform to the C
546     Style Guide.
548     See the C Style Guide for more information about correct C source
549     formatting.
550     (http://hub.opensolaris.org/bin/download/Community+Group+on/WebHome/cstyle.m)
552     Files can be excluded from this check using the cstyle.NOT file.
553     See NOT Files in the extension documentation ('hg help cdm').
554     '''
556     filelist = buildfilelist(wslst[repo], opts.get('parent'), args)
557     exclude = not_check(repo, 'cstyle')
558     ret = 0
560     ui.write('C style check:\n')
562     for f, e in filelist:
563         if e and e.is_removed():
564             continue
565         elif not (f.endswith('.c') or f.endswith('.h')):
566             continue
567         elif (e or opts.get('honour_not')) and exclude(f):
568             ui.status('Skipping %s...\n' % f)
569             continue
571         fh = open(f, 'r')
572         ret |= CStyle.cstyle(fh, output=ui,
573                             picky=True, check_posix_types=True,
574                             check_continuation=True)
575         fh.close()
576     return ret
579 def cdm_jstyle(ui, repo, *args, **opts):
580     '''check active Java source files for common stylistic errors
582     Files can be excluded from this check using the jstyle.NOT file.
583     See NOT Files in the extension documentation ('hg help cdm').
584     '''
586     filelist = buildfilelist(wslst[repo], opts.get('parent'), args)
587     exclude = not_check(repo, 'jstyle')
588     ret = 0

```

```

590     ui.write('Java style check:\n')
592     for f, e in filelist:
593         if e and e.is_removed():
594             continue
595         elif not f.endswith('.java'):
596             continue
597         elif (e or opts.get('honour_not')) and exclude(f):
598             ui.status('Skipping %s...\n' % f)
599             continue
601         fh = open(f, 'r')
602         ret |= JStyle.jstyle(fh, output=ui, picky=True)
603         fh.close()
604     return ret
607 def cdm_permchk(ui, repo, *args, **opts):
608     '''check the permissions of each active file
610     Check that the file permissions of each added or modified file do not
611     contain the executable bit.
613     Files can be excluded from this check using the permchk.NOT file.
614     See NOT Files in the extension documentation ('hg help cdm').
615     '''
617     filelist = buildfilelist(wslst[repo], opts.get('parent'), args)
618     exclude = not_check(repo, 'permchk')
619     exeFiles = []
621     ui.write('File permission check:\n')
623     for f, e in filelist:
624         if e and e.is_removed():
625             continue
626         elif (e or opts.get('honour_not')) and exclude(f):
627             ui.status('Skipping %s...\n' % f)
628             continue
630         mode = stat.S_IMODE(os.stat(f)[stat.ST_MODE])
631         if mode & stat.S_IEXEC:
632             exeFiles.append(f)
634     if len(exeFiles) > 0:
635         ui.write('Warning: the following active file(s) have executable mode '
636                '({+x} permission set,\nremove unless intentional:\n')
637         for fname in exeFiles:
638             ui.write(" %s\n" % fname)
640     return len(exeFiles) > 0
643 def cdm_tagchk(ui, repo, **opts):
644     '''check modification of workspace tags
646     Check for any modification of the repository's .hgtags file.
648     With the exception of the gatekeepers, nobody should introduce or
649     modify a repository's tags.
650     '''
652     active = wslst[repo].active(opts.get('parent'))
654     ui.write('Checking for new tags:\n')

```

```

656 if ".hgtags" in active:
657     tfile = wslst[repo].filepath('.hgtags')
658     ptip = active.parenttip.rev()

660     ui.write('Warning: Workspace contains new non-local tags.\n'
661             'Only gatekeepers should add or modify such tags.\n'
662             'Use the following commands to revert these changes:\n'
663             ' hg revert -r%d %s\n'
664             ' hg commit %s\n'
665             'You should also recommit before integration\n' %
666             (ptip, tfile, tfile))

668     return 1

670 return 0

673 def cdm_branchchk(ui, repo, **opts):
674     '''check for changes in number or name of branches

676     Check that the workspace contains only a single head, that it is
677     on the branch 'default', and that no new branches have been
678     introduced.
679     '''

681     ui.write('Checking for multiple heads (or branches):\n')

683     heads = set(repo.heads())
684     parents = set([x.node() for x in wslst[repo].workingctx().parents()])

686     #
687     # We care if there's more than one head, and those heads aren't
688     # identical to the dirstate parents (if they are identical, it's
689     # an uncommitted merge which mergechk will catch, no need to
690     # complain twice).
691     #
692     if len(heads) > 1 and heads != parents:
693         ui.write('Workspace has multiple heads (or branches):\n')
694         for head in [repo.changectx(head) for head in heads]:
695             ui.write(" %d:%s\t%s\n" %
696                     (head.rev(), str(head), head.description().splitlines()[0]))
697         ui.write('You must merge and recommit.\n')
698         return 1

700     ui.write('\nChecking for branch changes:\n')

702     if repo.dirstate.branch() != 'default':
703         ui.write("Warning: Workspace tip has named branch: '%s'\n"
704               "Only gatekeepers should push new branches.\n"
705               "Use the following commands to restore the branch name:\n"
706               " hg branch [-f] default\n"
707               " hg commit\n"
708               "You should also recommit before integration\n" %
709               (repo.dirstate.branch()))
710         return 1

712     branches = repo.branchtags().keys()
713     if len(branches) > 1:
714         ui.write('Warning: Workspace has named branches:\n')
715         for t in branches:
716             if t == 'default':
717                 continue
718             ui.write("\t%s\n" % t)

720     ui.write("Only gatekeepers should push new branches.\n")

```

```

721         "Use the following commands to remove extraneous branches.\n"
722         " hg branch [-f] default\n"
723         " hg commit"
724         "You should also recommit before integration\n")
725     return 1

