



man pages section 3: Extended Library Functions

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A.

Part No: 806-0631-10
February 2000

Copyright 2000 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2000 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, Californie 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Contents

Preface 23

aclcheck(3SEC) 29

aclsort(3SEC) 31

actomode(3SEC) 32

aclfrommode(3SEC) 32

acltotext(3SEC) 33

aclfromtext(3SEC) 33

acos(3M) 35

acosh(3M) 36

asinh(3M) 36

atanh(3M) 36

asin(3M) 37

atan2(3M) 38

atan(3M) 39

au_open(3BSM) 40

au_close(3BSM) 40

au_write(3BSM) 40

au_preselect(3BSM) 41

au_to(3BSM) 43

au_to_arg(3BSM) 43
au_to_attr(3BSM) 43
au_to_data(3BSM) 43
au_to_groups(3BSM) 43
au_to_in_addr(3BSM) 43
au_to_ipc(3BSM) 43
au_to_ipc_perm(3BSM) 43
au_to_iport(3BSM) 43
au_to_me(3BSM) 43
au_to_new_in_addr(3BSM) 43
au_to_new_process(3BSM) 43
au_to_new_socket(3BSM) 43
au_to_new_subject(3BSM) 43
au_to_opaque(3BSM) 43
au_to_path(3BSM) 43
au_to_process(3BSM) 43
au_to_return(3BSM) 43
au_to_socket(3BSM) 43
au_to_subject(3BSM) 43
au_to_text(3BSM) 43
au_user_mask(3BSM) 47
bgets(3GEN) 49
bufsplit(3GEN) 50
cbrt(3M) 51
ceil(3M) 52
config_admin(3CFGADM) 53
config_change_state(3CFGADM) 53
config_private_func(3CFGADM) 53

| | |
|----------------------------|----|
| config_test(3CFGADM) | 53 |
| config_stat(3CFGADM) | 53 |
| config_list(3CFGADM) | 53 |
| config_list_ext(3CFGADM) | 53 |
| config_ap_id_cmp(3CFGADM) | 53 |
| config_unload(3CFGADM) | 53 |
| config_strerror(3CFGADM) | 53 |
| ConnectToServer(3DMI) | 62 |
| copylist(3GEN) | 63 |
| copysign(3M) | 64 |
| cos(3M) | 65 |
| cosh(3M) | 66 |
| cpc(3CPC) | 67 |
| cpc_access(3CPC) | 70 |
| cpc_bind_event(3CPC) | 71 |
| cpc_take_sample(3CPC) | 71 |
| cpc_rele(3CPC) | 71 |
| cpc_count_usr_events(3CPC) | 77 |
| cpc_count_sys_events(3CPC) | 77 |
| cpc_event(3CPC) | 79 |
| cpc_event_diff(3CPC) | 81 |
| cpc_event_accum(3CPC) | 81 |
| cpc_getcpuver(3CPC) | 83 |
| cpc_getcciname(3CPC) | 83 |
| cpc_getcpuref(3CPC) | 83 |
| cpc_getusage(3CPC) | 83 |
| cpc_getnpic(3CPC) | 83 |
| cpc_walk_names(3CPC) | 83 |

cpc_pctx_bind_event(3CPC) 85
cpc_pctx_take_sample(3CPC) 85
cpc_pctx_rele(3CPC) 85
cpc_pctx_invalidate(3CPC) 85
cpc_seterrfn(3CPC) 87
cpc_shared_open(3CPC) 88
cpc_shared_bind_event(3CPC) 88
cpc_shared_take_sample(3CPC) 88
cpc_shared_rele(3CPC) 88
cpc_shared_close(3CPC) 88
cpc_strtoevent(3CPC) 90
cpc_eventtostr(3CPC) 90
cpc_version(3CPC) 93
demangle(3EXT) 94
cplus_demangle(3EXT) 94
devid_get(3DEVID) 95
devid_free(3DEVID) 95
devid_get_minor_name(3DEVID) 95
devid_deviceid_to_nmlist(3DEVID) 95
devid_free_nmlist(3DEVID) 95
devid_compare(3DEVID) 95
devid_sizeof(3DEVID) 95
di_binding_name(3DEVINFO) 98
di_bus_addr(3DEVINFO) 98
di_compatible_names(3DEVINFO) 98
di_devid(3DEVINFO) 98
di_driver_name(3DEVINFO) 98
di_driver_ops(3DEVINFO) 98

| | |
|---------------------------------------|-----|
| di_instance(3DEVINFO) | 98 |
| di_nodeid(3DEVINFO) | 98 |
| di_node_name(3DEVINFO) | 98 |
| di_child_node(3DEVINFO) | 100 |
| di_parent_node(3DEVINFO) | 100 |
| di_sibling_node(3DEVINFO) | 100 |
| di_drv_first_node(3DEVINFO) | 100 |
| di_drv_next_node(3DEVINFO) | 100 |
| di_devfs_path(3DEVINFO) | 102 |
| di_devfs_path_free(3DEVINFO) | 102 |
| di_init(3DEVINFO) | 103 |
| di_fini(3DEVINFO) | 103 |
| di_minor_devt(3DEVINFO) | 106 |
| di_minor_name(3DEVINFO) | 106 |
| di_minor_nodetype(3DEVINFO) | 106 |
| di_minor_spectype(3DEVINFO) | 106 |
| di_minor_next(3DEVINFO) | 107 |
| di_prom_init(3DEVINFO) | 108 |
| di_prom_fini(3DEVINFO) | 108 |
| di_prom_prop_data(3DEVINFO) | 109 |
| di_prom_prop_next(3DEVINFO) | 109 |
| di_prom_prop_name(3DEVINFO) | 109 |
| di_prom_prop_lookup_bytes(3DEVINFO) | 111 |
| di_prom_prop_lookup_ints(3DEVINFO) | 111 |
| di_prom_prop_lookup_strings(3DEVINFO) | 111 |
| di_prop_bytes(3DEVINFO) | 113 |
| di_prop_devt(3DEVINFO) | 113 |
| di_prop_ints(3DEVINFO) | 113 |

di_prop_name(3DEVINFO) 113
di_prop_strings(3DEVINFO) 113
di_prop_type(3DEVINFO) 113
di_prop_lookup_bytes(3DEVINFO) 116
di_prop_lookup_ints(3DEVINFO) 116
di_prop_lookup_strings(3DEVINFO) 116
di_prop_next(3DEVINFO) 118
DisconnectToServer(3DMI) 119
di_walk_minor(3DEVINFO) 120
di_walk_node(3DEVINFO) 122
DmiAddComponent(3DMI) 124
DmiAddGroup(3DMI) 124
DmiAddLanguage(3DMI) 124
DmiDeleteComponent(3DMI) 124
DmiDeleteGroup(3DMI) 124
DmiDeleteLanguage(3DMI) 124
DmiAddRow(3DMI) 129
DmiDeleteRow(3DMI) 129
DmiGetAttribute(3DMI) 129
DmiGetMultiple(3DMI) 129
DmiSetAttribute(3DMI) 129
DmiSetMultiple(3DMI) 129
dmi_error(3DMI) 134
DmiGetConfig(3DMI) 135
DmiGetVersion(3DMI) 135
DmiRegister(3DMI) 135
DmiSetConfig(3DMI) 135
DmiUnregister(3DMI) 135

| | |
|--------------------------------|-----|
| DmiListAttributes(3DMI) | 139 |
| DmiListClassNames(3DMI) | 139 |
| DmiListComponents(3DMI) | 139 |
| DmiListComponentsByClass(3DMI) | 139 |
| DmiListGroups(3DMI) | 139 |
| DmiListLanguages(3DMI) | 139 |
| DmiRegisterCi(3DMI) | 145 |
| DmiUnRegisterCi(3DMI) | 145 |
| DmiOriginateEvent(3DMI) | 145 |
| elf32_checksum(3ELF) | 147 |
| elf64_checksum(3ELF) | 147 |
| elf32_fsize(3ELF) | 148 |
| elf64_fsize(3ELF) | 148 |
| elf32_getehdr(3ELF) | 149 |
| elf32_newehdr(3ELF) | 149 |
| elf64_getehdr(3ELF) | 149 |
| elf64_newehdr(3ELF) | 149 |
| elf32_getphdr(3ELF) | 151 |
| elf32_newphdr(3ELF) | 151 |
| elf64_getphdr(3ELF) | 151 |
| elf64_newphdr(3ELF) | 151 |
| elf32_getshdr(3ELF) | 153 |
| elf64_getshdr(3ELF) | 153 |
| elf32_xlatetof(3ELF) | 155 |
| elf32_xlatetom(3ELF) | 155 |
| elf64_xlatetof(3ELF) | 155 |
| elf64_xlatetom(3ELF) | 155 |
| elf(3ELF) | 157 |

elf_begin(3ELF) 163
elf_end(3ELF) 163
elf_memory(3ELF) 163
elf_next(3ELF) 163
elf_rand(3ELF) 163
elf_cntl(3ELF) 168
elf_errmsg(3ELF) 170
elf_errno(3ELF) 170
elf_fill(3ELF) 171
elf_flagdata(3ELF) 172
elf_flagehdr(3ELF) 172
elf_flagelf(3ELF) 172
elf_flagphdr(3ELF) 172
elf_flagscn(3ELF) 172
elf_flagshdr(3ELF) 172
elf_getarhdr(3ELF) 174
elf_getarsym(3ELF) 176
elf_getbase(3ELF) 177
elf_getdata(3ELF) 178
elf_newdata(3ELF) 178
elf_rawdata(3ELF) 178
elf_getident(3ELF) 182
elf_getscn(3ELF) 184
elf_ndxscn(3ELF) 184
elf_newscn(3ELF) 184
elf_nextscn(3ELF) 184
elf_hash(3ELF) 186
elf_kind(3ELF) 187

elf_rawfile(3ELF) 188
elf_strptr(3ELF) 189
elf_update(3ELF) 190
elf_version(3ELF) 194
erf(3M) 195
erfc(3M) 195
exp(3M) 196
expm1(3M) 197
fabs(3M) 198
floor(3M) 199
fmod(3M) 200
freeDmiString(3DMI) 201
gelf(3ELF) 202
gelf_checksum(3ELF) 202
gelf_fsize(3ELF) 202
gelf_getclass(3ELF) 202
gelf_getdyn(3ELF) 202
gelf_getehdr(3ELF) 202
gelf_getphdr(3ELF) 202
gelf_getrel(3ELF) 202
gelf_getrela(3ELF) 202
gelf_getshdr(3ELF) 202
gelf_getsym(3ELF) 202
gelf_getsyminfo(3ELF) 202
gelf_newehdr(3ELF) 202
gelf_newphdr(3ELF) 202
gelf_update_dyn(3ELF) 202
gelf_update_ehdr(3ELF) 202

gelf_update_phdr(3ELF) 202
gelf_update_rel(3ELF) 202
gelf_update_rela(3ELF) 202
gelf_update_shdr(3ELF) 202
gelf_update_sym(3ELF) 202
gelf_update_syminfo(3ELF) 202
gelf_xlatetof(3ELF) 202
gelf_xslatetom(3ELF) 202
getacinfo(3BSM) 207
getacdir(3BSM) 207
getacflg(3BSM) 207
getacmin(3BSM) 207
getacna(3BSM) 207
setac(3BSM) 207
endac(3BSM) 207
getauclassent(3BSM) 209
getauclassnam(3BSM) 209
setauclass(3BSM) 209
endauclass(3BSM) 209
getauclassnam_r(3BSM) 209
getauclassent_r(3BSM) 209
getauditflags(3BSM) 211
getauditflagsbin(3BSM) 211
getauditflagschar(3BSM) 211
getauevent(3BSM) 212
getauevnam(3BSM) 212
getauevnum(3BSM) 212
getauevnonam(3BSM) 212

| | |
|------------------------|-----|
| setauevent(3BSM) | 212 |
| endauevent(3BSM) | 212 |
| getauevent_r(3BSM) | 212 |
| getauevnam_r(3BSM) | 212 |
| getauevnum_r(3BSM) | 212 |
| getauthattr(3SECDB) | 215 |
| getauthnam(3SECDB) | 215 |
| free_authattr(3SECDB) | 215 |
| setauthattr(3SECDB) | 215 |
| endauthattr(3SECDB) | 215 |
| chkauthattr(3SECDB) | 215 |
| getausernam(3BSM) | 218 |
| getauserent(3BSM) | 218 |
| setauser(3BSM) | 218 |
| endauser(3BSM) | 218 |
| getexecattr(3SECDB) | 220 |
| free_execattr(3SECDB) | 220 |
| setexecattr(3SECDB) | 220 |
| endexecattr(3SECDB) | 220 |
| getexecuser(3SECDB) | 220 |
| getexecprof(3SECDB) | 220 |
| match_execattr(3SECDB) | 220 |
| getfauditflags(3BSM) | 223 |
| getprofattr(3SECDB) | 224 |
| getprofnam(3SECDB) | 224 |
| free_profattr(3SECDB) | 224 |
| setprofattr(3SECDB) | 224 |
| endprofattr(3SECDB) | 224 |

getuserattr(3SECDB) 226
getusernam(3SECDB) 226
getuseruid(3SECDB) 226
free_userattr(3SECDB) 226
setuserattr(3SECDB) 226
enduserattr(3SECDB) 226
gmatch(3GEN) 228
hypot(3M) 229
ilogb(3M) 230
isencrypt(3GEN) 231
isnan(3M) 232
j0(3M) 233
j1(3M) 233
jn(3M) 233
kstat(3EXT) 234
kstat(3KSTAT) 237
kstat_chain_update(3KSTAT) 243
kstat_lookup(3KSTAT) 244
kstat_data_lookup(3KSTAT) 244
kstat_open(3KSTAT) 245
kstat_close(3KSTAT) 245
kstat_read(3KSTAT) 246
kstat_write(3KSTAT) 246
kva_match(3SECDB) 247
kvm_getu(3KVM) 248
kvm_getcmd(3KVM) 248
kvm_nextproc(3KVM) 250
kvm_getproc(3KVM) 250

| | |
|--------------------------|-----|
| kvm_setproc(3KVM) | 250 |
| kvm_nlist(3KVM) | 252 |
| kvm_open(3KVM) | 253 |
| kvm_close(3KVM) | 253 |
| kvm_read(3KVM) | 255 |
| kvm_write(3KVM) | 255 |
| kvm_uread(3KVM) | 255 |
| kvm_uwrite(3KVM) | 255 |
| kvm_kread(3KVM) | 255 |
| kvm_kwrite(3KVM) | 255 |
| lgamma(3M) | 257 |
| lgamma_r(3M) | 257 |
| gamma(3M) | 257 |
| gamma_r(3M) | 257 |
| libdevinfo(3DEVINFO) | 259 |
| libtnfctl(3TNF) | 262 |
| log10(3M) | 267 |
| log1p(3M) | 268 |
| log(3M) | 269 |
| logb(3M) | 270 |
| maillock(3MAIL) | 271 |
| mailunlock(3MAIL) | 271 |
| touchlock(3MAIL) | 271 |
| matherr(3M) | 273 |
| m_create_layout(3LAYOUT) | 279 |
| md5(3EXT) | 281 |
| MD5Init(3EXT) | 281 |
| MD5Update(3EXT) | 281 |

MD5Final(3EXT) 281
md5_calc(3EXT) 281
m_destroy_layout(3LAYOUT) 283
media_findname(3VOLMGT) 284
media_getattr(3VOLMGT) 287
media_setattr(3VOLMGT) 287
media_getid(3VOLMGT) 290
m_getvalues_layout(3LAYOUT) 291
mkdirp(3GEN) 292
rmdirp(3GEN) 292
mp(3MP) 293
mp_madd(3MP) 293
mp_msub(3MP) 293
mp_mult(3MP) 293
mp_mdiv(3MP) 293
mp_mcmp(3MP) 293
mp_min(3MP) 293
mp_mout(3MP) 293
mp_pow(3MP) 293
mp_gcd(3MP) 293
mp_rpow(3MP) 293
mp_itom(3MP) 293
mp_xtom(3MP) 293
mp_mtox(3MP) 293
mp_mfree(3MP) 293
m_setvalues_layout(3LAYOUT) 295
m_transform_layout(3LAYOUT) 296
m_wtransform_layout(3LAYOUT) 301

| | |
|---------------------------|-----|
| newDmiOctetString(3DMI) | 307 |
| newDmiString(3DMI) | 308 |
| nextafter(3M) | 309 |
| nlist(3ELF) | 310 |
| NOTE(3EXT) | 311 |
| _NOTE(3EXT) | 311 |
| p2open(3GEN) | 313 |
| p2close(3GEN) | 313 |
| pam(3PAM) | 315 |
| pam_acct_mgmt(3PAM) | 318 |
| pam_authenticate(3PAM) | 320 |
| pam_chauthtok(3PAM) | 322 |
| pam_getenv(3PAM) | 324 |
| pam_getenvlist(3PAM) | 325 |
| pam_get_user(3PAM) | 326 |
| pam_open_session(3PAM) | 328 |
| pam_close_session(3PAM) | 328 |
| pam_putenv(3PAM) | 330 |
| pam_setcred(3PAM) | 332 |
| pam_set_data(3PAM) | 334 |
| pam_get_data(3PAM) | 334 |
| pam_set_item(3PAM) | 336 |
| pam_get_item(3PAM) | 336 |
| pam_sm(3PAM) | 338 |
| pam_sm_acct_mgmt(3PAM) | 342 |
| pam_sm_authenticate(3PAM) | 344 |
| pam_sm_chauthtok(3PAM) | 346 |
| pam_sm_open_session(3PAM) | 349 |

pam_sm_close_session(3PAM) 349
pam_sm_setcred(3PAM) 351
pam_start(3PAM) 353
pam_end(3PAM) 353
pam_strerror(3PAM) 356
pathfind(3GEN) 357
pctx_capture(3CPC) 359
pctx_create(3CPC) 359
pctx_run(3CPC) 359
pctx_release(3CPC) 359
pctx_set_events(3CPC) 361
pow(3M) 364
printDmiAttributeValues(3DMI) 365
printDmiDataUnion(3DMI) 366
printDmiString(3DMI) 367
read_vtoc(3EXT) 368
write_vtoc(3EXT) 368
reg_ci_callback(3DMI) 369
regexpr(3GEN) 370
compile(3GEN) 370
step(3GEN) 370
advance(3GEN) 370
remainder(3M) 373
rint(3M) 374
scalb(3M) 375
scalbn(3M) 376
significand(3M) 377
sin(3M) 378

| | |
|------------------------|-----|
| sinh(3M) | 379 |
| sqrt(3M) | 380 |
| SSAAgentIsAlive(3SNMP) | 381 |
| SSAGetTrapPort(3SNMP) | 381 |
| SSARegSubtable(3SNMP) | 381 |
| SSARegSubagent(3SNMP) | 381 |
| SSARegSubtree(3SNMP) | 381 |
| SSASendTrap(3SNMP) | 381 |
| SSASubagentOpen(3SNMP) | 381 |
| SSAOidCmp(3SNMP) | 384 |
| SSAOidCpy(3SNMP) | 384 |
| SSAOidDup(3SNMP) | 384 |
| SSAOidFree(3SNMP) | 384 |
| SSAOidInit(3SNMP) | 384 |
| SSAOidNew(3SNMP) | 384 |
| SSAOidString(3SNMP) | 384 |
| SSAOidStrToOid(3SNMP) | 384 |
| SSAOidZero(3SNMP) | 384 |
| SSAStringCpy(3SNMP) | 386 |
| SSAStringInit(3SNMP) | 386 |
| SSAStringToChar(3SNMP) | 386 |
| SSAStringZero(3SNMP) | 386 |
| strncpy(3GEN) | 387 |
| stredd(3GEN) | 387 |
| strcadd(3GEN) | 387 |
| strecpy(3GEN) | 387 |
| strfind(3GEN) | 389 |
| strrspn(3GEN) | 389 |

strtrns(3GEN) 389
str(3GEN) 389
tan(3M) 390
tanh(3M) 391
tnfctl_buffer_alloc(3TNF) 392
tnfctl_buffer_dealloc(3TNF) 392
tnfctl_close(3TNF) 394
tnfctl_indirect_open(3TNF) 396
tnfctl_check_libs(3TNF) 396
tnfctl_internal_open(3TNF) 399
tnfctl_kernel_open(3TNF) 401
tnfctl_pid_open(3TNF) 402
tnfctl_exec_open(3TNF) 402
tnfctl_continue(3TNF) 402
tnfctl_probe_apply(3TNF) 408
tnfctl_probe_apply_ids(3TNF) 408
tnfctl_probe_state_get(3TNF) 411
tnfctl_probe_enable(3TNF) 411
tnfctl_probe_disable(3TNF) 411
tnfctl_probe_trace(3TNF) 411
tnfctl_probe_untrace(3TNF) 411
tnfctl_probe_connect(3TNF) 411
tnfctl_probe_disconnect_all(3TNF) 411
tnfctl_register_funcs(3TNF) 416
tnfctl_strerror(3TNF) 417
tnfctl_trace_attrs_get(3TNF) 418
tnfctl_trace_state_set(3TNF) 421
tnfctl_filter_state_set(3TNF) 421

| | |
|---------------------------------|-----|
| tnfctl_filter_list_get(3TNF) | 421 |
| tnfctl_filter_list_add(3TNF) | 421 |
| tnfctl_filter_list_delete(3TNF) | 421 |
| TNF_DECLARE_RECORD(3TNF) | 423 |
| TNF_DEFINE_RECORD_1(3TNF) | 423 |
| TNF_DEFINE_RECORD_2(3TNF) | 423 |
| TNF_DEFINE_RECORD_3(3TNF) | 423 |
| TNF_DEFINE_RECORD_4(3TNF) | 423 |
| TNF_DEFINE_RECORD_5(3TNF) | 423 |
| TNF_PROBE(3TNF) | 426 |
| TNF_PROBE_0(3TNF) | 426 |
| TNF_PROBE_1(3TNF) | 426 |
| TNF_PROBE_2(3TNF) | 426 |
| TNF_PROBE_3(3TNF) | 426 |
| TNF_PROBE_4(3TNF) | 426 |
| TNF_PROBE_5(3TNF) | 426 |
| TNF_PROBE_0_DEBUG(3TNF) | 426 |
| TNF_PROBE_1_DEBUG(3TNF) | 426 |
| TNF_PROBE_2_DEBUG(3TNF) | 426 |
| TNF_PROBE_3_DEBUG(3TNF) | 426 |
| TNF_PROBE_4_DEBUG(3TNF) | 426 |
| TNF_PROBE_5_DEBUG(3TNF) | 426 |
| TNF_DEBUG(3TNF) | 426 |
| tnf_process_disable(3TNF) | 431 |
| tnf_process_enable(3TNF) | 431 |
| tnf_thread_disable(3TNF) | 431 |
| tnf_thread_enable(3TNF) | 431 |
| tracing(3TNF) | 433 |

volmgt_acquire(3VOLMGT) 437
volmgt_check(3VOLMGT) 440
volmgt_feature_enabled(3VOLMGT) 442
volmgt_inuse(3VOLMGT) 443
volmgt_ownspath(3VOLMGT) 444
volmgt_release(3VOLMGT) 445
volmgt_root(3VOLMGT) 447
volmgt_running(3VOLMGT) 448
volmgt_symname(3VOLMGT) 449
volmgt_symdev(3VOLMGT) 449
y0(3M) 451
y1(3M) 451
yn(3M) 451
Index 451

Preface

Both novice users and those familiar with the SunOS operating system can use online man pages to obtain information about the system and its features. A man page is intended to answer concisely the question “What does it do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

Overview

The following contains a brief description of each man page section and the information it references:

- Section 1 describes, in alphabetical order, commands available with the operating system.
- Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.
- Section 2 describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value.
- Section 3 describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2.
- Section 4 outlines the formats of various files. The C structure declarations for the file formats are given where applicable.
- Section 5 contains miscellaneous documentation such as character-set tables.
- Section 6 contains available games and demos.
- Section 7 describes various special files that refer to specific hardware peripherals and device drivers. STREAMS software drivers, modules and the STREAMS-generic set of system calls are also described.

- Section 9 provides reference information needed to write device drivers in the kernel environment. It describes two device driver interface specifications: the Device Driver Interface (DDI) and the Driver/Kernel Interface (DKI).
- Section 9E describes the DDI/DKI, DDI-only, and DKI-only entry-point routines a developer can include in a device driver.
- Section 9F describes the kernel functions available for use by device drivers.
- Section 9S describes the data structures used by drivers to share information between the driver and the kernel.