727 return 0

730 def cdm_keywords(ui, repo, *args, **opts):
731     '''check active files for SCCS keywords

733     Check that any added or modified files do not contain SCCS keywords
734     (#ident lines, etc.).

736     Files can be excluded from this check using the keywords.NOT file.
737     See NOT Files in the extension documentation ('hg help cdm').
738     '''

740     filelist = buildfilelist(wslst[repo], opts.get('parent'), args)
741     exclude = not_check(repo, 'keywords')
742     ret = 0

744     ui.write('Keywords check:\n')

746     for f, e in filelist:
747         if e and e.is_removed():
748             continue
749         elif (e or opts.get('honour_nots')) and exclude(f):
750             ui.status('Skipping %s...\n' % f)
751             continue

753         fh = open(f, 'r')
754         ret |= Keywords.keywords(fh, output=ui)
755         fh.close()
756     return ret

759 #
760 # NB:
761 # There's no reason to hook this up as an invokable command, since
762 # we have 'hg status', but it must accept the same arguments.
763 #
764 def cdm_outchk(ui, repo, **opts):
765     '''Warn the user if they have uncommitted changes'''

767     ui.write('Checking for uncommitted changes:\n')

769     st = wslst[repo].modified()
770     if st:
771         ui.write('Warning: the following files have uncommitted changes:\n')
772         for elt in st:
773             ui.write(' %s\n' % elt)
774         return 1
775     return 0

778 def cdm_mergechk(ui, repo, **opts):
779     '''Warn the user if their workspace contains merges'''

781     active = wslst[repo].active(opts.get('parent'))

783     ui.write('Checking for merges:\n')

785     merges = filter(lambda x: len(x.parents()) == 2 and x.parents()[1],
786                    active.revs)

```

```

788     if merges:
789         ui.write('Workspace contains the following merges:\n')
790         for rev in merges:
791             desc = rev.description().splitlines()
792             ui.write(' %s:%s\t%s\n' %
793                     (rev.rev() or "working", str(rev),
794                      desc and desc[0] or "**** uncommitted change ****"))
795         return 1
796     return 0

799 def run_checks(ws, cmds, *args, **opts):
800     '''Run CMDS (with OPTS) over active files in WS'''

802     ret = 0

804     for cmd in cmds:
805         name = cmd.func_name.split('_')[1]
806         if not ws.ui.configbool('cdm', name, True):
807             ws.ui.status('Skipping %s check...\n' % name)
808         else:
809             ws.ui.pushbuffer()
810             result = cmd(ws.ui, ws.repo, honour_not=True, *args, **opts)
811             output = ws.ui.popbuffer()

813             ret |= result

815             if not ws.ui.quiet or result != 0:
816                 ws.ui.write(output, '\n')
817     return ret

820 def cdm_nits(ui, repo, *args, **opts):
821     '''check for stylistic nits in active files

823     Check each active file for basic stylistic errors.

825     The following checks are run over each active file (see 'hg help
826     <check>' for more information about each):

828     - copyright (copyright statements)
829     - cstyle (C source style)
830     - hdrchk (C header style)
831     - jstyle (java source style)
832     - manlint (man page formatting)
833     - mapfilechk (link-editor mapfiles)
834     - permchk (file permissions)
835     - keywords (SCCS keywords)

837     With the global -q/--quiet option, only provide output for those
838     checks which fail.
839     '''

841     cmds = [cdm_copyright,
842             cdm_cstyle,
843             cdm_hdrchk,
844             cdm_jstyle,
845             cdm_manlintchk,
846             cdm_mapfilechk,
847             cdm_permchk,
848             cdm_keywords]

850     return run_checks(wslst[repo], cmds, *args, **opts)

```

```

853 def cdm_pbchk(ui, repo, **opts):
854     '''run pre-integration checks on this workspace

856     Check this workspace for common errors prior to integration.

858     The following checks are run over the active list (see 'hg help
859     <check>' for more information about each):

861     - branchchk (addition/modification of branches)
862     - comchk (changeset descriptions)
863     - copyright (copyright statements)
864     - cstyle (C source style)
865     - hdrchk (C header style)
866     - jstyle (java source style)
867     - keywords (SCCS keywords)
868     - manlint (man page formatting)
869     - mapfilechk (link-editor mapfiles)
870     - permchk (file permissions)
871     - tagchk (addition/modification of tags)

873     Additionally, the workspace is checked for outgoing merges (which
874     should be removed with 'hg recommit'), and uncommitted changes.

876     With the global -q/--quiet option, only provide output for those
877     checks which fail.
878     '''

880     #
881     # The current ordering of these is that the commands from cdm_nits
882     # run first in the same order as they would in cdm_nits, then the
883     # pbchk specifics are run.
884     #
885     cmds = [cdm_copyright,
886             cdm_cstyle,
887             cdm_hdrchk,
888             cdm_jstyle,
889             cdm_manlintchk,
890             cdm_mapfilechk,
891             cdm_permchk,
892             cdm_keywords,
893             cdm_comchk,
894             cdm_tagchk,
895             cdm_branchchk,
896             cdm_outchk,
897             cdm_mergechk]

899     return run_checks(wslst[repo], cmds, **opts)

902 def cdm_recommit(ui, repo, **opts):
903     '''replace outgoing changesets with a single equivalent changeset

905     Replace all outgoing changesets with a single changeset containing
906     equivalent changes. This removes uninteresting changesets created
907     during development that would only serve as noise in the gate.

909     Any changed file that is now identical in content to that in the
910     parent workspace (whether identical in history or otherwise) will
911     not be included in the new changeset. Any merges information will
912     also be removed.

914     If no files are changed in comparison to the parent workspace, the
915     outgoing changesets will be removed, but no new changeset created.

917     recommit will refuse to run if the workspace contains more than
918     one outgoing head, even if those heads are on the same branch. To

```

```

919 recommit with only one branch containing outgoing changesets, your
920 workspace must be on that branch and at that branch head.

922 recommit will prompt you to take a backup if your workspace has
923 been changed since the last backup was taken. In almost all
924 cases, you should allow it to take one (the default).

926 recommit cannot be run if the workspace contains any uncommitted
927 changes, applied Mq patches, or has multiple outgoing heads (or
928 branches).
929 '''

931 ws = wslist[repo]

933 if not os.getcwd().startswith(repo.root):
934     raise util.Abort('recommit is not safe to run with -R')

936 abort_if_dirty(ws)

938 wlock = repo.wlock()
939 lock = repo.lock()

941 try:
942     parent = ws.parent(opts['parent'])
943     between = repo.changelog.nodesbetween(ws.findoutgoing(parent))[2]
944     heads = set(between) & set(repo.heads())

946     if len(heads) > 1:
947         ui.warn('Workspace has multiple outgoing heads (or branches):\n')
948         for head in sorted(map(repo.changelog.rev, heads), reverse=True):
949             ui.warn('\t%d\n' % head)
950         raise util.Abort('you must merge before recommitting')

952     #
953     # We can safely use the worklist here, as we know (from the
954     # abort_if_dirty() check above) that the working copy has not been
955     # modified.
956     #
957     active = ws.active(parent)

959     if filter(lambda b: len(b.parents()) > 1, active.bases()):
960         raise util.Abort('Cannot recommit a merge of two non-outgoing '
961             'changesets')

963     if len(active.revs) <= 0:
964         raise util.Abort("no changes to recommit")

966     if len(active.files()) <= 0:
967         ui.warn("Recommitting %d active changesets, but no active files\n" %
968             len(active.revs))

970     #
971     # During the course of a recommit, any file bearing a name
972     # matching the source name of any renamed file will be
973     # clobbered by the operation.
974     #
975     # As such, we ask the user before proceeding.
976     #
977     bogosity = [f.parentname for f in active if f.is_renamed() and
978                 os.path.exists(repo.wjoin(f.parentname))]
979     if bogosity:
980         ui.warn("The following file names are the original name of a "
981             "rename and also present\n"
982             "in the working directory:\n")

984     for fname in bogosity:

```

```

985         ui.warn(" %s\n" % fname)

987         if not yes_no(ui, "These files will be removed by recommit."
988             " Continue?",
989             False):
990             raise util.Abort("recommit would clobber files")

992     user = opts['user'] or ui.username()
993     comments = '\n'.join(active.comments())

995     message = cmdutil.logmessage(opts) or ui.edit(comments, user)
996     if not message:
997         raise util.Abort('empty commit message')

999     bk = CdmBackup(ui, ws, backup_name(repo.root))
1000     if bk.need_backup():
1001         if yes_no(ui, 'Do you want to backup files first?', True):
1002             bk.backup()

1004     oldtags = repo.tags()
1005     clearedtags = [(name, nd, repo.changelog.rev(nd), local)
1006                   for name, nd, local in active.tags()]

1008     ws.squishdeltas(active, message, user=user)
1009     finally:
1010         lock.release()
1011         wlock.release()

1013     if clearedtags:
1014         ui.write("Removed tags:\n")
1015         for name, nd, rev, local in sorted(clearedtags,
1016             key=lambda x: x[0].lower()):
1017             ui.write(" %5s:%s:\t%s%s\n" % (rev, node.short(nd),
1018                 name, (local and ' (local)' or '')))

1020     for ntag, nnode in sorted(repo.tags().items(),
1021         key=lambda x: x[0].lower()):
1022         if ntag in oldtags and ntag != "tip":
1023             if oldtags[ntag] != nnode:
1024                 ui.write("tag '%s' now refers to revision %d:%s\n" %
1025                     (ntag, repo.changelog.rev(nnode),
1026                     node.short(nnode)))

1029 def do_eval(cmd, files, root, changedir=True):
1030     if not changedir:
1031         os.chdir(root)

1033     for path in sorted(files):
1034         dirn, base = os.path.split(path)

1036         if changedir:
1037             os.chdir(os.path.join(root, dirn))

1039         os.putenv('workspace', root)
1040         os.putenv('filepath', path)
1041         os.putenv('dir', dirn)
1042         os.putenv('file', base)
1043         os.system(cmd)

1046 def cdm_eval(ui, repo, *command, **opts):
1047     '''run specified command for each active file

1049     Run the command specified on the command line for each active
1050     file, with the following variables present in the environment:

```

```

1052     :$file:      - active file basename.
1053     :$dir:       - active file dirname.
1054     :$filepath:  - path from workspace root to active file.
1055     :$workspace: - full path to workspace root.

1057 For example:

1059     hg eval 'echo $dir; hg log -l3 $file'

1061 will show the last the 3 log entries for each active file,
1062 preceded by its directory.
1063 '''

1065 act = wslist[repo].active(opts['parent'])
1066 cmd = ' '.join(command)
1067 files = [x.name for x in act if not x.is_removed()]

1069 do_eval(cmd, files, repo.root, not opts['remain'])

1072 def cdm_apply(ui, repo, *command, **opts):
1073     '''apply specified command to all active files

1075     Run the command specified on the command line over each active
1076     file.

1078     For example 'hg apply "wc -l"' will output a count of the lines in
1079     each active file.
1080     '''

1082     act = wslist[repo].active(opts['parent'])

1084     if opts['remain']:
1085         appnd = ' $filepath'
1086     else:
1087         appnd = ' $file'

1089     cmd = ' '.join(command) + appnd
1090     files = [x.name for x in act if not x.is_removed()]

1092     do_eval(cmd, files, repo.root, not opts['remain'])