Below is a generic format for man pages. The man pages of each manual section generally follow this order, but include only needed headings. For example, if there are no bugs to report, there is no BUGS section. See the `intro` pages for more information and detail about each section, and `man(1)` for more information about man pages in general.

| | |
|----------|--|
| NAME | This section gives the names of the commands or functions documented, followed by a brief description of what they do. |
| SYNOPSIS | <p>This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full path name is shown. Options and arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.</p> <p>The following special characters are used in this section:</p> <ul style="list-style-type: none"> [] Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified. . . . Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename . . .". Separator. Only one of the arguments separated by this character can be specified at a time. { } Braces. The options and/or arguments enclosed within braces are |

| | |
|---------------|--|
| | interdependent, such that everything enclosed must be treated as a unit. |
| PROTOCOL | This section occurs only in subsection 3R to indicate the protocol description file. |
| DESCRIPTION | This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss OPTIONS or cite EXAMPLES. Interactive commands, subcommands, requests, macros, and functions are described under USAGE. |
| IOCTL | This section appears on pages in Section 7 only. Only the device class that supplies appropriate parameters to the <code>ioctl(2)</code> system call is called <code>ioctl</code> and generates its own heading. <code>ioctl</code> calls for a specific device are listed alphabetically (on the man page for that specific device). <code>ioctl</code> calls are used for a particular class of devices all of which have an <code>io</code> ending, such as <code>mtio(7I)</code> . |
| OPTIONS | This section lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied. |
| OPERANDS | This section lists the command operands and describes how they affect the actions of the command. |
| OUTPUT | This section describes the output – standard output, standard error, or output files – generated by the command. |
| RETURN VALUES | If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared void do not return values, so they are not discussed in RETURN VALUES. |
| ERRORS | On failure, most functions place an error code in the global variable <code>errno</code> indicating why they |

failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than one condition can cause the same error, each condition is described in a separate paragraph under the error code.

USAGE

This section lists special rules, features, and commands that require in-depth explanations. The subsections listed here are used to explain built-in functionality:

- Commands
- Modifiers
- Variables
- Expressions
- Input Grammar

EXAMPLES

This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command-line entry and machine response is shown. Whenever an example is given, the prompt is shown as `example%`, or if the user must be superuser, `example#`. Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS, and USAGE sections.

ENVIRONMENT VARIABLES

This section lists any environment variables that the command or function affects, followed by a brief description of the effect.

EXIT STATUS

This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion, and values other than zero for various error conditions.

FILES

This section lists all file names referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.

ATTRIBUTES

This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See `attributes(5)` for more information.

| | |
|-------------|---|
| SEE ALSO | This section lists references to other man pages, in-house documentation, and outside publications. |
| DIAGNOSTICS | This section lists diagnostic messages with a brief explanation of the condition causing the error. |
| WARNINGS | This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics. |
| NOTES | This section lists additional information that does not belong anywhere else on the page. It takes the form of an aside to the user, covering points of special interest. Critical information is never covered here. |
| BUGS | This section describes known bugs and, wherever possible, suggests workarounds. |

Introduction to Library Functions

| | | | | | |
|-------------------------|--|------------------------|--|-------------------------|--|
| NAME | aclcheck – check the validity of an ACL | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lsec [library ...] #include <sys/acl.h> int aclcheck(aclent_t *aclbufp, int nentries, int *which);</pre> | | | | |
| DESCRIPTION | <p>The <code>aclcheck()</code> function checks the validity of an ACL pointed to by <code>aclbufp</code>. The <code>nentries</code> argument is the number of entries contained in the buffer. The <code>which</code> parameter returns the index of the first entry that is invalid.</p> <p>The function verifies that an ACL pointed to by <code>aclbufp</code> is valid according to the following rules:</p> <ul style="list-style-type: none"> ■ There must be exactly one <code>GROUP_OBJ</code> ACL entry. ■ There must be exactly one <code>USER_OBJ</code> ACL entry. ■ There must be exactly one <code>OTHER_OBJ</code> ACL entry. ■ If there are any <code>GROUP</code> ACL entries, then the group ID in each group ACL entry must be unique. ■ If there are any <code>USER</code> ACL entries, then the user ID in each user ACL entry must be unique. ■ If there are any <code>GROUP</code> or <code>USER</code> ACL entries, then there must be exactly one <code>CLASS_OBJ</code> (ACL mask) entry. ■ If there are any default ACL entries, then the following apply: <ul style="list-style-type: none"> ■ There must be exactly one default <code>GROUP_OBJ</code> ACL entry. ■ There must be exactly one default <code>OTHER_OBJ</code> ACL entry. ■ There must be exactly one default <code>USER_OBJ</code> ACL entry. ■ If there are any <code>DEF_GROUP</code> entries, then the group ID in each <code>DEF_GROUP</code> ACL entry must be unique. ■ If there are any <code>DEF_USER</code> entries, then the user ID in each <code>DEF_USER</code> ACL entry must be unique. ■ If there are any <code>DEF_GROUP</code> or <code>DEF_USER</code> entries, then there must be exactly one <code>DEF_CLASS_OBJ</code> (default ACL mask) entry. ■ If any of the above rules are violated, then the function fails with <code>errno</code> set to <code>EINVAL</code>. | | | | |
| RETURN VALUES | <p>If the ACL is valid, <code>aclcheck()</code> will return 0. Otherwise <code>errno</code> is set to <code>EINVAL</code> and return code is set to one of the following:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;"><code>GRP_ERROR</code></td> <td>There is more than one <code>GROUP_OBJ</code> or <code>DEF_GROUP_OBJ</code> ACL entry.</td> </tr> <tr> <td><code>USER_ERROR</code></td> <td>There is more than one <code>USER_OBJ</code> or <code>DEF_USER_OBJ</code> ACL entry.</td> </tr> </table> | <code>GRP_ERROR</code> | There is more than one <code>GROUP_OBJ</code> or <code>DEF_GROUP_OBJ</code> ACL entry. | <code>USER_ERROR</code> | There is more than one <code>USER_OBJ</code> or <code>DEF_USER_OBJ</code> ACL entry. |
| <code>GRP_ERROR</code> | There is more than one <code>GROUP_OBJ</code> or <code>DEF_GROUP_OBJ</code> ACL entry. | | | | |
| <code>USER_ERROR</code> | There is more than one <code>USER_OBJ</code> or <code>DEF_USER_OBJ</code> ACL entry. | | | | |

CLASS_ERROR There is more than one CLASS_OBJ (ACL mask) or DEF_CLASS_OBJ (default ACL mask) entry.

OTHER_ERROR There is more than one OTHER_OBJ or DEF_OTHER_OBJ ACL entry.

DUPLICATE_ERROR Duplicate entries of USER, GROUP, DEF_USER, or DEF_GROUP.

ENTRY_ERROR The entry type is invalid.

MISS_ERROR Missing an entry. The *which* parameter returns -1 in this case.

MEM_ERROR The system cannot allocate any memory. The *which* parameter returns -1 in this case.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | Evolving |

SEE ALSO

`acl(2)`, `aclsort(3SEC)`

| NAME | aclsort – sort an ACL | | | | |
|----------------------|--|----------------|-----------------|---------------------|----------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lsec [library ...] #include <sys/acl.h> int aclsort(int nentries, int calclass, aclent_t *aclbufp);</pre> | | | | |
| DESCRIPTION | <p>The <i>aclbufp</i> argument points to a buffer containing ACL entries. The <i>nentries</i> argument specifies the number of ACL entries in the buffer. The <i>calclass</i> argument, if non-zero, indicates that the CLASS_OBJ (ACL mask) permissions should be recalculated. The union of the permission bits associated with all ACL entries in the buffer other than CLASS_OBJ, OTHER_OBJ, and USER_OBJ is calculated. The result is copied to the permission bits associated with the CLASS_OBJ entry.</p> <p>The <code>aclsort()</code> function sorts the contents of the ACL buffer as follows:</p> <ul style="list-style-type: none"> ■ Entries will be in the order USER_OBJ, USER, GROUP_OBJ, GROUP, CLASS_OBJ (ACL mask), OTHER_OBJ, DEF_USER_OBJ, DEF_USER, DEF_GROUP_OBJ, DEF_GROUP, DEF_CLASS_OBJ (default ACL mask), and DEF_OTHER_OBJ. ■ Entries of type USER, GROUP, DEF_USER, and DEF_GROUP will be sorted in increasing order by ID. <p>The <code>aclsort()</code> function will succeed if all of the following are true:</p> <ul style="list-style-type: none"> ■ There is exactly one entry each of type USER_OBJ, GROUP_OBJ, CLASS_OBJ (ACL mask), and OTHER_OBJ. ■ There is exactly one entry each of type DEF_USER_OBJ, DEF_GROUP_OBJ, DEF_CLASS_OBJ (default ACL mask), and DEF_OTHER_OBJ if there are any default entries. ■ Entries of type USER, GROUP, DEF_USER, or DEF_GROUP may not contain duplicate entries. A duplicate entry is one of the same type containing the same numeric ID. | | | | |
| RETURN VALUES | Upon successful completion, the the function returns 0. Otherwise, it returns -1. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | Interface Stability | Evolving |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| Interface Stability | Evolving | | | | |
| SEE ALSO | <code>acl(2)</code> , <code>aclcheck(3SEC)</code> | | | | |

NAME | acltomode, aclfrommode – convert an ACL to or from permission bits

SYNOPSIS | `cc [flag ...] file ... -lsec [library ...]`
 | `#include <sys/types.h>`
 | `#include <sys/acl.h>`
 | `int acltomode(aclent_t *aclbufp, int nentries, mode_t *modep);`
 | `int aclfrommode(aclent_t *aclbufp, int nentries, mode_t *modep);`

DESCRIPTION | The `acltomode()` function converts an ACL pointed to by `aclbufp` into the permission bits buffer pointed to by `modep`. If the `USER_OBJ` ACL entry, `GROUP_OBJ` ACL entry, or the `OTHER_OBJ` ACL entry cannot be found in the ACL buffer, then the function fails with `errno` set to `EINVAL`.

| The `USER_OBJ` ACL entry permission bits are copied to the file owner permission bits in the permission bits buffer. The `OTHER_OBJ` ACL entry permission bits are copied to the file other permission bits in the permission bits buffer. If there is a `CLASS_OBJ` (ACL mask) entry, then the `CLASS_OBJ` ACL entry permission bits are intersected (bitwise AND) with the `GROUP_OBJ` ACL entry permission bits and the result is copied to the file group permission bits in the permission bits buffer. Otherwise, the `GROUP_OWNER` ACL entry permission bits are copied to the file group permission bits in the permission bits buffer.

| The `aclfrommode()` function converts the permission bits pointed to by `modep` into an ACL pointed to by `aclbufp`. If the `USER_OBJ` ACL entry, `GROUP_OBJ` ACL entry, or the `OTHER_OBJ` ACL entry cannot be found in the ACL buffer, then the function fails with `errno` set to `EINVAL`.

| The file owner permission bits from the permission bits buffer are copied to the `USER_OBJ` ACL entry. The file other permission bits from the permission bits buffer are copied to the `OTHER_OBJ` ACL entry. The file group permissions bits from the permission bits buffer are copied to the `CLASS_OBJ` (ACL mask) entry, if available, and to the `GROUP_OBJ` ACL entry.

| The `nentries` argument represents the number of ACL entries in the buffer pointed to by `aclbufp`.

RETURN VALUES | Upon successful completion, the function returns 0. Otherwise, it returns -1 and sets `errno` to indicate the error.

ATTRIBUTES | See `attributes (5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | Evolving |

SEE ALSO | `acl(2)`

| | | | | | | | | | | | |
|---------------------------|--|-------------------|---|--------------------|--|--------------------|---|-------------------|---|---------------------------|---|
| NAME | acltotext, aclfromtext – convert internal representation to or from external representation | | | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lsec [library ...] #include <sys/acl.h> char *acltotext(aclent_t *aclbufp, int aclcnt); aclent_t *aclfromtext(char *acltextp, int *aclcnt);</pre> | | | | | | | | | | |
| DESCRIPTION | <p>The <code>acltotext()</code> function converts an internal ACL representation pointed to by <code>aclbufp</code> into an external ACL representation. The space for the external text string is obtained using <code>malloc(3C)</code>. The caller is responsible for freeing the space upon completion..</p> <p>The <code>aclfromtext()</code> function converts an external ACL representation pointed to by <code>acltextp</code> into an internal ACL representation. The space for the list of ACL entries is obtained using <code>malloc(3C)</code>. The caller is responsible for freeing the space upon completion. The <code>aclcnt</code> argument indicates the number of ACL entries found.</p> <p>An external ACL representation is defined as follows:</p> <pre><acl_entry>[,<acl_entry>]...</pre> <p>Each <code><acl_entry></code> contains one ACL entry. The external representation of an ACL entry contains two or three colon-separated fields. The first field contains the ACL entry tag type. The entry type keywords are defined as:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>user</code></td> <td>This ACL entry with no UID specified in the ACL entry ID field specifies the access granted to the owner of the object. Otherwise, this ACL entry specifies the access granted to a specific user-name or user-id number.</td> </tr> <tr> <td><code>group</code></td> <td>This ACL entry with no GID specified in the ACL entry ID field specifies the access granted to the owning group of the object. Otherwise, this ACL entry specifies the access granted to a specific group-name or group-id number.</td> </tr> <tr> <td><code>other</code></td> <td>This ACL entry specifies the access granted to any user or group that does not match any other ACL entry.</td> </tr> <tr> <td><code>mask</code></td> <td>This ACL entry specifies the maximum access granted to user or group entries.</td> </tr> <tr> <td><code>default:user</code></td> <td>This ACL entry with no uid specified in the ACL entry ID field specifies the default access granted to the owner of the object. Otherwise, this ACL entry specifies the default access granted to a specific user-name or user-ID number.</td> </tr> </table> | <code>user</code> | This ACL entry with no UID specified in the ACL entry ID field specifies the access granted to the owner of the object. Otherwise, this ACL entry specifies the access granted to a specific user-name or user-id number. | <code>group</code> | This ACL entry with no GID specified in the ACL entry ID field specifies the access granted to the owning group of the object. Otherwise, this ACL entry specifies the access granted to a specific group-name or group-id number. | <code>other</code> | This ACL entry specifies the access granted to any user or group that does not match any other ACL entry. | <code>mask</code> | This ACL entry specifies the maximum access granted to user or group entries. | <code>default:user</code> | This ACL entry with no uid specified in the ACL entry ID field specifies the default access granted to the owner of the object. Otherwise, this ACL entry specifies the default access granted to a specific user-name or user-ID number. |
| <code>user</code> | This ACL entry with no UID specified in the ACL entry ID field specifies the access granted to the owner of the object. Otherwise, this ACL entry specifies the access granted to a specific user-name or user-id number. | | | | | | | | | | |
| <code>group</code> | This ACL entry with no GID specified in the ACL entry ID field specifies the access granted to the owning group of the object. Otherwise, this ACL entry specifies the access granted to a specific group-name or group-id number. | | | | | | | | | | |
| <code>other</code> | This ACL entry specifies the access granted to any user or group that does not match any other ACL entry. | | | | | | | | | | |
| <code>mask</code> | This ACL entry specifies the maximum access granted to user or group entries. | | | | | | | | | | |
| <code>default:user</code> | This ACL entry with no uid specified in the ACL entry ID field specifies the default access granted to the owner of the object. Otherwise, this ACL entry specifies the default access granted to a specific user-name or user-ID number. | | | | | | | | | | |

`default:group` This ACL entry with no gid specified in the ACL entry ID field specifies the default access granted to the owning group of the object. Otherwise, this ACL entry specifies the default access granted to a specific group-name or group-ID number.

`default:other` This ACL entry specifies the default access for other entry.

`default:mask` This ACL entry specifies the default access for mask entry.

The second field contains the ACL entry ID , as follows:

`uid` This field specifies a user-name, or user-ID if there is no user-name associated with the user-ID number.

`gid` This field specifies a group-name, or group-ID if there is no group-name associated with the group-ID number.

`empty` This field is used by the user and group ACL entry types.

The third field contains the following symbolic discretionary access permissions:

`r` read permission

`w` write permission

`x` execute/search permission

`-` no access

RETURN VALUES

Upon successful completion, the `acltotext()` function returns a pointer to a text string. Otherwise, it returns `NULL`.

Upon successful completion, the `aclfromtext()` function returns a pointer to a list of ACL entries. Otherwise, it returns `NULL`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | Evolving |

SEE ALSO

`acl(2)` , `malloc(3C)`

| NAME | acos – arc cosine function | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> | | | | |
| DESCRIPTION | The <code>acos()</code> function computes the principal value of the arc cosine of x . The value of x should be in the range $[-1,1]$. | | | | |
| RETURN VALUES | Upon successful completion, <code>acos()</code> returns the arc cosine of x , in the range $[0,\pi]$ radians. If the value of x is not in the range $[-1,1]$, and is not $\pm\text{Inf}$ or NaN, either 0.0 or NaN is returned and <code>errno</code> is set to <code>EDOM</code> . If x is NaN, NaN is returned. If x is $\pm\text{Inf}$, either 0.0 is returned and <code>errno</code> is set to <code>EDOM</code> , or NaN is returned and <code>errno</code> may be set to <code>EDOM</code> . For exceptional cases, <code>matherr(3M)</code> tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The <code>acos()</code> function will fail if: <code>EDOM</code> The value x is not $\pm\text{Inf}$ or NaN and is not in the range $[-1,1]$. The <code>acos()</code> function may fail if: <code>EDOM</code> The value x is $\pm\text{Inf}$. | | | | |
| USAGE | An application wishing to check for error situations should set <code>errno</code> to 0 before calling <code>acos()</code> . If <code>errno</code> is non-zero on return, or the value NaN is returned, an error has occurred. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>cos(3M)</code> , <code>isnan(3M)</code> , <code>matherr(3M)</code> , <code>attributes(5)</code> , <code>standards(5)</code> | | | | |

| NAME | acosh, asinh, atanh – inverse hyperbolic functions | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double acosh (double <i>x</i>); double asinh (double <i>x</i>); double atanh (double <i>x</i>); | | | | |
| DESCRIPTION | The acosh (), asinh () and atanh () functions compute the inverse hyperbolic cosine, sine, and tangent of their argument, respectively. | | | | |
| RETURN VALUES | The acosh (), asinh () and atanh () functions return the inverse hyperbolic cosine, sine, and tangent of their argument, respectively. The acosh () function returns NaN and sets <code>errno</code> to <code>EDOM</code> when its argument is less than 1.0. The atanh () function returns NaN and sets <code>errno</code> to <code>EDOM</code> when its argument has absolute value greater than 1.0. The atanh () function returns <code>261Inf</code> and sets <code>errno</code> to <code>ERANGE</code> when its argument is <code>2611.0</code> . If <i>x</i> is NaN, the asinh (), acosh () and atanh () functions return NaN. For exceptional cases, matherr (3M) tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The acosh () function will fail if: EDOM The <i>x</i> argument is less than 1.0. The atanh () function will fail if: EDOM The <i>x</i> argument has an absolute value greater than 1.0. ERANGE The <i>x</i> argument has an absolute value equal to 1.0 | | | | |
| ATTRIBUTES | See attributes (5) for descriptions of the following attributes: <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | cosh (3M) , matherr (3M) , sinh (3M) , tanh (3M) , attributes (5) , standards (5) | | | | |

NAME | asin – arc sine function

SYNOPSIS | cc [*flag ...*] *file ...* -lm [*library ...*]
 | #include <math.h>

DESCRIPTION | double **asin**(double *x*);
 | The `asin()` function computes the principal value of the arc sine of *x*. The value of *x* should be in the range $[-1,1]$.

RETURN VALUES | Upon successful completion, `asin()` returns the arc sine of *x*, in the range $[-\pi/2, \pi/2]$ radians. If the value of *x* is not in the range $[-1,1]$ and is not $\pm\text{Inf}$ or NaN, either 0.0 or NaN is returned and `errno` is set to `EDOM`.
 | If *x* is NaN, NaN is returned.
 | If *x* is $\pm\text{Inf}$, either 0.0 is returned and `errno` is set to `EDOM` or NaN is returned and `errno` may be set to `EDOM`.
 | For exceptional cases, `matherr(3M)` tabulates the values to be returned as dictated by Standards other than XPG4.

ERRORS | The `asin()` function will fail if:
 | `EDOM` The value *x* is not $\pm\text{Inf}$ or NaN and is not in the range $[-1,1]$.
 | The `asin()` function may fail if:
 | `EDOM` The value of *x* is $\pm\text{Inf}$.

USAGE | An application wishing to check for error situations should set `errno` to 0, then call `asin()`. If `errno` is non-zero on return, or the return value is NaN, an error has occurred.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `isnan(3M)`, `matherr(3M)`, `sin(3M)`, `attributes(5)`, `standards(5)`

NAME atan2 – arc tangent function

SYNOPSIS cc [flag ...] file ... -lm [library ...]
#include <math.h>
double **atan2**(double y, double x);

DESCRIPTION The atan2() function computes the principal value of the arc tangent of y/x, using the signs of both arguments to determine the quadrant of the return value.

RETURN VALUES Upon successful completion, atan2() returns the arc tangent of y/x in the range [-pi,pi] radians. If both arguments are 0.0, 0.0 is returned and errno may be set to EDOM.

If x or y is NaN, NaN is returned.

In IEEE 754 mode atan2() handles the following exceptional arguments in the spirit of ANSI/IEEE Std 754-1985.

- atan2(±0, x) returns ±0 for x > 0 or x = +0;
- atan2(±0, x) returns ±pi for x < 0 or x = -0;
- atan2(y, ±0) returns pi/2 for y > 0;
- atan2(y, ±0) returns -pi/2 for y < 0;
- atan2(±y, Inf) returns ±0 for finite y > 0;
- atan2(±Inf, x) returns ±pi/2 for finite x;
- atan2(±y, -Inf) returns ±pi for finite y > 0;
- atan2(±Inf, Inf) returns ±pi/4;
- atan2(±Inf, -Inf) returns ±3pi/4.

For exceptional cases, matherr(3M) tabulates the values to be returned as dictated by Standards other than XPG4.

ERRORS The atan2() function may fail if:
EDOM Both arguments are 0.0.

USAGE An application wishing to check for error situations should set errno to 0 before calling atan2(). If errno is non-zero on return, or the return value is NaN, an error has occurred.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO atan(3M), isnan(3M), matherr(3M), tan(3M), attributes(5), standards(5)

NAME atan – arc tangent function

SYNOPSIS cc [*flag ...*] *file ...* -lm [*library ...*]
#include <math.h>

double **atan**(double *x*);

DESCRIPTION The `atan()` function computes the principal value of the arc tangent of *x*.

RETURN VALUES Upon successful completion, `atan()` returns the arc tangent of *x* in the range $[-\pi/2, \pi/2]$ radians.

If *x* is NaN, NaN is returned.

If *x* is $\pm\text{Inf}$, $\pm\pi/2$ is returned.

ERRORS No errors will occur.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO `atan2(3M)`, `isnan(3M)`, `tan(3M)`, `attributes(5)`

NAME au_open, au_close, au_write – construct and write audit records

SYNOPSIS

```
cc [ flag ... ] file ... -lbsm -lsocket -lnsl -lintl [ library ... ]
#include <bsm/libbsm.h>
int au_close(int d, int keep, short event);

int au_open(void);

int au_write(int d, token_t *m);
```

DESCRIPTION

au_open() returns an audit record descriptor to which audit tokens can be written using au_write(). The audit record descriptor is an integer value that identifies a storage area where audit records are accumulated.

au_close() terminates the life of an audit record *d* of type *event* started by au_open(). If the *keep* parameter is zero, the data contained therein is discarded and the memory used is given up by calling free(3C) . Otherwise, the additional parameters are used to create a header token. Depending on the audit policy information obtained by auditon(2) , additional tokens such as *sequence* and *trailer* tokens may be added to the record. au_close() finally writes the record to the audit trail by calling audit(2) .

au_write() adds the audit token pointed to by *m* to the audit record identified by the descriptor *d* . After this call is made the audit token is no longer available to the caller.

RETURN VALUES

A successful invocation of au_write() and au_close() will return a 0 .

A successful invocation of au_open() returns an audit record descriptor. au_open() returns -1 if a descriptor could not be allocated. au_write() returns -1 if *d* is not a valid descriptor or if audit(2) experienced an error. errno is set to indicate the error. au_write() will return -1 if *d* is an invalid descriptor or if *m* is an invalid token.

ATTRIBUTES

See attributes (5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Safe |

SEE ALSO

bsmconv(1M) , audit(2) , auditon(2) , au_preselect(3BSM) , au_to(3BSM) , free(3C) , attributes(5)

NOTES

The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See bsmconv(1M) for more information.

| | | | | | | | | | | | |
|--|---|--|--|----------------|--|-------------|--|---------------|--|-----------------|--|
| NAME | au_preselect – preselect an audit event | | | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file... -lbsm -lsocket -lnsl -lintl [library ...] #include <bsm/libbsm.h></pre> | | | | | | | | | | |
| DESCRIPTION | <p><code>int au_preselect(audit_event_t event, au_mask_t *mask_p, int sorf, int flag);</code></p> <p><code>au_preselect()</code> determines whether or not the audit event <i>event</i> is preselected against the binary preselection mask pointed to by <i>mask_p</i> (usually obtained by a call to <code>getaudit(2)</code>). <code>au_preselect()</code> looks up the classes associated with <i>event</i> in <code>audit_event(4)</code> and compares them with the classes in <i>mask_p</i>. If the classes associated with <i>event</i> match the classes in the specified portions of the binary preselection mask pointed to by <i>mask_p</i>, the event is said to be preselected.</p> <p><i>sorf</i> indicates whether the comparison is made with the success portion, the failure portion or both portions of the mask pointed to by <i>mask_p</i>.</p> <p>The following are the valid values of <i>sorf</i>:</p> <table border="0"> <tr> <td>AU_PRS_SUCCESS</td> <td>Compare the event class with the success portion of the preselection mask.</td> </tr> <tr> <td>AU_PRS_FAILURE</td> <td>Compare the event class with the failure portion of the preselection mask.</td> </tr> <tr> <td>AU_PRS_BOTH</td> <td>Compare the event class with both the success and failure portions of the preselection mask.</td> </tr> </table> <p><i>flag</i> tells <code>au_preselect()</code> how to read the <code>audit_event(4)</code> database. Upon initial invocation, <code>au_preselect()</code> reads the <code>audit_event(4)</code> database and allocates space in an internal cache for each entry with <code>malloc(3C)</code>. In subsequent invocations, the value of <i>flag</i> determines where <code>au_preselect()</code> obtains audit event information. The following are the valid values of <i>flag</i>:</p> <table border="0"> <tr> <td>AU_PRS_REREAD</td> <td>Get audit event information by searching the <code>audit_event(4)</code> database.</td> </tr> <tr> <td>AU_PRS_USECACHE</td> <td>Get audit event information from internal cache created upon the initial invocation. This option is much faster.</td> </tr> </table> | AU_PRS_SUCCESS | Compare the event class with the success portion of the preselection mask. | AU_PRS_FAILURE | Compare the event class with the failure portion of the preselection mask. | AU_PRS_BOTH | Compare the event class with both the success and failure portions of the preselection mask. | AU_PRS_REREAD | Get audit event information by searching the <code>audit_event(4)</code> database. | AU_PRS_USECACHE | Get audit event information from internal cache created upon the initial invocation. This option is much faster. |
| AU_PRS_SUCCESS | Compare the event class with the success portion of the preselection mask. | | | | | | | | | | |
| AU_PRS_FAILURE | Compare the event class with the failure portion of the preselection mask. | | | | | | | | | | |
| AU_PRS_BOTH | Compare the event class with both the success and failure portions of the preselection mask. | | | | | | | | | | |
| AU_PRS_REREAD | Get audit event information by searching the <code>audit_event(4)</code> database. | | | | | | | | | | |
| AU_PRS_USECACHE | Get audit event information from internal cache created upon the initial invocation. This option is much faster. | | | | | | | | | | |
| RETURN VALUES | <p><code>au_preselect()</code> returns:</p> <table border="0"> <tr> <td>0</td> <td><i>event</i> is not preselected.</td> </tr> <tr> <td>1</td> <td><i>event</i> is preselected.</td> </tr> <tr> <td>-1</td> <td>An error occurred. <code>au_preselect()</code> couldn't allocate memory or couldn't find <i>event</i> in the <code>audit_event(4)</code> database.</td> </tr> </table> | 0 | <i>event</i> is not preselected. | 1 | <i>event</i> is preselected. | -1 | An error occurred. <code>au_preselect()</code> couldn't allocate memory or couldn't find <i>event</i> in the <code>audit_event(4)</code> database. | | | | |
| 0 | <i>event</i> is not preselected. | | | | | | | | | | |
| 1 | <i>event</i> is preselected. | | | | | | | | | | |
| -1 | An error occurred. <code>au_preselect()</code> couldn't allocate memory or couldn't find <i>event</i> in the <code>audit_event(4)</code> database. | | | | | | | | | | |
| FILES | <table border="0"> <tr> <td><code>/etc/security/audit_class</code></td> <td>maps audit class number to audit class names and descriptions</td> </tr> </table> | <code>/etc/security/audit_class</code> | maps audit class number to audit class names and descriptions | | | | | | | | |
| <code>/etc/security/audit_class</code> | maps audit class number to audit class names and descriptions | | | | | | | | | | |

/etc/security/audit_event maps audit even number to audit event names and associates

ATTRIBUTES

See attributes(5) for a description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

bsmconv(1M), getaudit(2), au_open(3BSM), getauclassent(3BSM), getaevent(3BSM), malloc(3C), audit_class(4), audit_event(4), attributes(5)

NOTES

au_preselect() is normally called prior to constructing and writing an audit record. If the event is not preselected, the overhead of constructing and writing the record can be saved.

The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See bsmconv(1M) for more information.

NAME | au_to, au_to_arg, au_to_attr, au_to_data, au_to_groups, au_to_in_addr, au_to_ipc, au_to_ipc_perm, au_to_iport, au_to_me, au_to_new_in_addr, au_to_new_process, au_to_new_socket, au_to_new_subject, au_to_opaque, au_to_path, au_to_process, au_to_return, au_to_socket, au_to_subject, au_to_text – create audit record tokens

SYNOPSIS

```
cc [ flag ... ] file ... -lbsm -lsocket -lnsl -lintl [ library ... ]
#include <sys/types.h>
#include <sys/vnode.h>
#include <netinet/in.h>
#include <bsm/libbsm.h>
token_t *au_to_arg(char n, char *text, u_long v);

token_t *au_to_attr(struct vattr *attr);

token_t *au_to_cmd(u_long argc, char **argv, char **envp);

token_t *au_to_data(char unit_print, char unit_type, char unit_count, char *p);

token_t *au_to_groups(int *groups);

token_t *au_to_in_addr(struct inaddr *internet_addr);

token_t *au_to_new_in_addr(struct inaddr *internet_addr);

token_t *au_to_iport(u_short_t iport);

token_t *au_to_ipc(int id);

token_t *au_to_ipc_perm(struct ipc_perm *perm);

token_t *au_to_iport(u_short_t iport);

token_t *au_to_me(void);

token_t *au_to_newgroups(int n, int *groups);

token_t *au_to_opaque(char *data, short bytes);

token_t *au_to_path(char *path);

token_t *au_to_process(au_id_t auid, uid_t euid, gid_t egid, uid_t ruid, gid_t rgid, pid_t pid, au_asid_t sid, au_tid_t *tid);

token_t *au_to_new_process(au_id_t auid, uid_t euid, gid_t egid, uid_t ruid, gid_t rgid, pid_t pid, au_asid_t sid, au_tid_addr_t *tid);

token_t *au_to_return(char number, uint_t value);

token_t *au_to_socket(struct socket *so);

token_t *au_to_new_socket(struct socket *so);
```

DESCRIPTION

`token_t *au_to_subject(au_id_t auid, uid_t euid, gid_t egid, uid_t ruid, gid_t rgid, pid_t pid, au_asid_t sid, au_tid_t *tid);`

`token_t *au_to_new_subject(au_id_t auid, uid_t euid, gid_t egid, uid_t ruid, gid_t rgid, pid_t pid, au_asid_t sid, au_tid_addr_t *tid);`

`token_t *au_to_text(char *text);`

The `au_to_arg()` function formats the data in *v* into an “argument token.” The *n* argument indicates the argument number. The *text* argument is a null terminated string describing the argument.

The `au_to_attr()` function formats the data pointed to by *attr* into a “vnode attribute token.”

The `au_to_data()` function formats the data pointed to by *p* into an “arbitrary data token.” The *unit_print* parameter determines the preferred display base of the data and is one of `AUP_BINARY`, `AUP_OCTAL`, `AUP_DECIMAL`, `AUP_HEX`, or `AUP_STRING`. The *unit_type* parameter defines the basic unit of data and is one of `AUR_BYTE`, `AUR_CHAR`, `AUR_SHORT`, `AUR_INT`, or `AUR_LONG`. The *unit_count* parameter specifies the number of basic data units to be used and must be positive.

The `au_to_groups()` function formats the array of 16 integers pointed to by *groups* into a “groups token.”

The `au_to_in_addr()` function formats the data pointed to by *internet_addr* into an “internet address token.”

The `au_to_new_in_addr()` function formats the data pointed to by *internet_addr* into an “internet address token.” The *internet_addr* is one containing an IPv6 IP address.

The `au_to_ipc()` function formats the data in the *id* parameter into an “interprocess communications ID token.”

The `au_to_ipc_perm()` function formats the data pointed to by *perm* into an “interprocess communications permission token.”

The `au_to_iport()` function formats the data pointed to by *iport* into an “ip port address token.”

The `au_to_me()` function collects audit information from the current process and creates a “subject token” by calling `au_to_subject()`.

The `au_to_newgroups()` function formats the array of *n* integers pointed to by *groups* into a “newgroups token.”

The `au_to_subject()` function formats an *auid* (audit user ID), an *euid* (effective user ID), an *egid* (effective group ID), a *ruid* (real user ID), an *rgid*

(real group ID), a *pid* (process ID), an *sid* (audit session ID), an *tid* (audit terminal ID), into a “subject token.”

The `au_to_new_subject()` function formats an *audit user ID*, an *effective user ID*, an *effective group ID*, a *real user ID*, an *real group ID*, a *pid* (process ID), an *sid* (audit session ID), an *tid* (audit terminal ID), into a “subject token.” The audit terminal ID is one that contains an IPv6 IP address.

The `au_to_opaque()` function formats the *bytes* bytes pointed to by *data* into an “opaque token.” The value of *size* must be positive.

The `au_to_path()` function formats the path name pointed to by *path* into a “path token.”

The `au_to_process()` function formats an *audit user ID*, an *effective user ID*, an *effective group ID*, a *real user ID*, a *real group ID*, a *pid* (process ID), an *sid* (audit session ID), and a *tid* (audit terminal ID), into a “process token.” A process token should be used when the process is the object of an action (ie. when the process is the receiver of a signal).

The `au_to_new_process()` function formats an *audit user ID*, an *effective user ID*, an *effective group ID*, a *real user ID*, a *real group ID*, a *pid* (process ID), an *sid* (audit session ID), and a *tid* (audit terminal ID), into a “process token.” A process token should be used when the process is the object of an action (ie. when the process is the receiver of a signal). The audit terminal ID is one that contains an IPv6 IP address.

The `au_to_return()` function formats an error number *number* and a return value *value* into a “return value token.”

The `au_to_socket()` function format the data pointed to by *so* into a “socket token.”

The `au_to_new_socket()` function format the data pointed to by *so* into a “socket token.” The socket contains IPv6 IP addresses.

The `au_to_text()` function formats the null-terminated string pointed to by *text* into a “text token.”

RETURN VALUES

These functions return `NULL` if memory cannot be allocated to put the resultant token into, or if an error in the input is detected.

ATTRIBUTES

See `attributes(5)` for a description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`bsmconv(1M)`, `au_open(3BSM)`, `attributes(5)`

NOTES

The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See `bsmconv(1M)` for more information.

| NAME | au_user_mask – get user's binary preselection mask | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lbsm -lsocket -lnsl -lintl [library ...] #include <bsm/libbsm.h></pre> | | | | |
| DESCRIPTION | <p>int au_user_mask(char *username, au_mask_t *mask_p);</p> <p>au_user_mask() reads the default, system wide audit classes from audit_control(4), combines them with the per-user audit classes from the audit_user(4) database, and updates the binary preselection mask pointed to by mask_p with the combined value.</p> <p>The audit flags in the flags field of the audit_control(4) database and the always-audit-flags and never-audit-flags from the audit_user(4) database represent binary audit classes. These fields are combined by au_preselect(3BSM) as follows:</p> $\text{mask} = (\text{flags} + \text{always-audit-flags}) - \text{never-audit-flags}$ <p>au_user_mask() only fails if both the both the audit_control(4) and the audit_user(4) database entries could not be retrieved. This allows for flexible configurations.</p> | | | | |
| RETURN VALUES | <p>au_user_mask() returns:</p> <p>0 Success.</p> <p>-1 Failure. Both the audit_control(4) and the audit_user(4) database entries could not be retrieved.</p> | | | | |
| FILES | <p>/etc/security/audit_control contains default parameters read by the audit daemon, auditd(1M)</p> <p>/etc/security/audit_user stores per-user audit event mask</p> | | | | |
| ATTRIBUTES | <p>See attributes(5) for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | login(1), bsmconv(1M), getaudit(2), setaudit(2), au_preselect(3BSM), getacinfo(3BSM), getausernam(3BSM), audit_control(4), audit_user(4), attributes(5) | | | | |
| NOTES | au_user_mask() should be called by programs like login(1) which set a process's preselection mask with setaudit(2). getaudit(2) should be used to obtain audit characteristics for the current process. | | | | |

The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See [bsmconv\(1M\)](#) for more information.

| NAME | bgets – read stream up to next delimiter | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lgen [<i>library ...</i>] #include <libgen.h> char * bgets (char * <i>buffer</i> , size_t * <i>count</i> , FILE * <i>stream</i> , const char * <i>breakstring</i>); | | | | |
| DESCRIPTION | <p>The <code>bgets()</code> function reads characters from <i>stream</i> into <i>buffer</i> until either <i>count</i> is exhausted or one of the characters in <i>breakstring</i> is encountered in the stream. The read data is terminated with a null byte ('\0') and a pointer to the trailing null is returned. If a <i>breakstring</i> character is encountered, the last non-null is the delimiter character that terminated the scan.</p> <p>Note that, except for the fact that the returned value points to the end of the read string rather than to the beginning, the call</p> <pre>bgets(buffer, sizeof buffer, stream, "\n");</pre> <p>is identical to</p> <pre>fgets (buffer, sizeof buffer, stream);</pre> <p>There is always enough room reserved in the buffer for the trailing null character.</p> <p>If <i>breakstring</i> is a null pointer, the value of <i>breakstring</i> from the previous call is used. If <i>breakstring</i> is null at the first call, no characters will be used to delimit the string.</p> | | | | |
| RETURN VALUES | NULL is returned on error or end-of-file. Reporting the condition is delayed to the next call if any characters were read but not yet returned. | | | | |
| EXAMPLES | <p>EXAMPLE 1 Example of <code>bgets()</code> function.</p> <pre>#include <libgen.h> char buffer[8]; /* read in first user name from /etc/passwd */ fp = fopen("/etc/passwd", "r"); bgets(buffer, 8, fp, ":");</pre> | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>gets(3C)</code> , <code>attributes(5)</code> | | | | |
| NOTES | When compiling multi-thread applications, the <code>_REENTRANT</code> flag must be defined on the compile line. This flag should only be used in multi-thread applications. | | | | |

NAME bufsplit – split buffer into fields

SYNOPSIS cc [*flag ...*] *file ...* -lgen [*library ...*]
 #include <libgen.h>
 size_t **bufsplit**(char **buf*, size_t *n*, char ***a*);

DESCRIPTION bufsplit() examines the buffer, *buf*, and assigns values to the pointer array, *a*, so that the pointers point to the first *n* fields in *buf* that are delimited by TABs or NEWLINES.

To change the characters used to separate fields, call `bufsplit()` with *buf* pointing to the string of characters, and *n* and *a* set to zero. For example, to use colon (:), period (.), and comma (,), as separators along with TAB and NEWLINE:

```
bufsplit (":.,\t\n", 0, (char**)0 );
```

RETURN VALUES The number of fields assigned in the array *a*. If *buf* is zero, the return value is zero and the array is unchanged. Otherwise the value is at least one. The remainder of the elements in the array are assigned the address of the null byte at the end of the buffer.

EXAMPLES **EXAMPLE 1** Example of `bufsplit()` function.

```
/*
 * set a[0] = "This", a[1] = "is", a[2] = "a",
 * a[3] = "test"
 */
bufsplit("This\tis\ta\ttest\n", 4, a);
```

NOTES bufsplit() changes the delimiters to null bytes in *buf*.

When compiling multithreaded applications, the `_REENTRANT` flag must be defined on the compile line. This flag should only be used in multithreaded applications.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO attributes(5)

NAME | cbrt – cube root function

SYNOPSIS | cc [*flag ...*] *file ...* -lm [*library ...*]
 | #include <math.h>
 | double **cbrt**(double *x*);

DESCRIPTION | The `cbrt()` function computes the cube root of *x*.

RETURN VALUES | On successful completion, `cbrt()` returns the cube root of *x*. If *x* is NaN, `cbrt()` returns NaN.

ERRORS | No errors will occur.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `attributes(5)`

| NAME | ceil – ceiling value function | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double ceil (double <i>x</i>); | | | | |
| DESCRIPTION | The <code>ceil()</code> function computes the smallest integral value not less than <i>x</i> . | | | | |
| RETURN VALUES | Upon successful completion, <code>ceil()</code> returns the smallest integral value not less than <i>x</i> , expressed as a type <code>double</code> . If <i>x</i> is NaN, NaN is returned. If <i>x</i> is ±Inf or ±0, <i>x</i> is returned. | | | | |
| ERRORS | No errors will occur. | | | | |
| USAGE | The integral value returned by <code>ceil()</code> as a <code>double</code> may not be expressible as an <code>int</code> or <code>long int</code> . The return value should be tested before assigning it to an integer type to avoid the undefined results of an integer overflow. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>floor(3M)</code> , <code>isnan(3M)</code> , <code>attributes(5)</code> | | | | |

NAME | `config_admin`, `config_change_state`, `config_private_func`, `config_test`,
`config_stat`, `config_list`, `config_list_ext`, `config_ap_id_cmp`, `config_unload`,
`config_strerror` – configuration administration interface

SYNOPSIS

```
cc [ flag ] file -lcfgadm -ldevinfo -ldl [ library... ]
#include <config_admin.h>
cfga_err_t config_change_state(cfga_cmd_t state_change_cmd, int num_ap_ids, char *
const *ap_ids, const char *options, struct cfga_confirm *confp, struct cfga_msg *msgp, char
**errstring, cfga_flags_t flags);

cfga_err_t config_private_func(const char *function, int num_ap_ids, char * const
*ap_ids, const char *options, struct cfga_confirm *confp, struct cfga_msg *msgp, char
**errstring, cfga_flags_t flags);

cfga_err_t config_test(int num_ap_ids, char * const *ap_ids, const char *options, struct
cfga_msg *msgp, char **errstring, cfga_flags_t flags);

cfga_err_t config_list_ext(int num_ap_ids, char * const *ap_ids, struct cfga_list_data
**ap_id_list, int *nlist, const char *options, const char *listops, char **errstring, cfga_flags_t
flags);

int config_ap_id_cmp(const cfga_ap_id_t ap_id1, const cfga_ap_id_t ap_id2);

void config_unload_libs();
```

Deprecated Interfaces

The following interfaces have been deprecated and their use is strongly discouraged:

```
cfga_err_t config_stat(int num_ap_ids, char * const *ap_ids, struct cfga_stat_data
*buf, const char *options, char **errstring);

cfga_err_t config_list(struct cfga_stat_data **ap_id_list, int *nlist, const char *options,
char **errstring);
```

**HARDWARE
DEPENDENT
LIBRARY
SYNOPSIS**

The `config_admin` library is a generic interface that is used for dynamic configuration, (DR). Each piece of hardware that supports DR must supply a hardware-specific *plugin* library that contains the entry points listed in this subsection. The generic library will locate and link to the appropriate library to effect DR operations. The interfaces specified in this subsection are really "hidden" from users of the generic libraries. It is, however, necessary that writers of the hardware-specific plug in libraries know what these interfaces are.

```
cfga_err_t cfga_change_state(cfga_cmd_t state_change_cmd, const char *ap_id,
const char *options, struct cfga_confirm *confp, struct cfga_msg *msgp, char **errstring,
cfga_flags_t flags);
```

```
cfga_err_t cfga_private_func(const char *function, const char *ap_id, const char
*options, struct cfga_confirm *confp, struct cfga_msg *msgp, char **errstring, cfga_flags_t
flags);
```

```
cfga_err_t cfga_test(const char *ap_id, const char *options, struct cfga_msg *msgp, char
** errstring, cfga_flags_t flags);
```

```
cfga_err_t cfga_list_ext(const char *ap_id, struct cfga_list_data **ap_id_list, int *nlist,
const char *options, const char *listopts, char **errstring, cfga_flags_t flags);
```

```
cfga_err_t cfga_help(struct cfga_msg *msgp, const char *options, cfga_flags_t flags);
```

```
int cfga_ap_id_cmp(const cfga_ap_id_t ap_id1, const cfga_ap_id_t ap_id2);
```

Deprecated Interfaces

The following interfaces have been deprecated and their use is strongly discouraged:

```
cfga_err_t cfga_stat(const char *ap_id, struct cfga_stat_data *buf, const char *options,
char **errstring);
```

```
cfga_err_t cfga_list(const char *ap_id, struct cfga_stat_data **ap_id_list, int *nlist, const
char *options, char **errstring);
```

DESCRIPTION

The `config_*`() functions provide a hardware independent interface to hardware-specific system configuration administration functions. The `cfga_*`() functions are provided by hardware-specific libraries that are dynamically loaded to handle configuration administration functions in a hardware-specific manner.

The `libcfgadm` library is used to provide the services of the `cfgadm(1M)` command. The hardware-specific libraries are located in `/usr/platform/${machine}/lib/cfgadm`, `/usr/platform/${arch}/lib/cfgadm`, and `/usr/lib/cfgadm`. The hardware-specific library names are derived from the driver name or from class names in device tree nodes that identify attachment points.

The `config_change_state`() function performs operations that change the state of the system configuration. The `state_change_cmd` argument can be one of the following: `CFGA_CMD_INSERT`, `CFGA_CMD_REMOVE`, `CFGA_CMD_DISCONNECT`, `CFGA_CMD_CONNECT`, `CFGA_CMD_CONFIGURE`, or `CFGA_CMD_UNCONFIGURE`. The `state_change_cmd` `CFGA_CMD_INSERT` is used to prepare for manual insertion or to activate automatic hardware insertion of an occupant. The `state_change_cmd` `CFGA_CMD_REMOVE` is used to prepare for manual removal or activate automatic hardware removal of an occupant. The `state_change_cmd` `CFGA_CMD_DISCONNECT` is used to disable normal communication to or from an occupant in a receptacle. The `state_change_cmd` `CFGA_CMD_CONNECT` is used to enable communication to or from an occupant in a receptacle. The `state_change_cmd` `CFGA_CMD_CONFIGURE` is used to bring the hardware resources contained on, or attached to, an occupant into the realm of

Solaris, allowing use of the occupant's hardware resources by the system. The *state_change_cmd* `CFGA_CMD_UNCONFIGURE` is used to remove the hardware resources contained on, or attached to, an occupant from the realm of Solaris, disallowing further use of the occupant's hardware resources by the system.

The *flags* argument may contain one or both of the defined flags, `CFGA_FLAG_FORCE` and `CFGA_FLAG_VERBOSE`. If the `CFGA_FLAG_FORCE` flag is asserted certain safety checks will be overridden. For example, this may not allow an occupant in the failed condition to be configured, but might allow an occupant in the failing condition to be configured. Acceptance of a force is hardware dependent. If the `CFGA_FLAG_VERBOSE` flag is asserted hardware-specific details relating to the operation are output utilizing the `cfga_msg` mechanism.

The `config_private_func()` function invokes private hardware-specific functions.

The `config_test()` function is used to initiate testing of the specified attachment point.

The *num_ap_ids* argument specifies the number of *ap_id*s in the *ap_ids* array. The *ap_ids* argument points to an array of *ap_id*s.

The *ap_id* argument points to a single *ap_id*.

The *function* and *options* strings conform to the `getsubopt(3C)` syntax convention and are used to supply hardware-specific function or option information. No generic hardware-independent functions or options are defined.

The `cfga_confirm` structure referenced by *confp* provides a call-back interface to get permission to proceed should the requested operation require, for example, a noticeable service interruption. The `cfga_confirm` structure includes the following members:

```
int (*confirm)(void *appdata_ptr, const char *message);
void *appdata_ptr;
```

The `confirm()` function is called with two arguments: the generic pointer *appdata_ptr* and the message detailing what requires confirmation. The generic pointer *appdata_ptr* is set to the value passed in in the `cfga_confirm` structure member `appdata_ptr` and can be used in a graphical user interface to relate the `confirm` function call to the `config_*` call. The `confirm` function should return 1 to allow the operation to proceed and 0 otherwise.

The `cfga_msg` structure referenced by *msgp* provides a call-back interface to output messages from a hardware-specific library. In the presence of the `CFGA_FLAG_VERBOSE` flag, these messages can be informational; otherwise

they are restricted to error messages. The `cfga_msg` structure includes the following members:

```
void (*message_routine)(void *appdata_ptr, const char *message);
void *appdata_ptr;
```

The `message_routine()` function is called with two arguments: the generic pointer `appdata_ptr` and the message. The generic pointer `appdata_ptr` is set to the value passed in in the `cfga_confirm` structure member `appdata_ptr` and can be used in a graphical user interface to relate the `message_routine()` function call to the `config_*` call. The messages must be in the native language specified by the `LC_MESSAGES` locale category; see `setlocale(3C)`.

For some generic errors a hardware-specific error message can be returned. The storage for the error message string, including the terminating null character, is allocated by the `config_*` functions using `malloc(3C)` and a pointer to this storage returned through `errstring`. If `errstring` is `NULL` no error message will be generated or returned. If `errstring` is not `NULL` and no error message is generated, the pointer referenced by `errstring` will be set to `NULL`. It is the responsibility of the function calling `config_*` to deallocate the returned storage using `free(3C)`. The error messages must be in the native language specified by the `LC_MESSAGES` locale category; see `setlocale(3C)`.

The `config_list_ext()` function provides the listing interface. When supplied with a list of `ap_id`s through the first two arguments, it returns an array of `cfga_list_data_t` structures for each attachment point specified. If the first two arguments are 0 and `NULL` respectively, then all attachment points in the device tree will be listed. Additionally, dynamic expansion of an attachment point to list dynamic attachment points may also be requested by passing the `CFG_FLAG_LIST_ALL` flag through the `flags` argument. Storage for the returned array of `stat` structures is allocated by the `config_list_ext()` function using `malloc(3C)`. This storage must be freed by the caller of `config_list_ext()` by using `free(3C)`.