1095 def cdm_reparent(ui, repo, parent):
1096     '''reparent your workspace

1098     Update the 'default' path alias that is used as the default source
1099     for 'hg pull' and the default destination for 'hg push' (unless
1100     there is a 'default-push' alias). This is also the path all
1101     Cadmium commands treat as your parent workspace.
1102     '''

1104     def append_new_parent(parent):
1105         fp = None
1106         try:
1107             fp = repo.opener('hgrc', 'a', atomictemp=True)
1108             if fp.tell() != 0:
1109                 fp.write('\n')
1110                 fp.write('[paths]\n'
1111                        'default = %s\n\n' % parent)
1112             fp.rename()
1113         finally:
1114             if fp and not fp.closed:
1115                 fp.close()

```

```

1117     def update_parent(path, line, parent):
1118         line = line - 1 # The line number we're passed will be 1-based
1119         fp = None

1121         try:
1122             fp = open(path)
1123             data = fp.readlines()
1124         finally:
1125             if fp and not fp.closed:
1126                 fp.close()

1128         #
1129         # line will be the last line of any continued block, go back
1130         # to the first removing the continuation as we go.
1131         #
1132         while data[line][0].isspace():
1133             data.pop(line)
1134             line -= 1

1136         assert data[line].startswith('default')

1138         data[line] = "default = %s\n" % parent
1139         if data[-1] != '\n':
1140             data.append('\n')

1142         try:
1143             fp = util.atomictempfile(path, 'w', 0644)
1144             fp.writelines(data)
1145             fp.rename()
1146         finally:
1147             if fp and not fp.closed:
1148                 fp.close()

1150     from mercurial import config
1151     parent = ui.expandpath(parent)

1153     if not os.path.exists(repo.join('hgrc')):
1154         append_new_parent(parent)
1155         return

1157     cfg = config.config()
1158     cfg.read(repo.join('hgrc'))
1159     source = cfg.source('paths', 'default')

1161     if not source:
1162         append_new_parent(parent)
1163         return
1164     else:
1165         path, target = source.rsplit(':', 1)

1167         if path != repo.join('hgrc'):
1168             raise util.Abort("Cannot edit path specification not in repo hgrc\n"
1169                             "default path is from: %s" % source)

1171         update_parent(path, int(target), parent)

1174 def backup_name(fullpath):
1175     '''Create a backup directory name based on the specified path.

1177     In most cases this is the basename of the path specified, but
1178     certain cases are handled specially to create meaningful names'''

1180     special = ['usr/closed']

1182     fullpath = fullpath.rstrip(os.path.sep).split(os.path.sep)

```

```

1184 #
1185 # If a path is 'special', we append the basename of the path to
1186 # the path element preceding the constant, special, part.
1187 #
1188 # Such that for instance:
1189 # /foo/bar/onnv-fixes/usr/closed
1190 # has a backup name of:
1191 # onnv-fixes-closed
1192 #
1193 for elt in special:
1194     elt = elt.split(os.path.sep)
1195     pathpos = len(elt)
1197     if fullpath[-pathpos:] == elt:
1198         return "%s-%s" % (fullpath[-pathpos - 1], elt[-1])
1199 else:
1200     return fullpath[-1]
1203 def cdm_backup(ui, repo, if_newer=False):
1204     '''backup workspace changes and metadata
1206     Create a backup copy of changes made in this workspace as compared
1207     to its parent workspace, as well as important metadata of this
1208     workspace.
1210     NOTE: Only changes as compared to the parent workspace are backed
1211     up. If you lose this workspace and its parent, you will not be
1212     able to restore a backup into a clone of the grandparent
1213     workspace.
1215     By default, backups are stored in the cdm.backup/ directory in
1216     your home directory. This is configurable using the cdm.backupdir
1217     configuration variable, for example:
1219     hg backup --config cdm.backupdir=/net/foo/backups
1221     or place the following in an appropriate hgrc file::
1223     [cdm]
1224     backupdir = /net/foo/backups
1226     Backups have the same name as the workspace in which they were
1227     taken, with '-closed' appended in the case of O/N's usr/closed.
1228     '''
1230     name = backup_name(repo.root)
1231     bk = CdmBackup(ui, wslst[repo], name)
1233     wlock = repo.wlock()
1234     lock = repo.lock()
1236     try:
1237         if if_newer and not bk.need_backup():
1238             ui.status('backup is up-to-date\n')
1239         else:
1240             bk.backup()
1241     finally:
1242         lock.release()
1243         wlock.release()
1246 def cdm_restore(ui, repo, backup, **opts):
1247     '''restore workspace from backup

```

```

1249     Restore this workspace from a backup (taken by 'hg backup').
1251     If the specified backup directory does not exist, it is assumed to
1252     be relative to the cadmium backup directory (~/.cdm.backup/ by
1253     default).
1255     For example::
1257         % hg restore on-rfe - Restore the latest backup of ~/.cdm.backup/on-rfe
1258         % hg restore -g3 on-rfe - Restore the 3rd backup of ~/.cdm.backup/on-rfe
1259         % hg restore /net/foo/backup/on-rfe - Restore from an explicit path
1260     '''
1262     if not os.getcwd().startswith(repo.root):
1263         raise util.Abort('restore is not safe to run with -R')
1265     abort_if_dirty(wslst[repo])
1267     if opts['generation']:
1268         gen = int(opts['generation'])
1269     else:
1270         gen = None
1272     if os.path.exists(backup):
1273         backup = os.path.abspath(backup)
1275     wlock = repo.wlock()
1276     lock = repo.lock()
1278     try:
1279         bk = CdmBackup(ui, wslst[repo], backup)
1280         bk.restore(gen)
1281     finally:
1282         lock.release()
1283         wlock.release()
1286 def cdm_webrev(ui, repo, **opts):
1287     '''generate web-based code review and optionally upload it
1289     Generate a web-based code review using webrev(1) and optionally
1290     upload it. All known arguments are passed through to webrev(1).
1291     '''
1293     webrev_args = ""
1294     for key in opts.keys():
1295         if opts[key]:
1296             if type(opts[key]) == type(True):
1297                 webrev_args += '-' + key + ' '
1298             else:
1299                 webrev_args += '-' + key + ' ' + opts[key] + ' '
1301     retval = os.system('webrev ' + webrev_args)
1302     if retval != 0:
1303         return retval - 255
1305     return 0
1308 def cdm_debugcdmal(ui, repo, *pats, **opts):
1309     '''dump the active list for the sake of debugging/testing'''
1311     ui.write(wslst[repo].active(opts['parent']).as_text(pats))
1314 def cdm_changed(ui, repo, *pats, **opts):

```