The `cfga_list_data` structure includes the following members:

```
cfga_log_ext_t      ap_log_id;          /* Attachment point logical id */
cfga_phys_ext_t    ap_phys_id;        /* Attachment point physical id */
cfga_class_t       ap_class;          /* Attachment point class */
cfga_stat_t        ap_r_state;        /* Receptacle state */
cfga_stat_t        ap_o_state;        /* Occupant state */
cfga_cond_t        ap_cond;          /* Attachment point condition */
cfga_busy_t        ap_busy;          /* Busy indicator */
time_t             ap_status_time;    /* Attachment point last change*/
cfga_info_t        ap_info;          /* Miscellaneous information */
cfga_type_t        ap_type;          /* Occupant type */
```

The types are defined as follows:


```

typedef char  cfga_log_ext_t[CFGA_LOG_EXT_LEN];
typedef char  cfga_phys_ext_t[CFGA_PHYS_EXT_LEN];
typedef char  cfga_class_t[CFGA_CLASS_LEN];
typedef char  cfga_info_t[CFGA_INFO_LEN];
typedef char  cfga_type_t[CFGA_TYPE_LEN];
typedef enum  cfga_cond_t;
typedef enum  cfga_stat_t;
typedef enum  cfga_busy_t;
typedef int   cfga_flags_t;

```

The *listopts* argument to `config_list_ext()` conforms to the `getsubopt(3C)` syntax and is used to pass listing sub-options. Currently, only the sub-option `class=class_name` is supported. This list option restricts the listing to attachment points of class `class_name`.

The *listopts* argument to `cfga_list_ext()` is reserved for future use. Hardware-specific libraries should ignore this argument if it is `NULL`. If *listopts* is not `NULL` and is not supported by the hardware-specific library, an appropriate error code should be returned.

The `ap_log_id` and the `ap_phys_id` members give the hardware-specific logical and physical names of the attachment point. The `ap_busy` member indicates activity is present that may result in changes to state or condition. The `ap_status_time` member provides the time at which either the `ap_r_state`, `ap_o_state`, or `ap_cond` field of the attachment point last changed. The `ap_info` member is available for the hardware-specific code to provide additional information about the attachment point. The `ap_class` member contains the attachment point class (if any) for an attachment point. The `ap_class` member is filled in by the generic library. If the `ap_log_id` and `ap_phys_id` members are not filled in by the hardware-specific library, the generic library will fill in these members using a generic format. The remaining members are the responsibility of the corresponding hardware-specific library.

The `ap_log_id`, `ap_phys_id`, `ap_info`, `ap_class`, and `ap_type` members are null-terminated strings. When printing these fields, the following format is suggested:

```
printf("%.*s", sizeof(p->ap_log_id), p->ap_log_id);
```

The `config_stat()`, `config_list()`, `cfga_stat()`, and `cfga_list()` functions and the `cfga_stat_data` data structure are deprecated interfaces and are provided solely for backward compatibility. Use of these interfaces is strongly discouraged.

The `config_ap_id_cmp` function performs a hardware dependent comparison on two *ap_id*s, returning an equal to, less than or greater than indication in the manner of `strcmp(3C)`. Each argument is either a `cfga_ap_id_t` or can be

a null-terminated string. This function can be used when sorting lists of *ap_id* s, for example with `qsort(3C)` , or when selecting entries from the result of a `config_list` function call.

The `config_unload_libs` function unlinks all previously loaded hardware-specific libraries.

The `config_strerror` function can be used to map an error return value to an error message string. See RETURN VALUES . The returned string should not be overwritten. `config_strerror` returns NULL if *cfgerrnum* is out-of-range.

The `cfga_help` function can be used request that a hardware-specific library output it's localized help message.

RETURN VALUES

The `config_*`() and `cfga_*`() functions return the following values. Additional error information may be returned through *errstring* if the return code is not CFGA_OK . See DESCRIPTION for details.

| | |
|-----------------------------|---|
| CFGA_BUSY | The command was not completed due to an element of the system configuration administration system being busy. |
| CFGA_ATTR_INVALID | No attachment points with the specified attributes exists |
| CFGA_ERROR | An error occurred during the processing of the requested operation. This error code includes validation of the command arguments by the hardware-specific code. |
| CFGA_INSUFFICIENT_CONDITION | Operation failed due to attachment point condition. |
| CFGA_INVALID | The system configuration administration operation requested is not supported on the specified attachment point. |
| CFGA_LIB_ERROR | A procedural error occurred in the library, including failure to obtain process resources such as memory and file descriptors. |
| CFGA_NACK | The command was not completed due to a negative acknowledgement from the <i>confp ->confirm</i> function. |

| | |
|------------------|--|
| CFGA_NO_LIB | A hardware-specific library could not be located using the supplied <i>ap_id</i> . |
| CFGA_NOTSUPP | System configuration administration is not supported on the specified attachment point. |
| CFGA_OK | The command completed as requested. |
| CFGA_OPNOTSUPP | System configuration administration operation is not supported on this attachment point. |
| CFGA_PRIV | The caller does not have the required process privileges. For example, if configuration administration is performed through a device driver, the permissions on the device node would be used to control access. |
| CFGA_SYSTEM_BUSY | The command required a service interruption and was not completed due to a part of the system that could not be quiesced. |

ERRORS

Many of the errors returned by the system configuration administration functions are hardware-specific. The strings returned in *errstring* may include the following:

attachment point *ap_id* not known

The attachment point detailed in the error message does not exist.

unknown hardware option *option* for operation

An unknown option was encountered in the *options* string.

hardware option *option* requires a value

An option in the *options* string should have been of the form *option =value* .

listing option *list_option* requires a value

An option in the *listopts* string should have been of the form *option =value* .

hardware option *option* does not require a value

An option in the *options* string should have been a simple option.

attachment point *ap_id* is not configured

A *config_change_state* command to CFGA_CMD_UNCONFIGURE an occupant was made to an attachment point whose occupant was not in the CFGA_STAT_CONFIGURED state.

attachment point *ap_id* is not unconfigured
 A *config_change_state* command requiring an unconfigured occupant was made to an attachment point whose occupant was not in the CFGA_STAT_UNCONFIGURED state.

attachment point *ap_id* condition not satisfactory
 A *config_change_state* command was made to an attachment point whose condition prevented the operation.

attachment point *ap_id* in condition *condition* cannot be used
 A *config_change_state* operation with force indicated was directed to an attachment point whose condition fails the hardware dependent test.

ATTRIBUTES

See *attributes(5)* for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|------------------|
| Availability | SUNWcsu, SUNWkvm |
| MT-Level | Safe |

SEE ALSO

cfgadm(1M), *devinfo(1M)*, *dlopen(3DL)*, *dlsym(3DL)*, *free(3C)*, *getsubopt(3C)*, *malloc(3C)*, *qsort(3C)*, *setlocale(3C)*, *strcmp(3C)*, *libcfgadm(3LIB)*, *attributes(5)*

NOTES

Applications using this library should be aware that the underlying implementation may use system services which alter the contents of the external variable *errno* and may use file descriptor resources.

The following code shows the intended error processing when *config_*()* returns a value other than *CFGA_OK* :

```
void
emit_error(int cfgernum, char *estrp)
{
    const char *ep;
    ep = config_strerror(cfgernum);
    if (ep == NULL)
        ep = gettext("configuration administration unknown error");
    if (estrp != NULL && *estrp != '\\0') {
        (void) fprintf(stderr, "%s: %s\
", ep, estrp);
    } else {
        (void) fprintf(stderr, "%s\
", ep);
    }
    if (estrp != NULL)
        free((void *)estrp);
}
```

Reference should be made to the Hardware Specific Guide for details of System Configuration Administration support.

NAME | ConnectToServer – connect to a DMI service provider

SYNOPSIS | `cc [flag ...] file ... -ldmici -ldmimi [library ...]`
`#include <dmi/api.hh>`
`bool_t ConnectToServer(ConnectI *argp, DmiRpcHandle *dmi_rpc_handle);`

DESCRIPTION | The `ConnectToServer()` function enables a management application or a component instrumentation to connect to a DMI service provider.

The `argp` parameter is an input parameter that uses the following data structure:

```
struct ConnectIN {
    char      *host;
    const char *nettype;
    ServerType servertype;
    RpcType   rpctype;
}
```

The `host` member indicates the host on which the service provider is running. The default is `localhost`.

The `nettype` member specifies the type of transport RPC uses. The default is `netpath`.

The `servertype` member indicates whether the connecting process is a management application or a component instrumentation.

The `rpctype` member specifies the type of RPC, either ONC or DCE. Only ONC is supported in the Solaris 7 release.

The `dmi_rpc_handle` parameter is the output parameter that returns DMI RPC handle.

RETURN VALUES | The `ConnectToServer()` function returns TRUE if successful, otherwise FALSE.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | Safe |

SEE ALSO | `DisconnectToServer(3DMI)`, `attributes(5)`

NAME | copylist – copy a file into memory

SYNOPSIS | `cc [flag ...] file ... -lgen [library ...]`
`#include <libgen.h>`
`char *copylist(const char *filenm, off_t *szptr);`

DESCRIPTION | The `copylist()` function copies a list of items from a file into freshly allocated memory, replacing new-lines with null characters. It expects two arguments: a pointer `filenm` to the name of the file to be copied, and a pointer `szptr` to a variable where the size of the file will be stored.

Upon success, `copylist()` returns a pointer to the memory allocated. Otherwise it returns NULL if it has trouble finding the file, calling `malloc()`, or reading the file.

USAGE | The `copylist()` function has a transitional interface for 64-bit file offsets. See `lf64(5)`.

EXAMPLES | **EXAMPLE 1** Example of `copylist()` function.

```
/* read "file" into buf */
off_t size;
char *buf;
buf = copylist("file", &size);
if (buf) {
    for (i=0; i<size; i++)
        if (buf[i])
            putchar(buf[i]);
        else
            putchar('\n');
} else {
    fprintf(stderr, "%s: Copy failed for "file".\n", argv[0]);
    exit (1);
}
```

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `malloc(3C)`, `attributes(5)`, `lf64(5)`

NOTES | When compiling multithreaded applications, the `_REENTRANT` flag must be defined on the compile line. This flag should only be used in multithreaded applications.

NAME copysign – return magnitude of first argument and sign of second argument

SYNOPSIS cc [*flag ...*] *file ...* -lm [*library ...*]
 #include <math.h>
 double **copysign**(double *x*, double *y*);

DESCRIPTION The `copysign()` function returns a value with the magnitude of *x* and the sign of *y*. It produces a NaN with the sign of *y* if *x* is a NaN.

RETURN VALUES The `copysign()` function returns a value with the magnitude of *x* and the sign of *y*.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO `attributes(5)`

NAME | cos – cosine function

SYNOPSIS | cc [*flag ...*] *file ...* -lm [*library ...*]
 | #include <math.h>
 | double **cos**(double *x*);

DESCRIPTION | The `cos()` function computes the cosine of *x*, measured in radians.

RETURN VALUES | Upon successful completion, `cos()` returns the cosine of *x*.
 | If *x* is NaN or ±Inf, NaN is returned.

ERRORS | No errors will occur.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `acos(3M)`, `isnan(3M)`, `sin(3M)`, `tan(3M)`, `attributes(5)`

| NAME | cosh – hyperbolic cosine function | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double cosh (double <i>x</i>); | | | | |
| DESCRIPTION | The <code>cosh()</code> function computes the hyperbolic cosine of <i>x</i> . | | | | |
| RETURN VALUES | Upon successful completion, <code>cosh()</code> returns the hyperbolic cosine of <i>x</i> . If the result would cause an overflow, <code>HUGE_VAL</code> is returned and <code>errno</code> is set to <code>ERANGE</code> . If <i>x</i> is NaN, NaN is returned. For exceptional cases, <code>matherr(3M)</code> tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The <code>cosh()</code> function will fail if: <code>ERANGE</code> The result would cause an overflow. | | | | |
| USAGE | An application wishing to check for error situations should set <code>errno</code> to 0 before calling <code>cosh()</code> . If <code>errno</code> is non-zero on return, or the returned value is NaN, an error has occurred. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>acosh(3M)</code> , <code>isnan(3M)</code> , <code>matherr(3M)</code> , <code>sinh(3M)</code> , <code>tanh(3M)</code> , <code>attributes(5)</code> , <code>standards(5)</code> | | | | |

| | | | | | |
|---|---|--------------------------------|--|----------------------------------|--|
| NAME | cpc – hardware performance counters | | | | |
| DESCRIPTION | <p>The UltraSPARC and Pentium microprocessor families contain <i>hardware performance counters</i> that allow the measurement of many different hardware events related to CPU behavior, including instruction and data cache misses as well as various internal states of the processor. More recent processors allow a variety of events to be captured. The counters can be configured to count user events or system events, or both. The two processor families currently share the restriction that only two event types can be measured simultaneously.</p> <p>UltraSPARC III and Pentium II processors are able to generate an interrupt on counter overflow, allowing the counters to be used for various forms of profiling.</p> <p>This manual page describes a set of APIs that allow Solaris applications to use these counters. Applications can measure their own behavior, the behavior of other applications, or the behavior of the whole system.</p> | | | | |
| Shared counters or private counters? | <p>There are two principal models for using these performance counters. Some users of these statistics wish to observe system-wide behavior; others wish to view the performance counters as part of the register set exported by each LWP. On a machine performing more than one activity, these two models are in conflict because the counters represent a critical hardware resource that cannot simultaneously be both shared and private.</p> <p>To fully support the two-level threads model in Solaris, it would be necessary to virtualize the performance counters to each thread. This version of the library does not allow per-thread data to be captured unless bound threads are used. Even without bound threads, however, the counters can still be used to assess aggregate program behavior.</p> | | | | |
| Generic or specific events? | <p>Although some events are common to all processors, it is apparent that the counters expose a great deal of the specific implementation details of the processor architecture. For this reason, events are specified by name using a string-based hardware event specification language. The values of the tokens in the language vary from processor model to processor model, and can only be interpreted with reference to the relevant hardware documentation. The functions provided to specify the strings use environment variables or arguments so that the names do not have to be compiled in applications, thus extending their longevity and portability across platforms and processor generations.</p> | | | | |
| Configuration Interfaces | <p>The following configuration interfaces are provided:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>cpc_version(3CPC)</code></td> <td>check the version the application was compiled with against the version of the library</td> </tr> <tr> <td><code>cpc_getcpuver(3CPC)</code></td> <td>determine the performance counter version of the current CPU</td> </tr> </table> | <code>cpc_version(3CPC)</code> | check the version the application was compiled with against the version of the library | <code>cpc_getcpuver(3CPC)</code> | determine the performance counter version of the current CPU |
| <code>cpc_version(3CPC)</code> | check the version the application was compiled with against the version of the library | | | | |
| <code>cpc_getcpuver(3CPC)</code> | determine the performance counter version of the current CPU | | | | |

| | |
|-----------------------------------|---|
| <code>cpc_getcciname(3CPC)</code> | return the corresponding printable string to describe that interface |
| <code>cpc_getnpic(3CPC)</code> | return the number of valid counter registers in the <code>cpc_event(3CPC)</code> data structure |
| <code>cpc_getcpuref(3CPC)</code> | return a reference to the corresponding processor documentation |

Programming events

Events are specified using a `getsubopt(3C)`-style language for both the events and the additional control bits that determine what causes the counters to increment. The `cpc_strtoevent()` function translates a string to an event specification which can then be used to program the counters. The `cpc_eventtostr()` function returns the canonical form of the string that corresponds to a particular event. The `cpc_getusage(3CPC)` function returns a string that specifies the syntax of the string, while `cpc_walk_names(3CPC)` allows the caller to apply a function to each possible event supported on the relevant processor.

Performance counter context

Each processor on the system possesses its own set of performance counter registers. For a single process, it is often desirable to maintain the illusion that the counters are an intrinsic part of that process (whichever processors it runs on), since this allows the events to be directly attributed to the process without having to make passive all other activity on the system.

To achieve this behavior, the library associates *performance counter context* with each LWP in the process; the context consists of a small amount of kernel memory to hold the counter values when the LWP is not running, and some simple kernel functions to save and restore those counter values from and to the hardware registers when the LWP performs a normal context switch. A process can only observe and manipulate its own copy of the performance counter control and data registers.

Performance Counters In Other Processes

Though applications can be modified to instrument themselves as demonstrated above, it is frequently useful to be able to examine the behavior of an existing application without changing the source code. A separate library, `libpctx`, provides a simple set of interfaces that use the facilities of `proc(4)` to control a target process, and together with functions in `libcpc`, allow `truss`-like tools to be constructed to measure the performance counters in other applications. An example of one such application is `cputrack(1)`.

The functions in `libpctx` are independent of those in `libcpc`. These functions manage a process using an event-loop paradigm — that is, the execution of certain system calls by the controlled process cause the library to stop the controlled process and execute callback functions in the context of the controlling

process. These handlers can perform various operations on the target process using APIs in `libpctx` and `libcpc` that consume `pctx_t` handles.

SEE ALSO

`cputrack(1)`, `cpustat(1M)`, `cpc_access(3CPC)`, `cpc_bind_event(3CPC)`,
`cpc_count_usr_events(3CPC)`, `cpc_pctx_bind_event(3CPC)`,
`cpc_event(3CPC)`, `cpc_event_diff(3CPC)`, `cpc_getcpuver(3CPC)`,
`cpc_seterrfn(3CPC)`, `cpc_shared_bind_event(3CPC)`,
`cpc_strtoevent(3CPC)`, `cpc_version(3CPC)`, `pctx_capture(3CPC)`,
`pctx_set_events(3CPC)`, `proc(4)`.

| NAME | cpc_access – test access CPU performance counters | | | | | | | | |
|----------------------|---|----------------|-----------------|----------|---------|--------------|---|---------------------|----------|
| SYNOPSIS | cc [<i>flag...</i>] <i>file...</i> -lcpc [<i>library...</i>] #include <libcpc.h> int cpc_access (void); | | | | | | | | |
| DESCRIPTION | <p>Access to CPU performance counters is possible only on systems where the appropriate hardware exists and is correctly configured. The <code>cpc_access()</code> function <i>must</i> be used to determine if the hardware exists and is accessible on the platform before any of the interfaces that use the counters are invoked.</p> <p>When the hardware is available, access to the per-process counters is always allowed to the process itself, and allowed to other processes mediated using the existing security mechanisms of <code>/proc</code>.</p> | | | | | | | | |
| RETURN VALUES | <p>Upon successful completion, <code>cpc_access()</code> returns 0. Otherwise, it returns -1 and sets <code>errno</code> to indicate the error.</p> <p>By default, two common <code>errno</code> values are decoded and cause the library to print an error message using its reporting mechanism. See <code>cpc_seterrfn(3CPC)</code> to for a description of how this behavior can be modified.</p> | | | | | | | | |
| ERRORS | <p>The <code>cpc_access()</code> function will fail if:</p> <p>EAGAIN Another process may be sampling system-wide CPU statistics.</p> <p>ENOSYS CPU performance counters are inaccessible on this machine.</p> | | | | | | | | |
| ATTRIBUTES | <p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> <tr> <td>Availability</td> <td>SUNWcpcu (32-bit) SUNWcpcux (64-bit)</td> </tr> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe | Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) | Interface Stability | Evolving |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | | | | | |
| MT-Level | MT-Safe | | | | | | | | |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) | | | | | | | | |
| Interface Stability | Evolving | | | | | | | | |
| SEE ALSO | <code>cpc(3CPC)</code> , <code>cpc_seterrfn(3CPC)</code> , <code>proc(4)</code> , <code>attributes(5)</code> | | | | | | | | |

| | |
|----------------------|---|
| NAME | cpc_bind_event, cpc_take_sample, cpc_rele – use CPU performance counters on lwps |
| SYNOPSIS | <pre>cc [flag...] file... -lcpc [library...] #include <libcpc.h> int cpc_bind_event(cpc_event_t *event, int flags); int cpc_take_sample(cpc_event_t *event); int cpc_rele(void);</pre> |
| DESCRIPTION | <p>Once the events to be sampled have been selected using, for example, <code>cpc_strtoevent(3CPC)</code>, the event selections can be bound to the calling LWP using <code>cpc_bind_event()</code>. If <code>cpc_bind_event()</code> returns successfully, the system has associated performance counter context with the calling LWP. The context allows the system to virtualize the hardware counters to that specific LWP, and the counters are enabled.</p> <p>Two flags are defined that can be passed into the routine to allow the behavior of the interface to be modified, as described below.</p> <p>Counter values can be sampled at any time by calling <code>cpc_take_sample()</code>, and dereferencing the fields of the <code>ce_pic []</code> array returned. The <code>ce_hrt</code> field contains the timestamp at which the kernel last sampled the counters.</p> <p>To immediately remove the performance counter context on an LWP, the <code>cpc_rele()</code> interface should be used. Otherwise, the context will be destroyed after the LWP or process exits.</p> <p>The caller should take steps to ensure that the counters are sampled often enough to avoid the 32-bit counters wrapping. The events most prone to wrap are those that count processor clock cycles. If such an event is of interest, sampling should occur frequently so that less than 4 billion clock cycles can occur between samples. Practically speaking, this is only likely to be a problem for otherwise idle systems, or when processes are bound to processors, since normal context switching behavior will otherwise hide this problem.</p> |
| RETURN VALUES | Upon successful completion, <code>cpc_bind_event()</code> and <code>cpc_take_sample()</code> return 0. Otherwise, these functions return -1, and set <code>errno</code> to indicate the error. |
| ERRORS | <p>The <code>cpc_bind_event()</code> and <code>cpc_take_sample()</code> functions will fail if:</p> <p><code>EFAULT</code> The <i>event</i> argument specifies a bad address.</p> <p><code>ENOTSUP</code> The caller has attempted an operation that is illegal or not supported on the current platform, such as attempting to specify signal delivery on counter overflow on a CPU that doesn't generate an interrupt on counter overflow.</p> |

| | |
|--------|---|
| EAGAIN | Another process may be sampling system-wide CPU statistics. For <code>cpc_bind_event()</code> , this implies that no new contexts can be created. For <code>cpc_take_sample()</code> , this implies that the performance counter context has been invalidated and must be released with <code>cpc_rele()</code> . Robust programs should be coded to expect this behavior and recover from it by releasing the now invalid context by calling <code>cpc_rele()</code> sleeping for a while, then attempting to bind and sample the event once more. |
| EINVAL | The <code>cpc_take_sample()</code> function has been invoked before the context is bound. |

EXAMPLES

EXAMPLE 1 Use hardware performance counters to measure events in a process.

The example below shows how a standalone program can be instrumented with the `libcpc` routines to use hardware performance counters to measure events in a process. The program performs 20 iterations of a computation, measuring the counter values for each iteration. By default, the example makes the counters measure external cache references and external cache hits; these options are only appropriate for UltraSPARC processors. By setting the `PERFEVENTS` environment variable to other strings (a list of which can be gleaned from the `-h` flag of the `cpustat` or `cpustrack` utilities), other events can be counted. The `error()` routine below is assumed to be a user-provided routine analogous to the familiar `printf(3C)` routine from the C library but which also performs an `exit(2)` after printing the message.

```
#include <inttypes.h>
#include <stdlib.h>
#include <stdio.h>
#include <unistd.h>
#include <libcpc.h>
int
main(int argc, char *argv[])
{
    int cpuver, iter;
    char *setting = NULL;
    cpc_event_t event;

    if (cpc_version(CPC_VER_CURRENT) != CPC_VER_CURRENT)
        error("application:library cpc version mismatch!");

    if ((cpuver = cpc_getcpuver()) == -1)
        error("no performance counter hardware!");

    if ((setting = getenv("PERFEVENTS")) == NULL)
        setting = "pic0=EC_ref,pic1=EC_hit";

    if (cpc_strtoevent(cpuver, setting, &event) != 0)
        error("can't measure '%s' on this processor", setting);
    setting = cpc_eventtostr(&event);
```



```

if (cpc_access() == -1)
    error("can't access perf counters: %s", strerror(errno));

if (cpc_bind_event(&event, 0) == -1)
    error("can't bind lwp%d: %s", _lwp_self(), strerror(errno));

for (iter = 1; iter <= 20; iter++) {
    cpc_event_t before, after;

    if (cpc_take_sample(&before) == -1)
        break;

    /* ==> Computation to be measured goes here <== */

    if (cpc_take_sample(&after) == -1)
        break;
    (void) printf("%3d: %" PRId64 " %" PRId64 "\n",
        iter,
        after.ce_pic[0] - before.ce_pic[0],
        after.ce_pic[1] - before.ce_pic[1]);
}

if (iter != 20)
    error("can't sample '%s': %s", setting, strerror(errno));

free(setting);
return (0);
}

```

EXAMPLE 2 Write a signal handler to catch overflow signals.

This example builds on Example 1, but demonstrates how to write the signal handler to catch overflow signals. The counters are preset so that counter zero is 1000 counts short of overflowing, while counter one is set to zero. After 1000 counts on counter zero, the signal handler will be invoked.

First the signal handler:

```

#definePRESET0      (UINT64_MAX - 1000ull)
#definePRESET1      0
void
emt_handler(int sig, siginfo_t *sip, void *arg)
{
    ucontext_t *uap = arg;
    cpc_event_t sample;

    if (sig != SIGEMT || sip->si_code != EMT_CPCOVF) {
        psignal(sig, "example");
        psiginfo(sip, "example");
        return;
    }

    (void) printf("lwp%d - si_addr %p ucontext: %%pc %p %%sp %p\n",
        _lwp_self(), (void *)sip->si_addr,

```

```

        (void *)uap->uc_mcontext.gregs[PC],
        (void *)uap->uc_mcontext.gregs[USP]);

if (cpc_take_sample(&sample) == -1)
    error("can't sample: %s", strerror(errno));

(void) printf("0x%" PRIx64 " 0x%" PRIx64 "\n",
            sample.ce_pic[0], sample.ce_pic[1]);
(void) fflush(stdout);

sample.ce_pic[0] = PRESET0;
sample.ce_pic[1] = PRESET1;
if (cpc_bind_event(&sample, CPC_BIND_EMT_OVF) == -1)
    error("cannot bind lwp%d: %s", _lwp_self(), strerror(errno));
}

```

and second the setup code (this can be placed after the code that selects the event to be measured):

```

struct sigaction act;
cpc_event_t event;
...
act.sa_sigaction = emt_handler;
bzero(&act.sa_mask, sizeof (act.sa_mask));
act.sa_flags = SA_RESTART|SA_SIGINFO;
if (sigaction(SIGEMT, &act, NULL) == -1)
    error("sigaction: %s", strerror(errno));
event.ce_pic[0] = PRESET0;
event.ce_pic[1] = PRESET1;
if (cpc_bind_event(&event, CPC_BIND_EMT_OVF) == -1)
    error("cannot bind lwp%d: %s", _lwp_self(), strerror(errno));

for (iter = 1; iter <= 20; iter++) {
    /* ==> Computation to be measured goes here <== */
}

cpc_bind_event(NULL, 0);    /* done */

```

Note that a more general version of the signal handler would use `write(2)` directly instead of depending on the signal-unsafe semantics of `stderr` and `stdout`. Most real signal handlers will probably do more with the samples than just print them out.

NOTES

Sometimes, even the overhead of performing a system call will be too disruptive to the events being measured. Once a call to `cpc_bind_event()` has been issued, it is possible to directly access the performance hardware registers from within the application. If the performance counter context is active, then the counters will count on behalf of the current LWP.

SPARC

```

rd %pic, %rN      ! All UltraSPARC
wr %rN, %pic     ! (ditto, but see text)

```

IA

rdpmc ! Pentium II only

If the counter context is not active or has been invalidated, the `%pic` register (SPARC), and the `rdpmc` instruction (Pentium) will become unavailable.

Note that the two 32-bit UltraSPARC performance counters are kept in the single 64-bit `%pic` register so a couple of additional instructions are required to separate the values. Also note that when the `%pcr` register bit has been set that configures the `%pic` register as readable by an application, it is also writable. Any values written will be preserved by the context switching mechanism.

Pentium II processors support the non-privileged `rdpmc` instruction which requires [5] that the counter of interest be specified in `%ecx`, and returns a 40-bit value in the `%edx:%eax` register pair. There is no non-privileged access mechanism for Pentium I processors.

Handling counter overflow

As described above, when counting events, some processors allow their counter registers to silently overflow. However more recent CPUs such as UltraSPARC III and Pentium II are capable of generating an interrupt when the hardware counter overflows.

The most obvious use for this facility is to ensure that the full 64-bit counter values are maintained without repeated sampling. However, current hardware does not record which counter overflowed. A more subtle use for this facility is to preset the counter to a value to a little less than the maximum value, then use the resulting interrupt to catch the counter overflow associated with that event. The overflow can then be used as an indication of the frequency of the occurrence of that event.

Note that the interrupt generated by the processor may not be particularly precise. That is, the particular instruction that caused the counter overflow may be earlier in the instruction stream than is indicated by the program counter value in the `ucontext`.

When `cpc_bind_event()` is called with the `CPC_BIND_EMT_OVF` flag set, then as before, the control registers and counters are preset from the 64-bit values contained in `event`. However, when the flag is set, the kernel arranges to send the calling process a `SIGEMT` signal when the overflow occurs, with the `si_code` field of the corresponding `siginfo` structure set to `EMT_CPCOVF`, and the `si_addr` field is the program counter value at the time the overflow interrupt was delivered. Counting, and thus the subsequent delivery of the signal on overflow is disabled until the next call to `cpc_bind_event()`. Even in a multithreaded process, during execution of the signal handler, the thread behaves as if it is temporarily bound to the running LWP.

Different processors have different counter ranges available, though all processors supported by Solaris allow at least 31 bits to be specified as a counter preset value; thus portable preset values lie in the range `UINT64_MAX` to `UINT64_MAX - INT32_MAX`.

The appropriate preset value will often need to be determined experimentally. Typically, it will depend on the event being measured, as well as the desire to minimize the impact of the act of measurement on the event being measured; less frequent interrupts and samples lead to less perturbation of the system.

If the processor cannot detect counter overflow, this call will fail (`ENOTSUP`). Specifying a null event unbinds the context from the underlying LWP and disables signal delivery. Currently, only user events can be measured using this technique. See Example 2, above.

Inheriting events onto multiple LWP s

By default, the library binds the performance counter context to the current LWP only. If the `CPC_BIND_LWP_INHERIT` flag is set, then any subsequent LWP s created by that LWP will automatically inherit the same performance counter context. The counters will be initialized to 0 as if a `cpc_bind_event()` had just been issued. This automatic inheritance behavior can be useful when dealing with multithreaded programs to determine aggregate statistics for the program as a whole.

If the `CPC_BIND_EMT_OVF` flag is also set, the process will immediately dispatch a `SIGEMT` signal to the freshly created LWP so that it can preset its counters appropriately on the new LWP . This initialization condition can be detected using `cpc_take_sample()` to check that both `ce_pic []` values are set to `UINT64_MAX`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|---|
| MT-Level | MT-Safe |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) |
| Interface Stability | Evolving |

SEE ALSO

`cpustat(1)`, `cpc(3CPC)`, `cpc_strtoevent(3CPC)`, `attributes(5)`

| | | | | | |
|----------------------|--|--------|---|--------|--|
| NAME | cpc_count_usr_events, cpc_count_sys_events – enable and disable performance counters | | | | |
| SYNOPSIS | <pre>cc [flag...] file... -lcpc [library...] #include <libcpc.h> int cpc_count_usr_events(int enable); int cpc_count_sys_events(int enable);</pre> | | | | |
| DESCRIPTION | <p>In certain applications, it can be useful to explicitly enable and disable performance counters at different times so that the performance of a critical algorithm can be examined. The <code>cpc_count_usr_events()</code> function can be used to control whether events are counted on behalf of the application running in user mode, while <code>cpc_count_sys_events()</code> can be used to control whether events are counted on behalf of the application while it is running in the kernel, without otherwise disturbing the binding of events to the invoking LWP. If the <i>enable</i> argument is non-zero, counting of events is enabled, otherwise they are disabled.</p> | | | | |
| RETURN VALUES | <p>Upon successful completion, <code>cpc_count_usr_events()</code> and <code>cpc_count_sys_events()</code> return 0. Otherwise, the functions return -1 and set <code>errno</code> to indicate the error.</p> | | | | |
| ERRORS | <p>The <code>cpc_count_usr_events()</code> and <code>cpc_count_sys_events()</code> functions will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">EAGAIN</td> <td>The associated performance counter context has been invalidated by another process.</td> </tr> <tr> <td style="vertical-align: top;">EINVAL</td> <td>No performance counter context has been created, or an attempt was made to enable system events while delivering counter overflow signals.</td> </tr> </table> | EAGAIN | The associated performance counter context has been invalidated by another process. | EINVAL | No performance counter context has been created, or an attempt was made to enable system events while delivering counter overflow signals. |
| EAGAIN | The associated performance counter context has been invalidated by another process. | | | | |
| EINVAL | No performance counter context has been created, or an attempt was made to enable system events while delivering counter overflow signals. | | | | |
| EXAMPLES | <p>EXAMPLE 1 Use <code>cpc_count_usr_events()</code> to minimize code needed by application.</p> <p>In this example, the routine <code>cpc_count_usr_events()</code> is used to minimize the amount of code that needs to be added to the application. The <code>cputrack(1)</code> command can be used in conjunction with these interfaces to provide event programming, sampling, and reporting facilities.</p> <p>If the application is instrumented in this way and then started by <code>cputrack</code> with the <code>nouser</code> flag set in the event specification, counting of user events will only be enabled around the critical code section of interest. If the program is run normally, no harm will ensue.</p> <pre>int have_counters = 0; int main(int argc, char *argv[]) { if (cpc_version(CPC_VER_CURRENT) == CPC_VER_CURRENT &&</pre> | | | | |

```

    cpc_getcpuver() != -1 && cpc_access() == 0)
    have_counters = 1;

    /* ... other application code */

    if (have_counters)
        (void) cpc_count_usr_events(1);

    /* ==> Code to be measured goes here <== */

    if (have_counters)
        (void) cpc_count_usr_events(0);

    /* ... other application code */
}

```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|---|
| MT-Level | MT-Safe |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) |
| Interface Stability | Evolving |

SEE ALSO

`cputrack(1)`, `cpc(3CPC)`, `cpc_access(3CPC)`, `cpc_version(3CPC)`,
`cpc_getcpuver(3CPC)`, `cpc_bind_event(3CPC)`,
`cpc_pctx_bind_event(3CPC)`, `attributes(5)`

NAME | cpc_event – data structure to describe CPU performance counters

SYNOPSIS | #include <libcpc.h>

DESCRIPTION | The libcpc interfaces manipulate CPU performance counters using the cpc_event_t data structure. This structure contains several fields that are common to all processors, and some that are processor-dependent. These structures can be declared by a consumer of the API, thus the size and offsets of the fields and the entire data structure are fixed per processor for any particular version of the library. See cpc_version(3CPC) for details of library versioning.

SPARC | For UltraSPARC, the structure contains the following members:

```
typedef struct {
    int ce_cpuver;
    hrtime_t ce_hrt;
    uint64_t ce_tick;
    uint64_t ce_pic[2];
    uint64_t ce_pcr;
} cpc_event_t;
```

IA | For Pentium, the structure contains the following members:

```
typedef struct {
    int ce_cpuver;
    hrtime_t ce_hrt;
    uint64_t ce_tsc;
    uint64_t ce_pic[2];
    uint32_t ce_pes[2];
#define ce_cesr ce_pes[0]
} cpc_event_t;
```

The APIs are used to manipulate the highly processor-dependent control registers (the ce_pcr, ce_cesr, and ce_pes fields); the programmer is strongly advised not to reference those fields directly in portable code. The ce_pic array elements contain 64-bit accumulated counter values. The hardware registers are virtualized to 64-bit quantities even though the underlying hardware only supports 32-bits (UltraSPARC) or 40-bits (Pentium) before overflow.

The ce_hrt field is a high resolution timestamp taken at the time the counters were sampled by the kernel. This uses the same timebase as gethrtime(3C).

On SPARC V9 machines, the number of cycles spent running on the processor is computed from samples of the processor-dependent %tick register, and placed in the ce_tick field. On Pentium processors, the processor-dependent time-stamp counter register is similarly sampled and placed in the ce_tsc field.

ATTRIBUTES | See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWcpcu |

SEE ALSO

gethrtime(3C), cpc(3CPC), cpc_version(3CPC), attributes(5).

| NAME | cpc_event_diff, cpc_event_accum – simple difference and accumulate operations | | | | | | | | |
|--------------------------------|---|----------------|-----------------|----------|---------|--------------|---|---------------------|----------|
| SYNOPSIS | <pre>cc [flag...] file... -lcpc [library...] #include <libcpc.h> void cpc_event_accum(cpc_event_t *accum, cpc_event_t *event); void cpc_event_diff(cpc_event_t *diff, cpc_event_t *after, cpc_event_t *before);</pre> | | | | | | | | |
| DESCRIPTION | <p>The <code>cpc_event_accum()</code> and <code>cpc_event_diff()</code> functions perform common accumulate and difference operations on <code>cpc_event(3CPC)</code> data structures. Use of these functions increases program portability, since structure members are not referenced directly .</p> | | | | | | | | |
| <code>cpc_event_accum()</code> | <p>The <code>cpc_event_accum()</code> function adds the <code>ce_pic</code> fields of <code>event</code> into the corresponding fields of <code>accum</code> . The <code>ce_hrt</code> field of <code>accum</code> is set to the later of the times in <code>event</code> and <code>accum</code> .</p> <p>SPARC:</p> <p>The function adds the contents of the <code>ce_tick</code> field of <code>event</code> into the corresponding field of <code>accum</code> .</p> <p>IA:</p> <p>The function adds the contents of the <code>ce_tsc</code> field of <code>event</code> into the corresponding field of <code>accum</code> .</p> | | | | | | | | |
| <code>cpc_event_diff()</code> | <p>The <code>cpc_event_diff()</code> function places the difference between the <code>ce_pic</code> fields of <code>after</code> and <code>before</code> and places them in the corresponding field of <code>diff</code> . The <code>ce_hrt</code> field of <code>diff</code> is set to the <code>ce_hrt</code> field of <code>after</code> .</p> <p>SPARC:</p> <p>Additionally, the function computes the difference between the <code>ce_tick</code> fields of <code>after</code> and <code>before</code> , and places it in the corresponding field of <code>diff</code> .</p> <p>IA:</p> <p>Additionally, the function computes the difference between the <code>ce_tsc</code> fields of <code>after</code> and <code>before</code> , and places it in the corresponding field of <code>diff</code> .</p> | | | | | | | | |
| ATTRIBUTES | <p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> <tr> <td>Availability</td> <td>SUNWcpcu (32-bit) SUNWcpcux (64-bit)</td> </tr> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe | Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) | Interface Stability | Evolving |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | | | | | |
| MT-Level | MT-Safe | | | | | | | | |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) | | | | | | | | |
| Interface Stability | Evolving | | | | | | | | |

SEE ALSO | `cpc(3CPC)`, `cpc_event(3CPC)`, `attributes(5)`.

| | |
|--------------------|--|
| NAME | cpc_getcpuver, cpc_getcciname, cpc_getcpuref, cpc_getusage, cpc_getnpic, cpc_walk_names – determine CPU performance counter configuration |
| SYNOPSIS | <pre>cc [flag...] file... -lcpc [library...] #include <libcpc.h> int cpc_getcpuver(void); const char *cpc_getcciname(int cpuver); const char *cpc_getcpuref(int cpuver); const char *cpc_getusage(int cpuver); uint_t cpc_getnpic(int cpuver); void cpc_walk_names(int cpuver, int regno, void *arg, void (*action)(void *arg, int regno, const char *name, uint8_t bits));</pre> |
| DESCRIPTION | <p>The <code>cpc_getcpuver()</code> function returns an abstract integer that corresponds to the distinguished version of the underlying processor. The library distinguishes between processors solely on the basis of their support for performance counters, so the version returned should not be interpreted in any other way. The set of values returned by the library is unique across all processor implementations.</p> <p>The <code>cpc_getcpuver()</code> function returns -1 if the library cannot support CPU performance counters on the current architecture. This may be because the processor has no such counter hardware, or because the library is unable to recognize it. Either way, such a return value indicates that the configuration functions described on this manual page cannot be used.</p> <p>The <code>cpc_getcciname()</code> function returns a printable description of the processor performance counter interfaces—for example, the string <i>UltraSPARC I&II</i>. Note that this name should not be assumed to be the same as the name the manufacturer might otherwise ascribe to the processor. It simply names the performance counter interfaces as understood by the library, and thus names the set of performance counter events that can be described by that interface. If the <code>cpuver</code> argument is unrecognized, the function returns <code>NULL</code>.</p> <p>The <code>cpc_getcpuref()</code> function returns a string that describes a reference work that should be consulted to (allow a human to) understand the semantics of the performance counter events that are known to the library. If the <code>cpuver</code> argument is unrecognized, the function returns <code>NULL</code>.</p> <p>The <code>cpc_getusage()</code> function returns a compact description of the <code>getsubopt()</code>-oriented syntax that is consumed by <code>cpc_strtoevent(3CPC)</code>. It is returned as a space-separated set of tokens to allow the caller to wrap lines at convenient boundaries. If the <code>cpuver</code> argument is unrecognized, the function returns <code>NULL</code>.</p> |

The `cpc_getnpic()` function returns the number of valid fields in the `ce_pic[]` array of a `cpc_event_t` data structure.

The library maintains a list of events that it believes the processor capable of measuring, along with the bit patterns that must be set in the corresponding control register, and which counter the result will appear in. The `cpc_walk_names()` function calls the `action()` function on each element of the list so that an application can print appropriate help on the set of events known to the library. The `arg` parameter is passed uninterpreted from the caller on each invocation of the `action()` function.

If the parameters specify an invalid or unknown CPU or register number, the function silently returns without invoking the action function.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|---|
| MT-Level | MT-Safe |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) |
| Interface Stability | Evolving |

SEE ALSO

`cpc(3CPC)`, `attributes(5)`.

NOTES

Only SPARC processors are described by the SPARC version of the library, and only Intel processors are described by the Intel version of the library.

| NAME | cpc_pctx_bind_event, cpc_pctx_take_sample, cpc_pctx_rele, cpc_pctx_invalidate – access CPU performance counters in other processes | | | | | | |
|----------------------|---|----------------|-----------------|----------|--------|--------------|-------------------|
| SYNOPSIS | <pre>cc [flag...] file... -lcpc -lpctx [library...] #include <libpctx.h> #include <libcpc.h> int cpc_pctx_bind_event(pctx_t *pctx, id_t lwpid, cpc_event_t *event, int flags); int cpc_pctx_take_sample(pctx_t *pctx, id_t lwpid, cpc_event_t *event); int cpc_pctx_rele(pctx_t *pctx, id_t lwpid); int cpc_pctx_invalidate(pctx_t *pctx, id_t lwpid);</pre> | | | | | | |
| DESCRIPTION | <p>These functions are designed to be run in the context of an event handler created using the libpctx(3LIB) family of functions that allow the caller, also known as the <i>controlling process</i>, to manipulate the performance counters in the context of a <i>controlled process</i>. The controlled process is described by the <i>pctx</i> argument, which must be obtained from an invocation of pctx_capture(3CPC) or pctx_create(3CPC) and passed to the functions described on this page in the context of an event handler.</p> <p>The semantics of the functions cpc_pctx_bind_event(), cpc_pctx_take_sample(), and cpc_pctx_rele() are directly analogous to those of cpc_bind_event(), cpc_take_sample(), and cpc_rele() described on the cpc_bind_event(3CPC) manual page.</p> <p>The cpc_pctx_invalidate() function allows the performance context to be invalidated in an LWP in the controlled process.</p> | | | | | | |
| RETURN VALUES | These functions return 0 on success. On failure, they return -1 and set errno to indicate the error. | | | | | | |
| ERRORS | <p>The cpc_pctx_bind_event(), cpc_pctx_take_sample(), and cpc_pctx_rele() functions return the same errno values the analogous functions described on the cpc_bind_event(3CPC) manual page. In addition, these function may fail if:</p> <p>ESRCH The value of the <i>lwpid</i> argument is invalid in the context of the controlled process.</p> | | | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Unsafe</td> </tr> <tr> <td>Availability</td> <td>SUNWcpcu (32-bit)</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | Unsafe | Availability | SUNWcpcu (32-bit) |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | | | |
| MT-Level | Unsafe | | | | | | |
| Availability | SUNWcpcu (32-bit) | | | | | | |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|--------------------|
| | SUNWcpcux (64-bit) |
| Interface Stability | Evolving |

SEE ALSO

cpc(3CPC) , cpc_bind_event(3CPC) , pctx_capture(3CPC) ,
pctx_create(3CPC) , attributes(5) .

NOTES

The capability to create and analyze overflow events in other processes is not available, though it may be made available in a future version of this API. In the current implementation, the *flags* field must be specified as 0.

NAME | cpc_seterrfn – control libcpc error reporting

SYNOPSIS | cc [*flag...*] *file...* -lcpc [*library...*]#include <libcpc.h>
 typedef void (cpc_errfn_t)(const char *fn, const char *fmt, va_list ap);
 void cpc_seterrfn(cpc_errfn_t *errfn);

DESCRIPTION | For the convenience of programmers instrumenting their code, several libcpc functions automatically emit to `stderr` error messages that attempt to provide a more detailed explanation of their error return values. While this can be useful for simple programs, some applications may wish to report their errors differently—for example, to a window or to a log file.

The `cpc_seterrfn()` function allows the caller to provide an alternate function for reporting errors; the type signature is shown above. The *fn* argument is passed the library function name that detected the error, the format string *fmt* and argument pointer *ap* can be passed directly to `vsnprintf(3C)` or similar `varargs`-based routine for formatting.

The default printing routine can be restored by calling the routine with an *errfn* argument of `NULL`.

EXAMPLES | **EXAMPLE 1** Debugging example.

This example produces error messages only when debugging the program containing it, or when the `cpc_strtoevent()` function is reporting an error when parsing an event specification

```
int debugging;
void
myapp_errfn(const char *fn, const char *fmt, va_list ap)
{
    if (strcmp(fn, "strtoevent") != 0 && !debugging)
        return;
    (void) fprintf(stderr, "myapp: cpc_%s(): ", fn);
    (void) vfprintf(stderr, fmt, ap);
}
```

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|---|
| MT-Level | Unsafe |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) |
| Interface Stability | Evolving |

SEE ALSO | `cpc(3CPC)`, `vsnprintf(3C)`, `attributes(5)`.

| | |
|--------------------|--|
| NAME | cpc_shared_open, cpc_shared_bind_event, cpc_shared_take_sample, cpc_shared_rele, cpc_shared_close – use CPU performance counters on processors |
| SYNOPSIS | <pre>cc [flag...] file... -lcpc [library...] #include <libcpc.h> int cpc_shared_open(void); int cpc_shared_bind_event(int fd, cpc_event_t *event, int flags); int cpc_shared_take_sample(int fd, cpc_event_t *event); int cpc_shared_rele(int fd); void cpc_shared_close(int fd);</pre> |
| DESCRIPTION | <p>The <code>cpc_shared_open()</code> function allows the caller to access the hardware counters in such a way that the performance of the currently bound CPU can be measured. The function returns a file descriptor if successful. Only one such open can be active at a time on any CPU.</p> <p>The <code>cpc_shared_bind_event()</code>, <code>cpc_shared_take_sample()</code>, and <code>cpc_shared_rele()</code> functions are directly analogous to the corresponding <code>cpc_bind_event()</code>, <code>cpc_take_sample()</code>, and <code>cpc_rele()</code> functions described on the <code>cpc_bind_event(3CPC)</code> manual page, except that they operate on the counters of a particular processor.</p> |
| USAGE | <p>If a thread wishes to access the counters using this interface, it must do so using a thread bound to an lwp, (see the <code>THR_BOUND</code> flag to <code>thr_create(3THR)</code>), that has in turn bound itself to a processor using <code>processor_bind(2)</code>.</p> <p>Unlike the <code>cpc_bind_event(3CPC)</code> family of functions, no counter context is attached to those lwps, so the performance counter samples from the processors reflects the system-wide usage, instead of per-lwp usage.</p> <p>The first successful invocation of <code>cpc_shared_open()</code> will immediately invalidate <i>all</i> existing performance counter context on the system, and prevent <i>all</i> subsequent attempts to bind counter context to lwps from succeeding anywhere on the system until the last caller invokes <code>cpc_shared_close()</code>.</p> <p>This is because it is impossible to simultaneously use the counters to accurately measure per-lwp and system-wide events, so there is an exclusive interlock between these uses.</p> <p>Access to the shared counters is mediated by file permissions on a <code>cpc</code> pseudo device. As shipped, only the superuser is allowed to access the shared device; this is because doing so prevents use of the counters on a per-lwp basis to any other users.</p> |

The `CPC_BIND_LWP_INHERIT` and `CPC_BIND_EMT_OVF` flags are invalid for the shared interface.

RETURN VALUES

On success, the functions (apart from `cpc_shared_close()`) return 0. On failure, the functions return -1 and set `errno`, to indicate the reason.