```

1315     '''mark a file as changed in the working copy
1317
1317 Maintain a list of files checked for modification in the working
1318 copy. If the list exists, most cadmium commands will only check
1319 the working copy for changes to those files, rather than checking
1320 the whole workspace (this does not apply to committed changes,
1321 which are always seen).
1323
1323 Since this list functions only as a hint as to where in the
1324 working copy to look for changes, entries that have not actually
1325 been modified (in the working copy, or in general) are not
1326 problematic.
1329
1329 Note: If such a list exists, it must be kept up-to-date.
1332
1332 Renamed files can be added with reference only to their new name:
1333     $ hg mv foo bar
1334     $ hg changed bar
1336
1336 Without arguments, 'hg changed' will list all files recorded as
1337 altered, such that, for instance:
1338     $ hg status $(hg changed)
1339     $ hg diff $(hg changed)
1340 Become useful (generally faster than their unadorned counterparts)
1342
1342 To create an initially empty list:
1343     $ hg changed -i
1344 Until files are added to the list it is equivalent to saying
1345 "Nothing has been changed"
1347
1347 Update the list based on the current active list:
1348     $ hg changed -u
1349 The old list is emptied, and replaced with paths from the
1350 current active list.
1352
1352 Remove the list entirely:
1353     $ hg changed -d
1354     '''
1356
1356 def modded_files(repo, parent):
1357     out = wslist[repo].findoutgoing(wslist[repo].parent(parent))
1358     outnodes = repo.changelog.nodesbetween(out)[0]
1360
1360     files = set()
1361     for n in outnodes:
1362         files.update(repo.changectx(n).files())
1364
1364     files.update(wslist[repo].status().keys())
1365     return files
1367
1367 #
1368 # specced_pats is convenient to treat as a boolean indicating
1369 # whether any file patterns or paths were specified.
1370 #
1371 specced_pats = pats or opts['include'] or opts['exclude']
1372 if len(filter(None, [opts['delete'], opts['update'], opts['init'],
1373 specced_pats])) > 1:
1374     raise util.Abort("-d, -u, -i and patterns are mutually exclusive")
1376
1376 wl = WorkList(wslist[repo])
1378
1378 if (not wl and specced_pats) or opts['init']:
1379     wl.delete()
1380     if yes_no(ui, "Create a list based on your changes thus far?", True):

```

```

1381         map(wl.add, modded_files(repo, opts.get('parent')))
1383
1383 if opts['delete']:
1384     wl.delete()
1385 elif opts['update']:
1386     wl.delete()
1387     map(wl.add, modded_files(repo, opts.get('parent')))
1388     wl.write()
1389 elif opts['init']:           # Any possible old list was deleted above
1390     wl.write()
1391 elif specced_pats:
1392     sources = []
1394
1394     match = wslist[repo].matcher(pats=pats, opts=opts)
1395     for abso in repo.walk(match):
1396         if abso in repo.dirstate:
1397             wl.add(abso)
1398             #
1399             # Store the source name of any copy. We use this so
1400             # both the add and delete of a rename can be entered
1401             # into the WorkList with only the destination name
1402             # explicitly being mentioned.
1403             #
1404             fctx = wslist[repo].workingctx().filectx(abso)
1405             rn = fctx.renamed()
1406             if rn:
1407                 sources.append(rn[0])
1408             else:
1409                 ui.warn("%s is not version controlled -- skipping\n" %
1410 match.rel(abso))
1412
1412     if sources:
1413         for fname, chng in wslist[repo].status(files=sources).iteritems():
1414             if chng == 'removed':
1415                 wl.add(fname)
1416         wl.write()
1417     else:
1418         for elt in sorted(wl.list()):
1419             ui.write("%s\n" % wslist[repo].filepath(elt))
1422
1422 cmdtable = {
1423     'apply': (cdm_apply, [(('p', 'parent', '', 'parent workspace'),
1424 ('r', 'remain', None, 'do not change directory')),
1425 'hg apply [-p PARENT] [-r] command...'),
1426 '^backup|bu': (cdm_backup, [(('t', 'if-newer', None,
1427 'only backup if workspace files are newer')),
1428 'hg backup [-t]'),
1429 'branchchk': (cdm_branchchk, [(('p', 'parent', '', 'parent workspace')),
1430 'hg branchchk [-p PARENT]'),
1431 'bugs': (cdm_bugs, [(('p', 'parent', '', 'parent workspace')),
1432 'hg bugs [-p PARENT]'),
1433 'cddlchk': (cdm_cddlchk, [(('p', 'parent', '', 'parent workspace')),
1434 'hg cddlchk [-p PARENT]'),
1435 'changed': (cdm_changed, [(('d', 'delete', None, 'delete the file list'),
1436 ('u', 'update', None, 'mark all changed files'),
1437 ('i', 'init', None, 'create an empty file list'),
1438 ('p', 'parent', '', 'parent workspace'),
1439 ('I', 'include', [],
1440 'include names matching the given patterns'),
1441 ('X', 'exclude', [],
1442 'exclude names matching the given patterns')),
1443 'hg changed -d\n'
1444 'hg changed -u\n'
1445 'hg changed -i\n'
1446 'hg changed [-I PATTERN...] [-X PATTERN...] [FILE...]),

```



```

1447     'comchk': (cdm_comchk, [('p', 'parent', '', 'parent workspace'),
1448                          ('N', 'nocheck', None,
1449                           'do not compare comments with databases')],
1450              'hg comchk [-p PARENT]'),
1451     'comments': (cdm_comments, [('p', 'parent', '', 'parent workspace')],
1452                      'hg comments [-p PARENT]'),
1453     'copyright': (cdm_copyright, [('p', 'parent', '', 'parent workspace')],
1454                      'hg copyright [-p PARENT]'),
1455     'cstyle': (cdm_cstyle, [('p', 'parent', '', 'parent workspace')],
1456                  'hg cstyle [-p PARENT]'),
1457     'debugcdmal': (cdm_debugcdmal, [('p', 'parent', '', 'parent workspace')],
1458                      'hg debugcdmal [-p PARENT] [FILE...]',
1459                      ('r', 'remain', None, 'do not change directory')),
1460     'eval': (cdm_eval, [('p', 'parent', '', 'parent workspace'),
1461                       ('r', 'remain', None, 'do not change directory')],
1462              'hg eval [-p PARENT] [-r] command...'),
1463     'hdrchk': (cdm_hdrchk, [('p', 'parent', '', 'parent workspace')],
1464                 'hg hdrchk [-p PARENT]'),
1465     'jstyle': (cdm_jstyle, [('p', 'parent', '', 'parent workspace')],
1466                 'hg jstyle [-p PARENT]'),
1467     'keywords': (cdm_keywords, [('p', 'parent', '', 'parent workspace'),
1468                                'hg keywords [-p PARENT]'),
1469     '^list[active]': (cdm_list, [('p', 'parent', '', 'parent workspace'),
1470                                ('a', 'added', None, 'show added files'),
1471                                ('m', 'modified', None, 'show modified files'),
1472                                ('r', 'removed', None, 'show removed files')],
1473                      'hg list [-amrRu] [-p PARENT]'),
1474     'manlint': (cdm_manlintchk, [('p', 'parent', '', 'parent workspace')],
1475                 'hg manlint [-p PARENT]'),
1476     'mapfilechk': (cdm_mapfilechk, [('p', 'parent', '', 'parent workspace')],
1477                      'hg mapfilechk [-p PARENT]'),
1478     '^nits': (cdm_nits, [('p', 'parent', '', 'parent workspace')],
1479               'hg nits [-p PARENT]'),
1480     '^pbchk': (cdm_pbchk, [('p', 'parent', '', 'parent workspace'),
1481                          ('N', 'nocheck', None, 'skip database checks')],
1482               'hg pbchk [-N] [-p PARENT]'),
1483     'permchk': (cdm_permchk, [('p', 'parent', '', 'parent workspace')],
1484                 'hg permchk [-p PARENT]'),
1485     '^pdiffs': (cdm_pdiffs, [('p', 'parent', '', 'parent workspace'),
1486                             ('a', 'text', None, 'treat all files as text'),
1487                             ('g', 'git', None, 'use extended git diff format'),
1488                             ('w', 'ignore-all-space', None,
1489                              'ignore white space when comparing lines'),
1490                             ('b', 'ignore-space-change', None,
1491                              'ignore changes in the amount of white space'),
1492                             ('B', 'ignore-blank-lines', None,
1493                              'ignore changes whose lines are all blank'),
1494                             ('U', 'unified', 3,
1495                              'number of lines of context to show'),
1496                             ('I', 'include', [],
1497                              'include names matching the given patterns'),
1498                             ('X', 'exclude', [],
1499                              'exclude names matching the given patterns')],
1500               'hg pdiffs [OPTION...] [-p PARENT] [FILE...]',
1501     '^recommit[reci]': (cdm_recommit, [('p', 'parent', '', 'parent workspace'),
1502                                       ('m', 'message', '',
1503                                        'use <text> as commit message'),
1504                                       ('l', 'logfile', '',
1505                                        'read commit message from file'),
1506                                       ('u', 'user', '',
1507                                        'record user as committer')],
1507                       'hg recommit [-m TEXT] [-l FILE] [-u USER] [-p PARENT]'),
1508     'renamed': (cdm_renamed, [('p', 'parent', '', 'parent workspace')],
1509                 'hg renamed [-p PARENT]'),
1510     'reparent': (cdm_reparent, [], 'hg reparent PARENT'),
1511     '^restore': (cdm_restore, [('g', 'generation', '', 'generation number')],
1512                      'hg restore [-g GENERATION] BACKUP'),

```

```

1513     'tagchk': (cdm_tagchk, [('p', 'parent', '', 'parent workspace')],
1514                'hg tagchk [-p PARENT]'),
1515     'webrev': (cdm_webrev, [('C', 'C', '', 'ITS priority file'),
1516                            ('D', 'D', '', 'delete remote webrev'),
1517                            ('I', 'I', '', 'ITS configuration file'),
1518                            ('i', 'i', '', 'include file'),
1519                            ('N', 'N', None, 'suppress comments'),
1520                            ('n', 'n', None, 'do not generate webrev'),
1521                            ('O', 'O', None, 'OpenSolaris mode'),
1522                            ('o', 'o', '', 'output directory'),
1523                            ('p', 'p', '', 'use specified parent'),
1524                            ('t', 't', '', 'upload target'),
1525                            ('U', 'U', None, 'upload the webrev'),
1526                            ('w', 'w', '', 'use wx active file')],
1527              'hg webrev [WEBREV_OPTIONS]'),
1528 }

```

unchanged portion omitted

new/usr/src/tools/scripts/git-pbchk.py

1

```
*****
10980 Thu Jul 17 00:50:43 2014
new/usr/src/tools/scripts/git-pbchk.py
manpage lint.
*****
1 #!/usr/bin/python2.6
2 #
3 # This program is free software; you can redistribute it and/or modify
4 # it under the terms of the GNU General Public License version 2
5 # as published by the Free Software Foundation.
6 #
7 # This program is distributed in the hope that it will be useful,
8 # but WITHOUT ANY WARRANTY; without even the implied warranty of
9 # MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
10 # GNU General Public License for more details.
11 #
12 # You should have received a copy of the GNU General Public License
13 # along with this program; if not, write to the Free Software
14 # Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
15 #
17 #
18 # Copyright (c) 2008, 2010, Oracle and/or its affiliates. All rights reserved.
19 # Copyright 2008, 2012 Richard Lowe
20 # Copyright 2014 Garrett D'Amore <garrett@damore.org>
21 #
23 import getopt
24 import os
25 import re
26 import subprocess
27 import sys
28 import tempfile
30 from cStringIO import StringIO
32 # This is necessary because, in a fit of pique, we used hg-format ignore lists
33 # for NOT files.
34 from mercurial import ignore
36 #
37 # Adjust the load path based on our location and the version of python into
38 # which it is being loaded. This assumes the normal onbld directory
39 # structure, where we are in bin/ and the modules are in
40 # lib/python(version)?/onbld/Scm/. If that changes so too must this.
41 #
42 sys.path.insert(1, os.path.join(os.path.dirname(__file__), "..", "lib",
43                               "python%d.%d" % sys.version_info[:2]))
45 #
46 # Add the relative path to usr/src/tools to the load path, such that when run
47 # from the source tree we use the modules also within the source tree.
48 #
49 sys.path.insert(2, os.path.join(os.path.dirname(__file__), ".."))
51 from onbld.Checks import Comments, Copyright, CStyle, HdrChk
52 from onbld.Checks import JStyle, Keywords, ManLint, Mapfile
53 from onbld.Checks import JStyle, Keywords, Mapfile
55 class GitError(Exception):
56     pass
58 def git(command):
59     """Run a command and return a stream containing its stdout (and write its
60     stderr to its stdout)"""
```

new/usr/src/tools/scripts/git-pbchk.py

2

```
62 if type(command) != list:
63     command = command.split()
65 command = ["git"] + command
67 try:
68     tmpfile = tempfile.TemporaryFile(prefix="git-nits")
69 except EnvironmentError, e:
70     raise GitError("Could not create temporary file: %s\n" % e)
72 try:
73     p = subprocess.Popen(command,
74                          stdout=tmpfile,
75                          stderr=subprocess.STDOUT)
76 except OSError, e:
77     raise GitError("could not execute %s: %s\n" (command, e))
79 err = p.wait()
80 if err != 0:
81     raise GitError(p.stdout.read())
83 tmpfile.seek(0)
84 return tmpfile
87 def git_root():
88     """Return the root of the current git workspace"""
90 p = git('rev-parse --git-dir')
92 if not p:
93     sys.stderr.write("Failed finding git workspace\n")
94     sys.exit(err)
96 return os.path.abspath(os.path.join(p.readlines()[0],
97                                     os.path.pardir))
100 def git_branch():
101     """Return the current git branch"""
103 p = git('branch')
105 if not p:
106     sys.stderr.write("Failed finding git branch\n")
107     sys.exit(err)
109 for elt in p:
110     if elt[0] == '*':
111         if elt.endswith('(no branch)'):
112             return None
113         return elt.split()[1]
116 def git_parent_branch(branch):
117     """Return the parent of the current git branch.
119 If this branch tracks a remote branch, return the remote branch which is
120 tracked. If not, default to origin/master."""
122 if not branch:
123     return None
125 p = git("for-each-ref --format=%(refname:short) %(upstream:short) " +
126         "refs/heads/")
```

```

128     if not p:
129         sys.stderr.write("Failed finding git parent branch\n")
130         sys.exit(err)

132     for line in p:
133         # Git 1.7 will leave a ' ' trailing any non-tracking branch
134         if ' ' in line and not line.endswith('\n'):
135             local, remote = line.split()
136             if local == branch:
137                 return remote
138     return 'origin/master'

141 def git_comments(parent):
142     """Return a list of any checkin comments on this git branch"""

144     p = git('log --pretty=tformat:%%B:SEP: %s..' % parent)

146     if not p:
147         sys.stderr.write("Failed getting git comments\n")
148         sys.exit(err)

150     return [x.strip() for x in p.readlines() if x != ':SEP:\n']

153 def git_file_list(parent, paths=None):
154     """Return the set of files which have ever changed on this branch.

156     NB: This includes files which no longer exist, or no longer actually
157     differ."""

159     p = git("log --name-only --pretty=format: %s.. %s" %
160            (parent, ' '.join(paths)))

162     if not p:
163         sys.stderr.write("Failed building file-list from git\n")
164         sys.exit(err)

166     ret = set()
167     for fname in p:
168         if fname and not fname.isspace() and fname not in ret:
169             ret.add(fname.strip())

171     return ret

174 def not_check(root, cmd):
175     """Return a function which returns True if a file given as an argument
176     should be excluded from the check named by 'cmd'"""

178     ignorefiles = filter(os.path.exists,
179                          [os.path.join(root, ".git", "%s.NOT" % cmd),
180                           os.path.join(root, "exception_lists", cmd)])
181     if len(ignorefiles) > 0:
182         return ignore.ignore(root, ignorefiles, sys.stderr.write)
183     else:
184         return lambda x: False

187 def gen_files(root, parent, paths, exclude):
188     """Return a function producing file names, relative to the current
189     directory, of any file changed on this branch (limited to 'paths' if
190     requested), and excluding files for which exclude returns a true value """

192     # Taken entirely from Python 2.6's os.path.relpath which we would use if we

```

```

193     # could.
194     def relpath(path, here):
195         c = os.path.abspath(os.path.join(root, path)).split(os.path.sep)
196         s = os.path.abspath(here).split(os.path.sep)
197         l = len(os.path.commonprefix((s, c)))
198         return os.path.join(*[os.path.pardir] * (len(s)-l) + c[l:])

200     def ret(select=None):
201         if not select:
202             select = lambda x: True

204         for f in git_file_list(parent, paths):
205             f = relpath(f, '.')
206             if (os.path.exists(f) and select(f) and not exclude(f)):
207                 yield f
208     return ret

211 def comchk(root, parent, flist, output):
212     output.write("Comments:\n")

214     return Comments.comchk(git_comments(parent), check_db=True,
215                            output=output)

218 def mapfilechk(root, parent, flist, output):
219     ret = 0

221     # We are interested in examining any file that has the following
222     # in its final path segment:
223     # - Contains the word 'mapfile'
224     # - Begins with 'map.'
225     # - Ends with '.map'
226     # We don't want to match unless these things occur in final path segment
227     # because directory names with these strings don't indicate a mapfile.
228     # We also ignore files with suffixes that tell us that the files
229     # are not mapfiles.
230     MapfileRE = re.compile(r'.*((mapfile[^\/*])|(/map\.[^\/*])|(\.map))$',
231                            re.IGNORECASE)
232     NotMapSuffixRE = re.compile(r'.*\.[ch]$', re.IGNORECASE)

234     output.write("Mapfile comments:\n")

236     for f in flist(lambda x: MapfileRE.match(x) and not
237                   NotMapSuffixRE.match(x)):
238         fh = open(f, 'r')
239         ret |= Mapfile.mapfilechk(fh, output=output)
240         fh.close()
241     return ret

244 def copyright(root, parent, flist, output):
245     ret = 0
246     output.write("Copyrights:\n")
247     for f in flist():
248         fh = open(f, 'r')
249         ret |= Copyright.copyright(fh, output=output)
250         fh.close()
251     return ret

254 def hdrchk(root, parent, flist, output):
255     ret = 0
256     output.write("Header format:\n")
257     for f in flist(lambda x: x.endswith('.h')):
258         fh = open(f, 'r')

```

```

259         ret |= HdrChk.hdrchk(fh, lenient=True, output=output)
260         fh.close()
261     return ret

264 def cstyle(root, parent, flist, output):
265     ret = 0
266     output.write("C style:\n")
267     for f in flist(lambda x: x.endswith('.c') or x.endswith('.h')):
268         fh = open(f, 'r')
269         ret |= CStyle.cstyle(fh, output=output, picky=True,
270                             check_posix_types=True,
271                             check_continuation=True)
272         fh.close()
273     return ret

276 def jstyle(root, parent, flist, output):
277     ret = 0
278     output.write("Java style:\n")
279     for f in flist(lambda x: x.endswith('.java')):
280         fh = open(f, 'r')
281         ret |= JStyle.jstyle(fh, output=output, picky=True)
282         fh.close()
283     return ret

286 def manlint(root, parent, flist, output):
287     ret = 0
288     output.write("Man page format:\n")
289     ManfileRE = re.compile(r'.*\.[0-9][a-z]*$', re.IGNORECASE)
290     for f in flist(lambda x: ManfileRE.match(x)):
291         fh = open(f, 'r')
292         ret |= ManLint.manlint(fh, output=output, picky=True)
293         fh.close()
294     return ret

296 def keywords(root, parent, flist, output):
297     ret = 0
298     output.write("SCCS Keywords:\n")
299     for f in flist():
300         fh = open(f, 'r')
301         ret |= Keywords.keywords(fh, output=output)
302         fh.close()
303     return ret

306 def run_checks(root, parent, cmds, paths='', opts={}):
307     """Run the checks given in 'cmds', expected to have well-known signatures,
308     and report results for any which fail.

310     Return failure if any of them did.

312     NB: the function name of the commands passed in is used to name the NOT
313     file which excepts files from them."""

315     ret = 0

317     for cmd in cmds:
318         s = StringIO()

320         exclude = not_check(root, cmd.func_name)
321         result = cmd(root, parent, gen_files(root, parent, paths, exclude),
322                    output=s)
323         ret |= result

```

```

325         if result != 0:
326             print s.getvalue()

328     return ret

331 def nits(root, parent, paths):
332     cmds = [copyright,
333            cstyle,
334            hdrchk,
335            jstyle,
336            keywords,
337            manlint,
338            mapfilechk]
339     run_checks(root, parent, cmds, paths)

342 def pbchk(root, parent, paths):
343     cmds = [comchk,
344            copyright,
345            cstyle,
346            hdrchk,
347            jstyle,
348            keywords,
349            manlint,
350            mapfilechk]
351     run_checks(root, parent, cmds)

354 def main(cmd, args):
355     parent_branch = None

357     try:
358         opts, args = getopt.getopt(args, 'b:')
359     except getopt.GetoptError, e:
360         sys.stderr.write(str(e) + '\n')
361         sys.stderr.write("Usage: %s [-b branch] [path...]\n" % cmd)
362         sys.exit(1)

364     for opt, arg in opts:
365         if opt == '-b':
366             parent_branch = arg

368     if not parent_branch:
369         parent_branch = git_parent_branch(git_branch())

371     func = nits
372     if cmd == 'git-pbchk':
373         func = pbchk
374     if args:
375         sys.stderr.write("only complete workspaces may be pbchk'd\n");
376         sys.exit(1)

378     func(git_root(), parent_branch, args)

380 if __name__ == '__main__':
381     try:
382         main(os.path.basename(sys.argv[0]), sys.argv[1:])
383     except GitError, e:
384         sys.stderr.write("failed to run git:\n %s\n" % str(e))
385         sys.exit(1)

```