ERRORS

- `ENXIO` The current machine either has no performance counters, or has been configured to disallow access to them system-wide.
- `EACCES` The caller does not have appropriate privilege to access the CPU performance counters system-wide.
- `EAGAIN` For `cpc_shared_open()`, this value implies that the counters on the bound cpu are busy because they are already being used to measure system-wide events by some other caller.
- `EAGAIN` Otherwise, this return value implies that the counters are not available because the thread has been unbound from the processor it was bound to at open time. Robust programs should be coded to expect this behavior, and should invoke `cpc_shared_close()`, before retrying the operation.
- `EINVAL` The counters cannot be accessed on the current CPU because the calling thread is not bound to that CPU using `processor_bind(2)`.
- `EFAULT` The *event* argument specifies a bad address.
- `ENOTSUP` The caller has attempted an operation that is illegal or not supported on the current platform.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|--------------------|
| MT-Level | MT-Safe |
| Availability | SUNWcpcu (32-bit) |
| | SUNWcpcux (64-bit) |
| Interface Stability | Evolving |

SEE ALSO

`processor_bind(2)`, `cpc(3CPC)`, `cpc_bind_event(3CPC)`, `thr_create(3THR)`, `attributes(5)`

| | |
|----------------------|--|
| NAME | cpc_strtoevent, cpc_eventtostr – translate strings to and from events |
| SYNOPSIS | <pre>cc [flag...] file... -lcpc [library...] #include <libcpc.h> int cpc_strtoevent(int <i>cpuver</i>, const char *<i>spec</i>, cpc_event_t *<i>event</i>); char *cpc_eventtostr(cpc_event_t *<i>event</i>);</pre> |
| DESCRIPTION | <p>The <code>cpc_strtoevent()</code> function translates an event specification to the appropriate collection of control bits in a <code>cpc_event_t</code> structure pointed to by the <code>event</code> argument. The event specification is a <code>getsubopt(3C)</code>-style string that describes the event and any attributes that the processor can apply to the event or events. If successful, the function returns 0, the <code>ce_cpuver</code> field and the ISA-dependent control registers of event are initialized appropriately, and the rest of the <code>cpc_event_t</code> structure is initialized to 0.</p> <p>The <code>cpc_eventtostr()</code> function takes an event and constructs a compact canonical string representation for that event.</p> |
| RETURN VALUES | <p>Upon successful completion, <code>cpc_strtoevent()</code> returns 0. If the string cannot be decoded, a non-zero value is returned and a message is printed using the library's error-reporting mechanism (see <code>cpc_seterrfn(3CPC)</code>).</p> <p>Upon successful completion, <code>cpc_eventtostr()</code> returns a pointer to a string. The string returned must be freed by the caller using <code>free(3C)</code>. If <code>cpc_eventtostr()</code> a null pointer is returned.</p> |
| USAGE | <p>The event selection syntax used is processor architecture-dependent. The supported processor families allow variations on how events are counted as well as what events can be counted. This information is available in compact form from the <code>cpc_getusage()</code> function (see <code>cpc_getcpuver(3CPC)</code>), but is explained in further detail below.</p> |
| UltraSPARC | <p>On UltraSPARC processors, the syntax for setting options is as follows:</p> <pre>pic0=<eventspec>,pic1=<eventspec> [,sys] [,nouser]</pre> <p>This syntax, which reflects the simplicity of the options available using the <code>%pcr</code> register, forces both counter events to be selected. By default only user events are counted; however, the <code>sys</code> keyword allows system (kernel) events to be counted as well. User event counting can be disabled by specifying the <code>nouser</code> keyword.</p> <p>The keywords <code>pic0</code> and <code>pic1</code> may be omitted; they can be used to resolve ambiguities if they exist.</p> |
| Pentium I | <p>On Pentium processors, the syntax for setting counter options is as follows:</p> <pre>pic0=<eventspec>,pic1=<eventspec> [,sys[[0 1]]] [,nouser[[0 1]]] [,noedge[[0 1]]] [,pc[[0 1]]]</pre> |

The syntax and semantics are the same as UltraSPARC, except that it is possible to specify whether a particular counter counts user or system events. If unspecified, the specification is presumed to apply to both counters.

There are some additional keywords. The `noedge` keyword specifies that the counter should count clocks (duration) instead of events. The `pc` keyword allows the external pin control pins to be set high (defaults to low). When the pin control register is set high, the external pin will be asserted when the associated register overflows. When the pin control register is set low, the external pin will be asserted when the counter has been incremented. The electrical effect of driving the pin is dependent upon how the motherboard manufacturer has chosen to connect it, if it is connected at all.

Pentium II

For Pentium II processors, the syntax is substantially more complex, reflecting the complex configuration options available:

```
pic0=<eventspec>,pic1=<eventspec> [,sys[[0|1]]]
[,nouser[[0|1]]] [,noedge[[0|1]]] [,pc[[0|1]]] [,inv[[0|1]]] [,int[[0|1]]]
[,cmask[0|1]=<maskspec>] [,umask[0|1]=<maskspec>]
```

This syntax is a straightforward extension of the earlier syntax. The additional `inv`, `int`, `cmask0`, `cmask1`, `umask0`, and `umask1` keywords allow extended counting semantics. The mask specification is a number between 0 and 255, expressed in hexadecimal, octal or decimal notation.

**EXAMPLES
SPARC****EXAMPLE 1 SPARC Example.**

```
cpc_event_t event;
char *setting = "pic0=EC_ref,pic1=EC_hit"; /* UltraSPARC-specific */

if (cpc_strtoevent(cpuver, setting, &event) != 0)
    /* can't measure 'setting' on this processor */
else
    setting = cpc_eventtostr(&event);
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|---|
| MT-Level | MT-Safe |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) |
| Interface Stability | Evolving |

SEE ALSO

`cpc(3CPC)`, `cpc_getcpuver(3CPC)`, `cpc_seterrfn(3CPC)`, `free(3C)`, `getsubopt(3C)`, `attributes(5)`

NOTES

These functions are provided as a convenience only. As new processors are usually released asynchronously with software, the library allows the `pic0` and `pic1` keywords to interpret numeric values specified directly in hexadecimal, octal, or decimal.

| NAME | cpc_version – coordinate CPC library and application versions | | | | | | | | |
|----------------------|---|----------------|-----------------|----------|--------|--------------|---|---------------------|----------|
| SYNOPSIS | <pre>cc [flag...] file... -lcpc [library...] #include <libcpc.h> uint_t cpc_version(uint_t version);</pre> | | | | | | | | |
| DESCRIPTION | The <code>cpc_version()</code> function takes an interface version as an argument and returns an interface version as a result. Usually, the argument will be the value of <code>CPC_VER_CURRENT</code> bound to the application when it was compiled. | | | | | | | | |
| RETURN VALUES | <p>If the version requested is still supported by the implementation, <code>cpc_version()</code> returns the requested version number and the application can use the facilities of the library on that platform. If the implementation cannot support the version needed by the application, <code>cpc_version()</code> returns <code>CPC_VER_NONE</code>, indicating that the application will at least need to be recompiled to operate correctly on the new platform, and may require further changes.</p> <p>If <code>version</code> is <code>CPC_VER_NONE</code>, <code>cpc_version()</code> returns the most current version of the library.</p> | | | | | | | | |
| EXAMPLES | <p>EXAMPLE 1 Protect an application from using an incompatible library.</p> <p>The following lines of code protect an application from using an incompatible library:</p> <pre>if (cpc_version(CPC_VER_CURRENT) == CPC_VER_NONE) { /* version mismatch - library cannot translate */ exit(1); }</pre> | | | | | | | | |
| ATTRIBUTES | <p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Unsafe</td> </tr> <tr> <td>Availability</td> <td>SUNWcpcu (32-bit) SUNWcpcux (64-bit)</td> </tr> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | Unsafe | Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) | Interface Stability | Evolving |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | | | | | |
| MT-Level | Unsafe | | | | | | | | |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) | | | | | | | | |
| Interface Stability | Evolving | | | | | | | | |
| SEE ALSO | <code>cpc(3CPC)</code> , <code>attributes(5)</code> | | | | | | | | |
| NOTES | The version number is used only to express incompatible semantic changes in the performance counter interfaces on the given platform within a single instruction set architecture, for example, when a new set of performance counter registers are added to an existing processor family that cannot be specified in the existing <code>cpc_event_t</code> data structure. | | | | | | | | |

| | | | | | | | |
|----------------------|--|---|--|----------------|---|-----------------|---|
| NAME | demangle, cplus_demangle – decode a C++ encoded symbol name | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file[library ...] -ldemangle #include <demangle.h> int cplus_demangle(const char *symbol, char *prototype, size_t size);</pre> | | | | | | |
| DESCRIPTION | <p>The <code>cplus_demangle()</code> function decodes (demangles) a C++ linker symbol name (mangled name) into a (partial) C++ prototype, if possible. C++ mangled names may not have enough information to form a complete prototype.</p> <p>The <i>symbol</i> string argument points to the input mangled name.</p> <p>The <i>prototype</i> argument points to a user-specified output string buffer, of <i>size</i> bytes.</p> <p>The <code>cplus_demangle()</code> function operates on mangled names generated by SPARCompilers C++ 3.0.1, 4.0.1, 4.1 and 4.2.</p> <p>The <code>cplus_demangle()</code> function improves and replaces the <code>demangle()</code> function.</p> <p>Refer to the <code>CC.1</code>, <code>dem.1</code>, and <code>c++filt.1</code> manual pages in the <code>/opt/SUNWspro/man/man1</code> directory. These pages are only available with the SPROcc package.</p> | | | | | | |
| RETURN VALUES | <p>The <code>cplus_demangle()</code> function returns the following values:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">0</td> <td>The <i>symbol</i> argument is a valid mangled name and <i>prototype</i> contains a (partial) prototype for the symbol.</td> </tr> <tr> <td>DEMANGLE_ENAME</td> <td>The <i>symbol</i> argument is not a valid mangled name and the content of <i>prototype</i> is a copy of the symbol.</td> </tr> <tr> <td>DEMANGLE_ESPACE</td> <td>The <i>prototype</i> output buffer is too small to contain the prototype (or the symbol), and the content of <i>prototype</i> is undefined.</td> </tr> </table> | 0 | The <i>symbol</i> argument is a valid mangled name and <i>prototype</i> contains a (partial) prototype for the symbol. | DEMANGLE_ENAME | The <i>symbol</i> argument is not a valid mangled name and the content of <i>prototype</i> is a copy of the symbol. | DEMANGLE_ESPACE | The <i>prototype</i> output buffer is too small to contain the prototype (or the symbol), and the content of <i>prototype</i> is undefined. |
| 0 | The <i>symbol</i> argument is a valid mangled name and <i>prototype</i> contains a (partial) prototype for the symbol. | | | | | | |
| DEMANGLE_ENAME | The <i>symbol</i> argument is not a valid mangled name and the content of <i>prototype</i> is a copy of the symbol. | | | | | | |
| DEMANGLE_ESPACE | The <i>prototype</i> output buffer is too small to contain the prototype (or the symbol), and the content of <i>prototype</i> is undefined. | | | | | | |

| | |
|--------------------|---|
| NAME | devid_get, devid_free, devid_get_minor_name, devid_deviceid_to_nmlist, devid_free_nmlist, devid_compare, devid_sizeof – device ID interfaces for user applications |
| SYNOPSIS | <pre>cc [flag ...] file ... -ldevid [library ...] #include <devid.h> int devid_get(int fd, ddi_devid_t *retdevid); void devid_free(ddi_devid_t devid); int devid_get_minor_name(int fd, char **retminor_name); int devid_deviceid_to_nmlist(char *search_path, ddi_devid_t devid, char *minor_name, devid_nmlist_t **retlist); void devid_free_nmlist(devid_nmlist_t *list); int devid_compare(ddi_devid_t devid1, ddi_devid_t devid2); size_t devid_sizeof(ddi_devid_t devid);</pre> |
| DESCRIPTION | <p>These functions provide unique identifiers (device ID s) for devices. Applications and device drivers use these functions to identify and locate devices, independent of the device's physical connection or its logical device name or number.</p> <p>The <code>devid_get()</code> function returns in <code>retdevid</code> the device ID for the device associated with the open file descriptor <code>fd</code>, which refers to any device. It returns an error if the device does not have an associated device ID. The caller must free the memory allocated for <code>retdevid</code> using the <code>devid_free()</code> function.</p> <p>The <code>devid_free()</code> function frees the space that was allocated for <code>devid</code> by <code>devid_get()</code>.</p> <p>The <code>devid_get_minor_name()</code> function returns the minor name, in <code>retminor_name</code>, for the device associated with the open file descriptor <code>fd</code>. This name is specific to the particular minor number, but is "instance number" specific. The caller of this function must free the memory allocated for the returned string in <code>retminor_name</code> using the <code>devid_free()</code> function.</p> <p>The <code>devid_deviceid_to_nmlist()</code> function returns an array of <code>devid_nmlist</code> structures, where each entry matches the <code>devid</code> and minor name passed in. The <code>devid_nmlist</code> structure contains the device name and device number. The last entry of the array contains a null pointer for the <code>devname</code> and <code>NODEV</code> for the device number. This function traverses the file tree, starting at <code>search_path</code>. For each device with a matching device ID and minor name tuple, a device name and device number are added to the <code>retlist</code>. If no matches are found, an error is returned. The caller of this function must free the memory allocated for the returned array with the <code>devid_free_nmlist()</code> function.</p> |

The `devid_free_nmlist()` function frees the memory allocated by the `devid_deviceid_to_nmlist()` function.

The `devid_compare()` function compares two device IDs byte-by-byte and determines both equality and sort order. The function returns an integer greater than 0 if the device ID pointed to by `devid1` is greater than the device ID pointed to by `devid2`. It returns 0 if the device ID pointed to by `devid1` is equal to the device ID pointed to by `devid2`. It returns an integer less than 0 if the device ID pointed to by `devid1` is less than the device ID pointed to by `devid2`.

The `devid_sizeof()` function returns the size in number of bytes allocated for the `devid`.

RETURN VALUES

Upon successful completion, the `devid_get()`, `devid_get_minor_name()`, and `devid_deviceid_to_nmlist()` functions return 0. Otherwise, they return -1 and set `errno` to indicate the error.

The `devid_compare()` function returns the following values:

- <=-1 The device ID pointed to by `devid1` is less than the device ID pointed to by `devid2`.
- 0 The device ID pointed to by `devid1` is equal to the device ID pointed to by `devid2`.
- >=1 The device ID pointed to by `devid1` is greater than the device ID pointed to by `devid2`.

The `devid_sizeof()` function returns the size in number of bytes allocated for the `devid`.

EXAMPLES

EXAMPLE 1 Using `devid_get()` and `devid_get_minor_name()`

The following example shows the proper use of `devid_get()` and `devid_get_minor_name()` to free the space allocated for `devid` and `minor_name`.

```
int fd;
ddi_devid_t   devid;
char         *minor_name;
if ((fd = open("/dev/dsk/c0t3d0s0", O_RDONLY|O_NDELAY)) < 0) {
    ...
}
if (devid_get(fd, &devid) != 0) {
    ...
}
if (devid_get_minor_name(fd, &minor_name) != 0) {
    ...
}
< process devid and minor_name >
devid_free(devid);
free(minor_name);
```


EXAMPLE 2 Using `devid_deviceid_to_nmlist()` and `devid_free_nmlist()`

The following example shows the proper use of `devid_deviceid_to_nmlist()` and `devid_free_nmlist()`

```

:
  devid_nmlist_t *list = NULL;
  int    err;
  err = devid_deviceid_to_nmlist("/dev/rdsk", devid, minor_name, &list);
  if (err)
    return (err);
  < loop through list and process device names and device numbers >
  devid_free_nmlist(list);

```

FILES

| | |
|-------------------------------------|--|
| <code>/usr/lib/libdevid.so.1</code> | location of the device ID library interfaces |
| <code>/usr/lib/libdevid.so</code> | symlink to <code>/usr/lib/libdevid.so.1</code> |

ATTRIBUTES See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT Level | MT-Safe |

SEE ALSO `libdevid(3LIB)`, `attributes(5)`, `ddi_devid_compare(9F)`, `ddi_devid_free(9F)`, `ddi_devid_init(9F)`, `ddi_devid_register(9F)`, `ddi_devid_sizeof(9F)`, `ddi_devid_unregister(9F)`, `ddi_devid_valid(9F)`

| | |
|----------------------------|---|
| NAME | di_binding_name, di_bus_addr, di_compatible_names, di_devid, di_driver_name, di_driver_ops, di_instance, di_nodeid, di_node_name – return libdevinfo node information |
| SYNOPSIS | <pre>#include <libdevinfo.h> char *di_binding_name(di_node_t node); char *di_bus_addr(di_node_t node); int di_compatible_names(di_node_t node, char **names); ddi_devid_t di_devid(di_node_t node); char *di_driver_name(di_node_t node); uint_t di_driver_ops(di_node_t node); int di_instance(di_node_t node); int di_nodeid(di_node_t node); char *di_node_name(di_node_t node);</pre> |
| DESCRIPTION | These interfaces are used to extract information associated with a device node. |
| PARAMETERS | |
| All Interfaces | <i>node</i> A handle to a device node. |
| di_compatible_names | <i>names</i> The address of a pointer. |
| RETURN VALUES | |
| di_binding_name() | di_binding_name() returns a pointer to the binding name. The binding name is the name used by the system to select a driver for the device. |
| di_bus_addr() | di_bus_addr() returns a pointer to a null-terminated string containing the assigned bus address for the device. NULL is returned if a bus address has not been assigned to the device. A zero-length string may be returned and is considered a valid bus address. |
| di_compatible_names | <p>The return value of di_compatible_names() is the number of compatible names. <i>names</i> is updated to point to a buffer contained within the snapshot. The buffer contains a concatenation of null-terminated strings, for example:</p> <pre><name1>X<name2>X...<namen>X</pre> <p>See the discussion of generic names in <i>Writing Device Drivers</i> for a description of how compatible names are used by Solaris to achieve driver binding for the node.</p> |

di_devid() | **di_devid()** returns the device ID for *node*, if it is registered. Otherwise, a null pointer is returned. Interfaces in the `libdevid(3LIB)` library may be used to manipulate the handle to the device id.

di_driver_name() | **di_driver_name()** returns the name of the driver bound to the *node*. A null pointer is returned if *node* is not bound to any driver.

di_driver_ops() | **di_driver_ops()** returns a bit array of device driver entry points that are supported by the driver bound to this *node*. Possible bit fields supported by the driver are `DI_CB_OPS`, `DI_BUS_OPS`, `DI_STREAM_OPS`.

di_instance() | **di_instance()** returns the instance number of the device. A value of -1 indicates an instance number has not been assigned to the device by the system.

di_nodeid() | **di_nodeid()** returns the type of device, which may be one of the following possible values: `DI_PSEUDO_NODEID`, `DI_PROM_NODEID`, and `DI_SID_NODEID`. Devices of type `DI_PROM_NODEID` may have additional properties that are defined by the PROM. See `di_prom_prop_data(3DEVINFO)` and `di_prom_prop_lookup_bytes(3DEVINFO)`.

di_node_name() | **di_node_name()** returns a pointer to a null-terminated string containing the node name.

EXAMPLES

See `di_init(3DEVINFO)` for an example showing typical use of these interfaces.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

`di_init(3DEVINFO)`, `di_prom_init(3DEVINFO)`,
`di_prom_prop_data(3DEVINFO)`,
`di_prom_prop_lookup_bytes(3DEVINFO)`,
`libdevinfo(3DEVINFO)`, `libdevid(3LIB)`, `attributes(5)`

Writing Device Drivers

| | |
|--------------------|--|
| NAME | di_child_node, di_parent_node, di_sibling_node, di_drv_first_node, di_drv_next_node – libdevinfo node traversal functions |
| SYNOPSIS | <pre>#include <libdevinfo.h> di_node_t di_child_node(di_node_t node); di_node_t di_parent_node(di_node_t node); di_node_t di_sibling_node(di_node_t node); di_node_t di_drv_first_node(const char *drv_name, di_node_t root); di_node_t di_drv_next_node(di_node_t node);</pre> |
| DESCRIPTION | <p>The kernel device configuration data may be viewed in two ways, either as a tree of device configuration nodes or as a list of nodes associated with each driver. In the tree view, each node may contain references to its parent, the next sibling in a list of siblings, and the first child of a list of children. In the per-driver view, each node contains a reference to the next node associated with the same driver.</p> <p>Both views are captured in the snapshot, and the interfaces are provided for node access.</p> <p>di_child_node() obtains a handle to the first child of <i>node</i>. DI_NODE_NIL is returned and <i>errno</i> is set to ENXIO or ENOTSUP , if no child node exists in the snapshot.</p> <p>di_parent_node() obtains a handle to the parent node of <i>node</i>. DI_NODE_NIL is returned and <i>errno</i> is set to ENXIO or ENOTSUP , if no parent node exists in the snapshot.</p> <p>di_sibling_node() obtains a handle to the next sibling node of <i>node</i>. A DI_NODE_NIL is returned and <i>errno</i> is set to ENXIO or ENOTSUP , if no next sibling node exists in the snapshot.</p> <p>di_drv_first_node() obtains a handle to the first node associated with the driver specified by <i>drv_name</i> . If there is no such driver, DI_NODE_NIL is returned with <i>errno</i> is set to EINVAL . If the driver exists, but there is no node associated with this driver, DI_NODE_NIL is returned and <i>errno</i> is set to ENXIO or ENOTSUP .</p> <p>di_drv_next_node() returns a handle to the next node bound to the same driver. DI_NODE_NIL is returned if no more nodes exist.</p> |
| PARAMETERS | <p>The following parameter descriptions apply to di_child_node() , di_drv_next_node() , di_parent_node() , and di_sibling_node() :</p> <p><i>node</i> A handle to any node in the snapshot.</p> <p>The following parameter descriptions apply to di_drv_first_node() :</p> |

drv_name The name of the driver of interest.

root The handle of the root node for the snapshot returned by `di_init()`.

RETURN VALUES

Upon successful completion, a handle is returned. Otherwise, `DI_NODE_NIL` is returned and `errno` is set to indicate the error.

ERRORS

These functions set `errno` as listed for the following conditions:

`EINVAL` The argument is invalid.

`ENXIO` The requested node does not exist.

`ENOTSUP` The node was not found in the snapshot, but it may exist in the kernel. This error may occur if the snapshot contains a partial device tree.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

`libdevinfo(3DEVINFO)`, `attributes(5)`

Writing Device Drivers

NAME | di_devfs_path, di_devfs_path_free – generate and free physical path names

SYNOPSIS | #include <libdevinfo.h>
char *di_devfs_path(di_node_t node);
void di_devfs_path_free(char *path_buf);

DESCRIPTION | di_devfs_path() generates the physical path of the device *node*. The caller is responsible for freeing the memory allocated to store the physical path by calling di_devfs_path_free().
di_devfs_path_free() frees memory that was allocated by di_devfs_path().

PARAMETERS

di_devfs_path() | *node* | Handle to a device node in the snapshot.

di_devfs_path_free() | *path_buf* | Pointer returned by di_devfs_path().

RETURN VALUES

di_devfs_path() | Pointer to the string containing the physical path of *node*.

ERRORS

EINVAL | *node* is not a valid handle.
di_devfs_path() also return any error code from malloc(3C).

ATTRIBUTES | See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO | malloc(3C), libdevinfo(3DEVINFO), attributes(5)
Writing Device Drivers

| | |
|----------------------|--|
| NAME | di_init, di_fini – create and destroy a snapshot of kernel device tree |
| SYNOPSIS | <pre>#include <libdevinfo.h> di_node_t di_init(const char *phys_path, uint_t flags); void di_fini(di_node_t root);</pre> |
| DESCRIPTION | <p>di_init() creates a snapshot of the kernel device tree and returns a handle of the <i>root</i> node. The caller specifies the contents of the snapshot by providing <i>flag</i> and <i>phys_path</i> .</p> <p>di_fini() destroys the snapshot of the kernel device tree and frees the associated memory. All handles associated with this snapshot become invalid after the call to di_fini() .</p> |
| PARAMETERS | |
| di_init() | <p><i>phys_path</i> Physical path of the <i>root</i> node of the snapshot. See di_devfs_path(3DEVINFO) .</p> <p><i>flags</i> Snapshot content specification. The possible values may be a bitwise OR of the following:</p> <p style="margin-left: 40px;">DINFOSUBTREE Include subtree.</p> <p style="margin-left: 40px;">DINFOPROP Include properties.</p> <p style="margin-left: 40px;">DINFOMINOR Include minor data.</p> <p style="margin-left: 40px;">DINFOCPYALL Include all of above.</p> <p>If <i>flags</i> is 0 , the snapshot contains only a single node without properties or minor nodes.</p> |
| di_fini() | <p><i>root</i> Handle obtained by calling di_init() .</p> |
| RETURN VALUES | |
| di_init() | <p>Upon success, a handle is returned. Otherwise, DI_NODE_NIL is returned and errno is set to indicate the error.</p> |
| ERRORS | <p>di_init() may set errno to any error code that may also be set by open(2) , ioctl(2) or mmap(2) . The most common error codes include:</p> <p>EACCESS Insufficient privilege for accessing device configuration data.</p> <p>ENXIO Either the device named by <i>phys_path</i> is not present in the system, or the devinfo(7D) driver is not installed properly.</p> <p>EINVAL Either <i>phys_path</i> is incorrectly formed or the <i>flags</i> argument is invalid.</p> |

EXAMPLES**EXAMPLE 1** Using the libdevinfo() Interfaces To Print All Device Tree Node Names

The following is an example using the libdevinfo() interfaces to print all device tree node names:

```

/*
 * Code to print all device tree node names
 */

#include <stdio.h>
#include <libdevinfo.h>

int
prt_nodename(di_node_t node, void *arg)
{
    printf("%s\n", di_node_name(node));
    return (DI_WALK_CONTINUE);
}

main()
{
    di_node_t root_node;
    if((root_node = di_init("/", DINFOSUBTREE)) == DI_NODE_NIL) {
        fprintf(stderr, "di_init() failed\n");
        exit(1);
    }
    di_walk_node(root_node, DI_WALK_CLDFIRST, NULL, prt_nodename);
    di_fini(root_node);
}

```

EXAMPLE 2 Using the libdevinfo() Interfaces To Print The Physical Path Of SCSI Disks

The following example uses the libdevinfo() interfaces to print the physical path of SCSI disks:

```

/*
 * Code to print physical path of scsi disks
 */

#include <stdio.h>
#include <libdevinfo.h>
#define DISK_DRIVER "sd" /* driver name */

void
prt_diskinfo(di_node_t node)
{
    int instance;
    char *phys_path;

    /*
     * If the device node exports no minor nodes,
     * there is no physical disk.
     */
}

```



```

    */
    if (di_minor_next(node, DI_MINOR_NIL) == DI_MINOR_NIL) {
        return;
    }

    instance = di_instance(node);
    phys_path = di_devfs_path(node);
    printf("%s%d: %s\n", DISK_DRIVER, instance, phys_path);
    di_devfs_path_free(phys_path);
}

void
walk_disknodes(di_node_t node)
{
    node = di_drv_first_node(DISK_DRIVER, node);
    while (node != DI_NODE_NIL) {
        prt_diskinfo(node);
        node = di_drv_next_node(node);
    }
}

main()
{
    di_node_t root_node;
    if ((root_node = di_init("/", DINFOCOPYALL)) == DI_NODE_NIL) {
        fprintf(stderr, "di_init() failed\n");
        exit(1);
    }
    walk_disknodes(root_node);
    di_fini(root_node);
}

```

ATTRIBUTES

See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

[open\(2\)](#), [ioctl\(2\)](#), [mmap\(2\)](#), [libdevinfo\(3DEVINFO\)](#), [attributes\(5\)](#)

Writing Device Drivers

| NAME | di_minor_devt, di_minor_name, di_minor_nodetype, di_minor_spectype – return libdevinfo minor node information | | | | | | |
|----------------------|--|----------------|-----------------|----------|------|---------------------|----------|
| SYNOPSIS | <pre>#include <libdevinfo.h> dev_t di_minor_devt(di_minor_t minor); char *di_minor_name(di_minor_t minor); char *di_minor_nodetype(di_minor_t minor); int di_minor_spectype(di_minor_t minor);</pre> | | | | | | |
| DESCRIPTION | These interfaces are used to return libdevinfo minor node information. | | | | | | |
| PARAMETERS | <i>minor</i> A handle to minor data node. | | | | | | |
| RETURN VALUES | | | | | | | |
| di_minor_name() | di_minor_name() returns the minor <i>name</i> . See ddi_create_minor_node(9F) for a description of the <i>name</i> parameter. | | | | | | |
| di_minor_devt() | The function di_minor_devt() returns the dev_t value of the minor node that is specified by SYS V ABI. See getmajor(9F), getminor(9F), and ddi_create_minor_node(9F) for more information. | | | | | | |
| di_minor_spectype() | di_minor_spectype() returns the <i>spec_type</i> of the file, either S_IFCHR or S_IFBLK. See ddi_create_minor_node(9F) for a description of the <i>spec_type</i> parameter. | | | | | | |
| di_minor_nodetype() | di_minor_nodetype() returns the minor <i>node_type</i> of the minor node. See ddi_create_minor_node(9F) for a description of the <i>node_type</i> parameter. | | | | | | |
| ERRORS | No error codes are returned. | | | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT Level</td> <td>Safe</td> </tr> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT Level | Safe | Interface Stability | Evolving |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | | | |
| MT Level | Safe | | | | | | |
| Interface Stability | Evolving | | | | | | |
| SEE ALSO | attributes(5), ddi_create_minor_node(9F), getmajor(9F), getminor(9F) <i>Writing Device Drivers</i> | | | | | | |

NAME | `di_minor_next` – libdevinfo minor node traversal functions

SYNOPSIS | `#include <libdevinfo.h>`
`di_minor_t di_minor_next(di_node_t node, di_minor_t minor);`

DESCRIPTION | `di_minor_next()` returns a handle to the next minor node for the device node `node`. If `minor` is `DI_MINOR_NIL`, a handle to the first minor node is returned.

PARAMETERS | `node` Device node with which the minor node is associated.
`minor` Handle to the current minor node or `DI_MINOR_NIL`.

RETURN VALUES | Upon successful completion, a handle to the next minor node is returned. Otherwise, `DI_MINOR_NIL` is returned and `errno` is set to indicate the error.

ERRORS | `errno` is set as listed for the following conditions:
`EINVAL` Invalid argument.
`ENXIO` End of minor node list.
`ENOTSUP` Minor node information is not available in snapshot.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO | `libdevinfo(3DEVINFO)`, `attributes(5)`
Writing Device Drivers

| NAME | di_prom_init, di_prom_fini – create and destroy a handle to the PROM device information | | | | | | |
|--|--|----------------|-----------------|----------|------|---------------------|----------|
| SYNOPSIS | <pre>#include <libdevinfo.h> di_prom_handle_t di_prom_init(); void di_prom_fini(di_prom_handle_t ph);</pre> | | | | | | |
| DESCRIPTION | For device nodes whose nodeid value is DI_PROM_NODEID (see di_nodeid(3DEVINFO)), additional properties may be retrieved from the PROM . di_prom_init() returns a handle that is used to retrieve such properties. This handle is passed to di_prom_prop_lookup_bytes(3DEVINFO) and di_prom_prop_next(3DEVINFO) . di_prom_fini() destroys the handle and all handles to PROM device information obtained from that handle. | | | | | | |
| PARAMETERS | <i>ph</i> Handle to prom returned by di_prom_init() . | | | | | | |
| RETURN VALUES di_prom_init() | Upon successful completion, a handle is returned. Otherwise, DI_PROM_HANDLE_NIL is returned and errno is set to indicate the error. | | | | | | |
| ERRORS | di_prom_init() sets errno to any error code that may also be set by openprom(7D) or malloc(3C) . | | | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | | | |
| | <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT Level</td> <td>Safe</td> </tr> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT Level | Safe | Interface Stability | Evolving |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | | | |
| MT Level | Safe | | | | | | |
| Interface Stability | Evolving | | | | | | |
| SEE ALSO | di_nodeid(3DEVINFO) , di_prom_prop_next(3DEVINFO) , di_prom_prop_lookup_bytes(3DEVINFO) libdevinfo(3DEVINFO) , malloc(3C) , attributes(5) , openprom(7D) | | | | | | |

| | |
|-----------------------|---|
| NAME | di_prom_prop_data, di_prom_prop_next, di_prom_prop_name – access PROM device information |
| SYNOPSIS | <pre>#include <libdevinfo.h> di_prom_prop_t di_prom_prop_next(di_prom_handle_t ph, di_node_t node, di_prom_prop_t prom_prop); char *di_prom_prop_name(di_prom_prop_t prom_prop); int di_prom_prop_data(di_prom_prop_t prom_prop, uchar_t **prop_data);</pre> |
| DESCRIPTION | <p>di_prom_prop_next() obtains a handle to the next property on the PROM property list associated with <i>node</i>. If <i>prom_prop</i> is DI_PROM_PROP_NIL, the first property associated with <i>node</i> is returned.</p> <p>di_prom_prop_name() returns the name of the <i>prom_prop</i> property.</p> <p>di_prom_prop_data() returns the value of the <i>prom_prop</i> property. The return value is a non-negative integer specifying the size in number of bytes in <i>prop_data</i>.</p> <p>All memory allocated by these functions is managed by the library and must not be freed by the caller.</p> |
| PARAMETERS | |
| All Interfaces | <i>prom_prop</i> Handle to a PROM property. |
| di_prom_prop_data() | <i>prop_data</i> Address of a pointer. |
| di_prom_prop_next() | <i>ph</i> PROM handle |
| | <i>node</i> Handle to a device node in the snapshot of kernel device tree. |
| RETURN VALUES | |
| di_prom_prop_data() | di_prom_prop_data() returns the number of bytes in <i>prop_data</i> and <i>prop_data</i> is updated to point to a byte array containing the property value. If 0 is returned, the property is a boolean property, and the existence of this property indicates the value is true. |
| di_prom_prop_name() | di_prom_prop_name() returns a pointer to a string that contains the name of <i>prom_prop</i> . |
| di_prom_prop_next() | di_prom_prop_next() returns a handle to the next PROM property. DI_PROM_PROP_NIL is returned if no additional properties exist. |

ERRORS

See `openprom(7D)` for a description of possible errors.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

`attributes(5)`, `openprom(7D)`

Writing Device Drivers

| | |
|----------------------|--|
| NAME | di_prom_prop_lookup_bytes, di_prom_prop_lookup_ints, di_prom_prop_lookup_strings – search for a PROM property |
| SYNOPSIS | <pre>#include <libdevinfo.h> int di_prom_prop_lookup_bytes(di_prom_handle_t ph, di_node_t node, const char *prop_name, uchar_t **prop_data); int di_prom_prop_lookup_ints(di_prom_handle_t ph, di_node_t node, const char *prop_name, int **prop_data); int di_prom_prop_lookup_strings(di_prom_handle_t ph, di_node_t node, const char *prop_name, char **prop_data);</pre> |
| DESCRIPTION | These functions are used for returning the value of a known PROM property name and value type. These functions will update the <i>prop_data</i> pointer to reference memory that contains the property value. All memory allocated by these functions is managed by the library and must not be freed by the caller. |
| PARAMETERS | <p>The following parameter descriptions apply to all interfaces:</p> <p><i>node</i> Handle to device node in snapshot created by di_init(3DEVINFO) .</p> <p><i>ph</i> Handle returned by di_prom_init(3DEVINFO) .</p> <p><i>prop_name</i> Name of the property being searched.</p> <p>The following parameter description applies to di_prom_prop_lookup_bytes() only:</p> <p><i>prop_data</i> The address of a pointer to an array of unsigned characters.</p> <p>The following parameter description applies to di_prom_prop_lookup_ints() only:</p> <p><i>prop_data</i> The address of a pointer to an integer.</p> <p>The following parameter description applies to di_prom_prop_lookup_strings() only:</p> <p><i>prop_data</i> The address of pointer to a buffer.</p> |
| RETURN VALUES | If the property is found, the number of entries in <i>prop_data</i> is returned. If the property is a boolean type, 0 is returned, and the existence of this property indicates the value is true. Otherwise, -1 is returned with <i>errno</i> set to indicate the error condition. |

For `di_prom_prop_lookup_bytes()`, the number of entries is the number of unsigned characters contained in the buffer pointed to by `prop_data`.

For `di_prom_prop_lookup_ints()`, the number of entries is the number of integers contained in the buffer pointed to by `prop_data`.

For `di_prom_prop_lookup_strings()`, the number of entries is the number of null-terminated strings contained in the buffer. The strings are stored in a concatenated format in the buffer.

ERRORS

These functions set `errno` as listed for the following conditions:

- `EINVAL` Invalid argument.
- `ENXIO` The property does not exist.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

`di_init(3DEVINFO)`, `di_prom_prop_next(3DEVINFO)`, `libdevinfo(3DEVINFO)`, `attributes(5)`, `openprom(7D)`

Writing Device Drivers

| | | | | | | | | | | | | | |
|----------------------|---|----------------------|---|------------------|--|---------------------|---|-------------------|---|----------------------|---|-------------------|---|
| NAME | di_prop_bytes, di_prop_devt, di_prop_ints, di_prop_name, di_prop_strings, di_prop_type – access property values and attributes | | | | | | | | | | | | |
| SYNOPSIS | <pre>#include <libdevinfo.h> int di_prop_bytes(di_prop_t prop, uchar_t **prop_data); dev_t di_prop_devt(di_prop_t prop); int di_prop_ints(di_prop_t prop, int **prop_data); char *di_prop_name(di_prop_t prop); int di_prop_strings(di_prop_t prop, char **prop_data); int di_prop_type(di_prop_t prop);</pre> | | | | | | | | | | | | |
| DESCRIPTION | <p>These interfaces are used to access information associated with property values and attributes.</p> <p>All memory allocated by these functions is managed by the library and must not be freed by the caller.</p> <p>di_prop_name() returns the name of the property.</p> <p>di_prop_type() returns the type of the property. The type determines the appropriate interface to access property values. The following is a list of possible types:</p> <table border="0"> <tr> <td style="padding-right: 20px;">DI_PROP_TYPE_BOOLEAN</td> <td>There is no interface to call since there is no property data associated with boolean properties. The existence of the property defines a TRUE value.</td> </tr> <tr> <td>DI_PROP_TYPE_INT</td> <td>Use di_prop_ints() to access property data.</td> </tr> <tr> <td>DI_PROP_TYPE_STRING</td> <td>Use di_prop_strings() to access property data.</td> </tr> <tr> <td>DI_PROP_TYPE_BYTE</td> <td>Use di_prop_bytes() to access property data.</td> </tr> <tr> <td>DI_PROP_TYPE_UNKNOWN</td> <td>Use di_prop_bytes() to access property data. Since the type of property is unknown, the caller is responsible for interpreting the contents of the data.</td> </tr> <tr> <td>DI_PROP_UNDEFINED</td> <td>The property has been undefined by the driver. No property data is available.</td> </tr> </table> | DI_PROP_TYPE_BOOLEAN | There is no interface to call since there is no property data associated with boolean properties. The existence of the property defines a TRUE value. | DI_PROP_TYPE_INT | Use di_prop_ints() to access property data. | DI_PROP_TYPE_STRING | Use di_prop_strings() to access property data. | DI_PROP_TYPE_BYTE | Use di_prop_bytes() to access property data. | DI_PROP_TYPE_UNKNOWN | Use di_prop_bytes() to access property data. Since the type of property is unknown, the caller is responsible for interpreting the contents of the data. | DI_PROP_UNDEFINED | The property has been undefined by the driver. No property data is available. |
| DI_PROP_TYPE_BOOLEAN | There is no interface to call since there is no property data associated with boolean properties. The existence of the property defines a TRUE value. | | | | | | | | | | | | |
| DI_PROP_TYPE_INT | Use di_prop_ints() to access property data. | | | | | | | | | | | | |
| DI_PROP_TYPE_STRING | Use di_prop_strings() to access property data. | | | | | | | | | | | | |
| DI_PROP_TYPE_BYTE | Use di_prop_bytes() to access property data. | | | | | | | | | | | | |
| DI_PROP_TYPE_UNKNOWN | Use di_prop_bytes() to access property data. Since the type of property is unknown, the caller is responsible for interpreting the contents of the data. | | | | | | | | | | | | |
| DI_PROP_UNDEFINED | The property has been undefined by the driver. No property data is available. | | | | | | | | | | | | |

`di_prop_devt()` returns the `dev_t` with which this property is associated. If the value is `DDI_DEV_T_NONE`, the property is not defined for a specific minor node.

`di_prop_bytes()` returns the property data as a series of unsigned characters.

`di_prop_ints()` returns the property data as a series of integers.

`di_prop_strings()` returns the property data as a concatenation of null-terminated strings.

PARAMETERS

| | | |
|------------------------|------------------|--|
| All Interfaces | <i>prop</i> | Handle to a property returned by <code>di_prop_next(3DEVINFO)</code> . |
| di_prop_bytes | <i>prop_data</i> | The address of a pointer to an unsigned character. |
| di_prop_ints | <i>prop_data</i> | The address of a pointer to an integer. |
| di_prop_strings | <i>prop_data</i> | The address of pointer to a character. |

RETURN VALUES

di_prop_bytes,
di_prop_ints,
di_prop_strings

Upon successful completion, these interfaces return a non-negative value, indicating the number of entries in the property value buffer. See `di_prop_lookup_bytes(3DEVINFO)` for a description of the return values. Otherwise, `-1` is returned and `errno` is set to indicate the error condition.

di_prop_devt `di_prop_devt()` returns the `dev_t` value associated with the property.

di_prop_name `di_prop_name()` returns a pointer to a string containing the name of the property.

di_prop_type `di_prop_type()` may return one of various types described in the DESCRIPTION section.

ERRORS

These functions set `errno` as listed for the following conditions:

`EINVAL` Invalid argument. For example, the property type does not match the interface.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

`di_prom_prop_lookup_bytes(3DEVINFO)`, `di_prop_next(3DEVINFO)`,
`libdevinfo(3DEVINFO)`, `attributes(5)`

Writing Device Drivers

| | | |
|-------------------------------|---|---|
| NAME | di_prop_lookup_bytes, di_prop_lookup_ints, di_prop_lookup_strings – search for a property | |
| SYNOPSIS | <pre>#include <libdevinfo.h> int di_prop_lookup_bytes(dev_t dev, di_node_t node, const char *prop_name, uchar_t **prop_data); int di_prop_lookup_ints(dev_t dev, di_node_t node, const char *prop_name, int **prop_data); int di_prop_lookup_strings(dev_t dev, di_node_t node, const char *prop_name, char **prop_data);</pre> | |
| DESCRIPTION | <p>These functions are used for returning the value of a known property name type and dev_t value.</p> <p>All memory allocated by these functions is managed by the library and must not be freed by the caller.</p> | |
| PARAMETERS | | |
| All Interfaces | <i>dev</i> | dev_t of minor node with which the property is associated. DDI_DEV_T_ANY is a wild card that matches all dev_t 's, including DDI_DEV_T_NONE . |
| | <i>node</i> | Handle to the device node with which the property is associated. |
| | <i>prop_name</i> | Name of the property for which to search. |
| di_prop_lookup_bytes | <i>prop_data</i> | Address to a pointer to an array of unsigned characters containing the property data. |
| di_prop_lookup_ints | <i>prop_data</i> | Address to a pointer to an array of integers containing the property data. |
| di_prop_lookup_strings | <i>prop_data</i> | Address to a pointer to a buffer containing a concatenation of null-terminated strings containing the property data. |
| RETURN VALUES | If the property is found, the number of entries in <i>prop_data</i> is returned. If the property is a boolean type, 0 is returned, and the existence of this property indicates the value is true. Otherwise, -1 is returned with <i>errno</i> set to indicate the error condition. | |
| ERRORS | These functions set <i>errno</i> as listed for the following conditions: | |

EINVAL Invalid argument.
 ENOTSUP The snapshot contains no property information.
 ENXIO The property does not exist; try
 di_prom_prop_lookup_*().

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

di_init(3DEVINFO), di_prom_prop_lookup_bytes(3DEVINFO),
 libdevinfo(3DEVINFO), attributes(5)
Writing Device Drivers

| NAME | di_prop_next – libdevinfo property traversal function | | | | | | |
|----------------------|---|----------------|-----------------|----------|------|---------------------|----------|
| SYNOPSIS | #include <libdevinfo.h> di_prop_t di_prop_next(di_node_t node, di_prop_t prop); | | | | | | |
| DESCRIPTION | The function di_prop_next () returns a handle to the next property on the property list. If prop is DI_PROP_NIL, the handle to the first property is returned. | | | | | | |
| PARAMETERS | node Handle to a device node. prop Handle to a property. | | | | | | |
| RETURN VALUES | Upon successful completion, di_prop_next () returns a handle. Otherwise DI_PROP_NIL is returned, and errno is set to indicate the error condition. | | | | | | |
| ERRORS | The di_prop_next () functions sets errno as listed for the following conditions: EINVAL Invalid argument. ENOTSUP Snapshot does not contain property information. ENXIO There are no more properties. | | | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT Level</td> <td>Safe</td> </tr> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT Level | Safe | Interface Stability | Evolving |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | | | |
| MT Level | Safe | | | | | | |
| Interface Stability | Evolving | | | | | | |
| SEE ALSO | di_init(3DEVINFO), libdevinfo(3DEVINFO), attributes(5) <i>Writing Device Drivers</i> | | | | | | |

| NAME | DisconnectToServer – disconnect from a DMI service provider | | | | |
|----------------------|--|----------------|-----------------|----------|------|
| SYNOPSIS | <pre>cc [flag ...] file ... -ldmici -ldmimi [library ...] #include <dmi/api.hh> bool_t DisconnectToServer(DmiRpcHandle *dmi_rpc_handle);</pre> | | | | |
| DESCRIPTION | The <code>DisconnectToServer()</code> function disconnects a management application or a component instrumentation from a DMI service provider. | | | | |
| RETURN VALUES | The <code>ConnectToServer()</code> function returns TRUE if successful, otherwise FALSE. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-level</td> <td>Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-level | Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-level | Safe | | | | |
| SEE ALSO | <code>ConnectToServer(3DMI)</code> , <code>attributes(5)</code> | | | | |

| | | | | | | | | | |
|-----------------------|---|------------------|---|-----------------------|--|-------------|-------------------------------------|------------|--|
| NAME | di_walk_minor – traverse libdevinfo minor nodes | | | | | | | | |
| SYNOPSIS | <pre>#include <libdevinfo.h> int di_walk_minor(di_node_t root, const char *minor_nodetype, uint_t flag, void *arg, int (*minor_callback)(di_node_t node, di_minor_t minor, void *arg));</pre> | | | | | | | | |
| DESCRIPTION | di_walk_minor() visits all minor nodes attached to device nodes in a subtree rooted at root. For each minor node that matches minor_nodetype, the caller-supplied function minor_callback() is invoked. The walk terminates immediately when minor_callback() returns DI_WALK_TERMINATE. | | | | | | | | |
| PARAMETERS | | | | | | | | | |
| di_walk_minor | <table border="0"> <tr> <td style="padding-right: 20px;"><i>root</i></td> <td>Root of subtree to visit.</td> </tr> <tr> <td><i>minor_nodetype</i></td> <td>A character string specifying the minor data type, which may be one of the types defined by the Solaris DDI framework, for example, DDI_NT_BLOCK. NULL matches all <i>minor_node</i> types. See ddi_create_minor_node(9F).</td> </tr> <tr> <td><i>flag</i></td> <td>Specify 0. Reserved for future use.</td> </tr> <tr> <td><i>arg</i></td> <td>Pointer to caller- specific user data.</td> </tr> </table> | <i>root</i> | Root of subtree to visit. | <i>minor_nodetype</i> | A character string specifying the minor data type, which may be one of the types defined by the Solaris DDI framework, for example, DDI_NT_BLOCK. NULL matches all <i>minor_node</i> types. See ddi_create_minor_node(9F). | <i>flag</i> | Specify 0. Reserved for future use. | <i>arg</i> | Pointer to caller- specific user data. |
| <i>root</i> | Root of subtree to visit. | | | | | | | | |
| <i>minor_nodetype</i> | A character string specifying the minor data type, which may be one of the types defined by the Solaris DDI framework, for example, DDI_NT_BLOCK. NULL matches all <i>minor_node</i> types. See ddi_create_minor_node(9F). | | | | | | | | |
| <i>flag</i> | Specify 0. Reserved for future use. | | | | | | | | |
| <i>arg</i> | Pointer to caller- specific user data. | | | | | | | | |
| minor_callback | <table border="0"> <tr> <td style="padding-right: 20px;"><i>node</i></td> <td>The device node with which to the minor node is associated.</td> </tr> <tr> <td><i>minor</i></td> <td>The minor node visited.</td> </tr> <tr> <td><i>arg</i></td> <td>Pointer to caller-specific data.</td> </tr> </table> | <i>node</i> | The device node with which to the minor node is associated. | <i>minor</i> | The minor node visited. | <i>arg</i> | Pointer to caller-specific data. | | |
| <i>node</i> | The device node with which to the minor node is associated. | | | | | | | | |
| <i>minor</i> | The minor node visited. | | | | | | | | |
| <i>arg</i> | Pointer to caller-specific data. | | | | | | | | |
| RETURN VALUES | | | | | | | | | |
| di_walk_minor | Upon successful completion, 0 is returned. Otherwise, -1 is returned and errno is set to indicate the error. | | | | | | | | |
| minor_callback | <p>The allowed return values are:</p> <table border="0"> <tr> <td style="padding-right: 20px;">DI_WALK_CONTINUE</td> <td>Continue to visit subsequent minor data nodes.</td> </tr> <tr> <td>DI_WALK_TERMINATE</td> <td>Terminate the walk immediately.</td> </tr> </table> | DI_WALK_CONTINUE | Continue to visit subsequent minor data nodes. | DI_WALK_TERMINATE | Terminate the walk immediately. | | | | |
| DI_WALK_CONTINUE | Continue to visit subsequent minor data nodes. | | | | | | | | |
| DI_WALK_TERMINATE | Terminate the walk immediately. | | | | | | | | |
| ERRORS | | | | | | | | | |
| di_walk_minor | EINVAL Invalid argument. | | | | | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | | | | | |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

di_minor_nodetype(3DEVINFO), libdevinfo(3DEVINFO),
attributes(5) ddi_create_minor_node(9F)

Writing Device Drivers

| | |
|----------------------|--|
| NAME | di_walk_node – traverse libdevinfo device nodes |
| SYNOPSIS | <pre>#include <libdevinfo.h> int di_walk_node(di_node_t root, uint_t flag, void *arg, int (*node_callback)(di_node_t node, void *arg));</pre> |
| DESCRIPTION | di_walk_node() visits all nodes in the subtree rooted at <i>root</i> . For each node found, the caller-supplied function <i>node_callback</i> () is invoked. The return value of <i>node_callback</i> () specifies subsequent walking behavior. |
| PARAMETERS | |
| di_walk_node | <p><i>root</i> Handle to the root node of the subtree to visit.</p> <p><i>flag</i> Specifies walking order, either <code>DI_WALK_CLDFIRST</code> (depth first) or <code>DI_WALK_SIBFIRST</code> (breadth first). <code>DI_WALK_CLDFIRST</code> is the default.</p> <p><i>arg</i> Pointer to caller-specific data.</p> |
| node_callback | <p><i>node</i> The node being visited.</p> <p><i>arg</i> Pointer to caller-specific data.</p> |
| RETURN VALUES | |
| di_walk_node | 0 is returned upon success. Otherwise, -1 is returned, and <code>errno</code> is set to indicate the error. |
| node_callback | <p>The allowed return values are:</p> <p><code>DI_WALK_CONTINUE</code> Continue walking.</p> <p><code>DI_WALK_PRUNESIB</code> Continue walking, but skip siblings and their child nodes.</p> <p><code>DI_WALK_PRUNECHILD</code> Continue walking, but skip subtree rooted at current node .</p> <p><code>DI_WALK_TERMINATE</code> Terminate the walk immediately.</p> |
| ERRORS | |
| di_walk_node | <code>EINVAL</code> Invalid argument. |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

`di_init(3DEVINFO)`, `libdevinfo(3DEVINFO)`, `attributes(5)`

Writing Device Drivers

| | |
|--------------------|---|
| NAME | DmiAddComponent, DmiAddGroup, DmiAddLanguage, DmiDeleteComponent, DmiDeleteGroup, DmiDeleteLanguage – Management Interface database administration functions |
| SYNOPSIS | <pre>cc [flag ...] file ... -ldmimi -ldmi -lnsl -lrwtool [library ...] #include <dmi/server.h> #include <dmi/miapi.h> bool_t DmiAddComponent(DmiAddComponentIN argin, DmiAddComponentOUT *result, DmiRpcHandle *dmi_rpc_handle); bool_t DmiAddGroup(DmiAddGroupIN argin, DmiAddGroupOUT *result, DmiRpcHandle *dmi_rpc_handle); bool_t DmiAddLanguage(DmiAddLanguageIN argin, DmiAddLanguageOUT*result, DmiRpcHandle *dmi_rpc_handle); bool_t DmiDeleteComponent(DmiDeleteComponentIN argin, DmiDeleteComponentOUT *result, DmiRpcHandle *dmi_rpc_handle); bool_t DmiDeleteGroup(DmiDeleteGroupIN argin, DmiDeleteGroupOUT *result, DmiRpcHandle *dmi_rpc_handle); bool_t DmiDeleteLanguage(DmiDeleteLanguageIN argin, DmiDeleteLanguageOUT *result, DmiRpcHandle *dmi_rpc_handle);</pre> |
| DESCRIPTION | <p>The database administration functions add a new component to the database or add a new language mapping for an existing component. You may also remove an existing component, remove a specific language mapping, or remove a group from a component.</p> <p>The <code>DmiAddComponent ()</code> function adds a new component to the DMI database. It takes the name of a file, or the address of memory block containing MIF data, checks the data for adherence to the DMI MIF grammar, and installs the MIF in the database. The procedure returns a unique component ID for the newly installed component. The <i>argin</i> parameter is an instance of a <code>DmiAddComponentIN</code> structure containing the following members:</p> <pre>DmiHandle_t handle; /* an open session handle */ DmiFileDataList_t *fileData; /* MIF data for component */</pre> <p>The <i>result</i> parameter is a pointer to a <code>DmiAddComponentOUT</code> structure containing the following members:</p> <pre>DmiErrorStatus_t error_status; DmiId_t compId; /* SP-allocated component ID */ DmiStringList_t *errors; /* installation error messages */</pre> |

The `DmiAddLanguage()` function adds a new language mapping for an existing component in the database. It takes the name of a file, or the address of memory block containing translated MIF data, checks the data for adherence to the DMI MIF grammar, and installs the language MIF in the database. The *argIn* parameter is an instance of a `DmiAddLanguageIN` structure containing the following members:

```
DmiHandle_t      handle;          /* an open session handle */
DmiFileDataList_t *fileData;     /* language mapping file */
DmiId_t         compId;          /* component to access */
```

The *result* parameter is a pointer to a `DmiAddLanguageOUT` structure containing the following members:

```
DmiErrorStatus_t error_status;
DmiStringList_t  *errors;        /* installation error messages */
```

The `DmiAddGroup()` function adds a new group to an existing component in the database. It takes the name of a file, or the address of memory block containing the group's MIF data, checks the data for adherence to the DMI MIF grammar, and installs the group MIF in the database. The *argIn* parameter is an instance of a `DmiAddGroupIN` structure containing the following members:

```
DmiHandle_t      handle;          /* an open session handle */
DmiFileDataList_t *fileData;     /* MIF file data for group */
DmiId_t         compId;          /* component to access */
```

The *result* parameter is a pointer to a `DmiAddGroupOUT` structure containing the following members:

```
DmiErrorStatus_t error_status;
DmiId_t         groupId;         /* SP-allocated group ID */
DmiStringList_t *errors;        /* installation error messages */
```

The `DmiDeleteComponent()` function removes an existing component from the database. The *argIn* parameter is an instance of a `DmiDeleteComponentIN` structure containing the following members:

```
DmiHandle_t      handle;          /* an open session handle */
DmiId_t         compId;          /* component to delete */
```

The *result* parameter is a pointer to a `DmiDeleteComponentOUT` structure containing the following members:

```
DmiErrorStatus_t error_status;
```

The `DmiDeleteLanguage()` function removes a specific language mapping for a component. You specify the language string and component ID. The *argIn* parameter is an instance of a `DmiDeleteLanguageIN` structure containing the following members:

```
DmiHandle_t      handle;          /* an open session handle */
DmiString_t     *language;       /* language to delete */
DmiId_t         compId;          /* component to access */
```

The *result* parameter is a pointer to a `DmiDeleteLanguageOUT` structure containing the following members:

```
DmiErrorStatus_t  error_status;
```

The `DmiDeleteGroup()` function removes a group from a component. The caller specifies the component and group IDs. The *argIn* parameter is an instance of a `DmiDeleteGroupIN` structure containing the following members:

```
DmiHandle_t      handle;          /* an open session handle */
DmiId_t         compId;          /* component containing group */
DmiId_t         groupId;         /* group to delete */
```

The *result* parameter is a pointer to a `DmiDeleteGroupOUT` structure containing the following members:

```
DmiErrorStatus_t  error_status;
```

RETURN VALUES

The `DmiAddComponent()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_FILE_ERROR
DMIERR_BAD_SCHEMA_DESCRIPTION_FILE
```

The `DmiAddGroup()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_INSUFFICIENT_PRIVILEGES
DMIERR_COMPONENT_NOT_FOUND
DMIERR_FILE_ERROR
DMIERR_BAD_SCHEMA_DESCRIPTION_FILE
```

The `DmiAddLanguage()` function returns the following possible values:

```

DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_COMPONENT_NOT_FOUND
DMIERR_FILE_ERROR
DMIERR_BAD_SCHEMA_DESCRIPTION_FILE
    
```

The `DmiDeleteComponent()` function returns the following possible values:

```

DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_INSUFFICIENT_PRIVILEGES
DMIERR_COMPONENT_NOT_FOUND
DMIERR_FILE_ERROR
    
```

The `DmiDeleteGroup()` function returns the following possible values:

```

DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_INSUFFICIENT_PRIVILEGES
DMIERR_COMPONENT_NOT_FOUND
DMIERR_FILE_ERROR
    
```

The `DmiDeleteLanguage()` function returns the following possible values:

```

DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_COMPONENT_NOT_FOUND
DMIERR_FILE_ERROR
    
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWsasdk |
| MT-level | Unsafe |

SEE ALSO

attributes(5)

NAME | DmiAddRow, DmiDeleteRow, DmiGetAttribute, DmiGetMultiple,
DmiSetAttribute, DmiSetMultiple – Management Interface operation functions

SYNOPSIS

```
cc [ flag ... ] file ... -ldmimi -ldmi -lnsl -lrwtool [ library ... ]
#include <server.h>
#include <miapi.h>
bool_t DmiAddRow(DmiAddRowIN argin, DmiAddRowOUT *result, DmiRpcHandle
*dmi_rpc_handle);

bool_t DmiDeleteRow(DmiDeleteRowIN argin, DmiDeleteRowOUT *result,
DmiRpcHandle *dmi_rpc_handle);

bool_t DmiGetAttribute(DmiGetAttributeIN argin, DmiGetAttributeOUT *result,
DmiRpcHandle *dmi_rpc_handle);

bool_t DmiGetMultiple(DmiGetMultipleIN argin, DmiGetMultipleOUT *result,
DmiRpcHandle *dmi_rpc_handle);

bool_t DmiSetAttribute(DmiSetAttributeIN argin, DmiSetAttributeOUT *result,
DmiRpcHandle *dmi_rpc_handle);

bool_t DmiSetMultiple(DmiSetMultipleIN argin, DmiSetMultipleOUT *result,
DmiRpcHandle *dmi_rpc_handle);
```

DESCRIPTION

The operation functions provide a method for retrieving a single value from the Service Provider and for setting a single attribute value. In addition, you may also retrieve attribute values from the Service Provider. You may perform a set operation on an attribute or a list of attributes and add or delete a row from an existing table.

The `DmiAddRow()` function adds a row to an existing table. The `rowData` parameter contains the full data, including key attribute values, for a row. It is an error for the key list to specify an existing table row. The `argin` parameter is an instance of a `DmiAddRowIN` structure containing the following members:

```
DmiHandle_t          handle;          /* An open session handle */
DmiRowData_t        *rowData;        /* Attribute values to set */
```

The `result` parameter is a pointer to a `DmiAddRowOUT` structure containing the following members:

```
DmiErrorStatus_t    error_status;
```

The `DmiDeleteRow()` function removes a row from an existing table. The key list must specify valid keys for a table row. The `argin` parameter is an instance of a `DmiDeleteRowIN` structure containing the following members:

```
DmiHandle_t          handle;          /* An open session handle */
```

```
DmiRowData_t      *rowData;      /* Row to delete */
```

The *result* parameter is a pointer to a `DmiDeleteRowOUT` structure containing the following members:

```
DmiErrorStatus_t  error_status;
```

The `DmiGetAttribute()` function provides a simple method for retrieving a single attribute value from the Service Provider. The `compId`, `groupId`, `attribId`, and `keyList` identify the desired attribute. The resulting attribute value is returned in a newly allocated `DmiDataUnion` structure. The address of this structure is returned through the `value` parameter. The *argIn* parameter is an instance of a `DmiListComponentsIN` structure containing the following members:

```
DmiHandle_t      handle;        /* an open session handle */
DmiId_t          compId;        /* Component to access */
DmiId_t          groupId;       /* Group within component */
DmiId_t          attribId;      /* Attribute within a group */
DmiAttributeValues_t *keyList;  /* Keylist to specify a table row */
```

The *result* parameter is a pointer to a `DmiGetAttributeOUT` structure containing the following members:

```
DmiErrorStatus_t  error_status;
DmiDataUnion_t    *value;       /* Attribute value returned */
```

The `DmiGetMultiple()` function retrieves attribute values from the Service Provider. This procedure may get the value for an individual attribute, or for multiple attributes across groups, components, or rows of a table.

The `DmiSetAttribute()` function provides a simple method for setting a single attribute value. The `compId`, `groupId`, `attribId`, and `keyList` identify the desired attribute. The `setMode` parameter defines the procedure call as a Set, Reserve, or Release operation. The new attribute value is contained in the `DmiDataUnion` structure whose address is passed in the `value` parameter. The *argIn* parameter is an instance of a `DmiSetAttributeIN` structure containing the following members:

```
DmiHandle_t      handle;
DmiId_t          compId;
DmiId_t          groupId;
DmiId_t          attribId;
DmiAttributeValues_t *keyList;
DmiSetMode_t     setMode;
DmiDataUnion_t    *value;
```

The *result* parameter is a pointer to a `DmiSetAttributeOUT` structure containing the following members:

```
DmiErrorStatus_t    error_status;
```

The `DmiSetMultiple()` function performs a set operation on an attribute or list of attributes. Set operations include actually setting the value, testing and reserving the attribute for future setting, or releasing the set reserve. These variations on the set operation are specified by the parameter `setMode`. The *argIn* parameter is an instance of a `DmiSetMultipleIN` structure containing the following members:

```
DmiHandle_t        handle;          /* An open session handle */
DmiSetMode_t       setMode;         /* set, reserve, or release */
DmiMultiRowData_t *rowData;        /* Attribute values to set */
```

The *result* parameter is a pointer to a `DmiSetMultipleOUT` structure containing the following members:

```
DmiErrorStatus_t    error_status;
```

The `rowData` array describes the attributes to set, and contains the new attribute values. Each element of `rowData` specifies a component, group, key list (for table accesses), and attribute list to set. No data is returned from this function.

RETURN VALUES

The `DmiAddRow()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_VALUE_UNKNOWN
DMIERR_COMPONENT_NOT_FOUND
DMIERR_GROUP_NOT_FOUND
DMIERR_ILLEGAL_KEYS
DMIERR_DIRECT_INTERFACE_NOT_REGISTERED
DMIERR_UNKNOWN_CI_REGISTRY
DMIERR_VALUE_UNKNOWN
DMIERR_UNABLE_TO_ADD_ROW
```

The `DmiDeleteRow()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_ATTRIBUTE_NOT_FOUND
DMIERR_COMPONENT_NOT_FOUND
```

```

DMIERR_GROUP_NOT_FOUND
DMIERR_ILLEGAL_KEYS
DMIERR_ILLEGAL_TO_GET
DMIERR_DIRECT_INTERFACE_NOT_REGISTERED
DMIERR_ROW_NOT_FOUND
DMIERR_UNKNOWN_CI_REGISTRY
DMIERR_VALUE_UNKNOWN
DMIERR_UNABLE_TO_DELETE_ROW

```

The `DmiGetAttribute()` function returns the following possible values:

```

DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_ATTRIBUTE_NOT_FOUND
DMIERR_COMPONENT_NOT_FOUND
DMIERR_GROUP_NOT_FOUND
DMIERR_ILLEGAL_KEYS
DMIERR_ILLEGAL_TO_GET
DMIERR_DIRECT_INTERFACE_NOT_REGISTERED
DMIERR_ROW_NOT_FOUND
DMIERR_UNKNOWN_CI_REGISTRY
DMIERR_FILE_ERROR
DMIERR_VALUE_UNKNOWN

```

The `DmiGetMultiple()` function returns the following possible values:

```

DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_RPC_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_ATTRIBUTE_NOT_FOUND
DMIERR_COMPONENT_NOT_FOUND
DMIERR_GROUP_NOT_FOUND
DMIERR_ILLEGAL_KEYS
DMIERR_ILLEGAL_TO_GET
DMIERR_DIRECT_INTERFACE_NOT_REGISTERED
DMIERR_ROW_NOT_FOUND
DMIERR_UNKNOWN_CI_REGISTRY
DMIERR_FILE_ERROR
DMIERR_VALUE_UNKNOWN

```

The `DmiSetAttribute()` function returns the following possible values:

```

DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_ATTRIBUTE_NOT_FOUND

```

DMIERR_COMPONENT_NOT_FOUND
 DMIERR_GROUP_NOT_FOUND
 DMIERR_ILLEGAL_KEYS
 DMIERR_ILLEGAL_TO_GET
 DMIERR_DIRECT_INTERFACE_NOT_REGISTERED
 DMIERR_ROW_NOT_FOUND
 DMIERR_UNKNOWN_CI_REGISTRY
 DMIERR_FILE_ERROR
 DMIERR_VALUE_UNKNOWN

The `DmiSetMultiple()` function returns the following possible values:

DMIERR_NO_ERROR
 DMIERR_ILLEGAL_RPC_HANDLE
 DMIERR_OUT_OF_MEMORY
 DMIERR_ILLEGAL_PARAMETER
 DMIERR_SP_INACTIVE
 DMIERR_ATTRIBUTE_NOT_FOUND
 DMIERR_COMPONENT_NOT_FOUND
 DMIERR_GROUP_NOT_FOUND
 DMIERR_ILLEGAL_KEYS
 DMIERR_ILLEGAL_TO_SET
 DMIERR_DIRECT_INTERFACE_NOT_REGISTERED
 DMIERR_ROW_NOT_FOUND
 DMIERR_UNKNOWN_CI_REGISTRY
 DMIERR_FILE_ERROR
 DMIERR_VALUE_UNKNOWN

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | Unsafe |

SEE ALSO

`attributes(5)`

NAME | dmi_error – print error in string form

SYNOPSIS | `cc [flag ...] file ... -ldmi -lnsl -lrwtool [library ...]`
`#include <dmi/dmi_error.hh>`
`void dmi_error(DmiErrorStatus_t error_status);`

DESCRIPTION | For the given *error_status*, the `dmi_error()` function prints the corresponding error in string form. The function prints "unknown dmi errors" if *error_status* is invalid.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | MT-Safe |

SEE ALSO | `libdmi(3LIB)`, `attributes(5)`

| | |
|--------------------|--|
| NAME | DmiGetConfig, DmiGetVersion, DmiRegister, DmiSetConfig, DmiUnregister – Management Interface initialization functions |
| SYNOPSIS | <pre>cc [flag ...] file ... -ldmimi -ldmi -lnsl -lrwtool [library ...] #include <server.h> #include <miapi.h> bool_t DmiGetConfig(DmiGetConfigIN argin, DmiGetConfigOUT *result, DmiRpcHandle *dmi_rpc_handle); bool_t DmiGetVersion(DmiGetVersionIN argin, DmiGetVersionOUT *result, DmiRpcHandle *dmi_rpc_handle); bool_t DmiRegister(DmiRegisterIN argin, DmiRegisterOUT *result, DmiRpcHandle *dmi_rpc_handle); bool_t DmiSetConfig(DmiSetConfigIN argin, DmiSetConfigOUT *result, DmiRpcHandle *dmi_rpc_handle); bool_t DmiUnregister(DmiUnregisterIN argin, DmiUnregisterOUT *result, DmiRpcHandle *dmi_rpc_handle);</pre> |
| DESCRIPTION | <p>The Management Interface initialization functions enable you to register management applications to the Service Provider. You may also retrieve information about the Service Provider, get and set session configuration information for your session.</p> <p>The <code>DmiGetConfig()</code> function retrieves the per-session configuration information. The configuration information consists of a string describing the current language being used for the session. The <i>argin</i> parameter is an instance of a <code>DmiGetConfigIN</code> structure containing the following member:</p> <pre>DmiHandle_t handle; /* an open session handle */</pre> <p>The <i>result</i> parameter is a pointer to a <code>DmiGetConfigOUT</code> structure containing the following members:</p> <pre>DmiErrorStatus_t error_status; DmiString_t *language; /* current session language */</pre> <p>The <code>DmiGetVersion()</code> function retrieves information about the Service Provider. The management application uses the <code>DmiGetVersion()</code> procedure to determine the DMI specification level supported by the Service Provider. This procedure also returns the service provided description string, and may contain version information about the Service Provider implementation. The <i>argin</i> parameter is an instance of a <code>DmiGetVersionIN</code> structure containing the following member:</p> <pre>DmiHandle_t handle; /* an open session handle */</pre> |

The *result* parameter is a pointer to a `DmiGetVersionOUT` structure containing the following members:

```
DmiErrorStatus_t  error_status;
DmiString_t       *dmiSpecLevel; /* DMI specification version */
DmiString_t       *description;  /* OS specific DMI SP version */
DmiFileTypeList_t *fileTypes;    /* file types for MIF installation */
```

The `DmiRegister()` function provides the management application with a unique per-session handle. The Service Provider uses this procedure to initialize to an internal state for subsequent procedure calls made by the application. This procedure must be the first command executed by the management application. *argIn* is an instance of a `DmiRegisterIN` structure containing the following member:

```
DmiHandle_t       handle;          /* an open session handle */
```

The *result* parameter is a pointer to a `DmiRegisterOUT` structure containing the following members:

```
DmiErrorStatus_t  error_status;
DmiHandle_t       *handle;        /* an open session handle */
```

The `DmiSetConfig()` function sets the per-session configuration information. The configuration information consists of a string describing the language required by the management application. The *argIn* parameter is an instance of a `DmiSetConfigIN` structure containing the following member:

```
DmiHandle_t       handle;          /* an open session handle */
DmiString_t       *language;      /* current language required */
```

The *result* parameter is a pointer to a `DmiSetConfigOUT` structure containing the following member:

```
DmiErrorStatus_t  error_status;
```

The `DmiUnregister()` function is used by the Service Provider to perform end-of-session cleanup actions. On return from this function, the session handle is no longer valid. This function must be the last DMI command executed by the management application. The *argIn* parameter is an instance of a `DmiUnregisterIN` structure containing the following member:

```
DmiHandle_t       handle;          /* an open session handle */
```


The *result* parameter is a pointer to a `DmiUnregisterOUT` structure containing the following members:

```
DmiErrorStatus_t    error_status;
```

RETURN VALUES

The `DmiGetConfig()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
```

The `DmiGetVersion()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_SP_INACTIVE
```

The `DmiRegister()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_SP_INACTIVE
```

The `DmiSetConfig()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_ILLEGAL_TO_SET
```

The `DmiUnRegister()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | Unsafe |

SEE ALSO

attributes(5)

NAME | DmiListAttributes, DmiListClassNames, DmiListComponents, DmiListComponentsByClass, DmiListGroups, DmiListLanguages – Management Interface listing functions

SYNOPSIS

```
cc [ flag ... ] file ... -ldmimi -ldmi -lnsl -lrwtool [ library ... ]
#include <server.h>
#include <miapi.h>
bool_t DmiListAttributes(DmiListAttributesIN argin, DmiListAttributesOUT *result,
DmiRpcHandle *dmi_rpc_handle);

bool_t DmiListClassNames(DmiListClassNamesIN argin, DmiListClassNamesOUT
*result, DmiRpcHandle *dmi_rpc_handle);

bool_t DmiListComponents(DmiListComponentsIN argin, DmiListComponentsOUT
*result, DmiRpcHandle *dmi_rpc_handle);

bool_t DmiListComponentsByClass(DmiListComponentsByClassIN argin,
DmiListComponentsByClassOUT *result, DmiRpcHandle *dmi_rpc_handle);

bool_t DmiListGroups(DmiListGroupsIN argin, DmiListGroupsOUT *result,
DmiRpcHandle *dmi_rpc_handle);

bool_t DmiListLanguages(DmiListLanguagesIN argin, DmiListLanguagesOUT *result,
DmiRpcHandle *dmi_rpc_handle);
```

DESCRIPTION

The listing functions enables you to retrieve the names and the description of components in a system. You may also list components by class that match a specified criteria. The listing functions retrieve the set of language mappings installed for a specified component, retrieve class name strings for all groups in a component, retrieve a list of groups within a component, and retrieve the properties for one or more attributes in a group.

The `DmiListComponents()` function retrieves the name and (optionally) the description of components in a system. Use this to interrogate a system to determine what components are installed. The *argin* parameter is an instance of a `DmiListComponentsIN` structure containing the following members:

```
DmiHandle_t      handle;          /* an open session handle */
DmiRequestMode_t requestMode;    /* Unique, first, or next */
DmiUnsigned_t   maxCount;        /* maximum number to return,
0 for all */
DmiBoolean_t    getPragma;       /* get optional pragma string */
DmiBoolean_t    getDescription; /* get optional component
description */
DmiId_t         compId;          /* component ID to start with */
```

The *result* parameter is a pointer to a `DmiListComponentsOUT` structure containing the following members:

```
DmiErrorStatus_t    error_status;
DmiComponentList_t *reply;          /* list of components */
```

An enumeration accesses a specific component or may be used to sequentially access all components in a system. The caller may choose not to retrieve the component description by setting the value `getDescription` to false. The caller may choose not to retrieve the pragma string by setting the value of `getPragma` to false. The `maxCount`, `requestMode`, and `compId` parameters allow the caller to control the information returned by the Service Provider. When the `requestMode` is `DMI_UNIQUE`, `compId` specifies the first component requested (or only component if `maxCount` is one). When the `requestMode` is `DMI_NEXT`, `compId` specifies the component just before the one requested. When `requestMode` is `DMI_FIRST`, `compId` is unused.

To control the amount of information returned, the caller sets `maxCount` to something other than zero. The service provider must honor this limit on the amount of information returned. When `maxCount` is 0 the service provider returns information for all components, subject to the constraints imposed by `requestMode` and `compId`.

The `DmiListComponentsByClass()` function lists components that match specified criteria. Use this function to determine if a component contains a certain group or a certain row in a table. A filter condition may be that a component contains a specified group class name or that it contains a specific row in a specific group. As with `DmiListComponents()`, the description and pragma strings are optional return values. *argIn* is an instance of a `DmiListComponentsByClassIN` structure containing the following members:

```
DmiHandle_t        handle;          /* an open session handle */
DmiRequestMode_t  requestMode;     /* Unique, first or next */
DmiUnsigned_t     maxCount;        /* maximum number to return,
                                     or 0 for all */
DmiBoolean_t      getPragma;       /* get the optional pragma
                                     string */
DmiBoolean_t      getDescription;  /* get optional component
                                     description */
DmiId_t           compId;          /* component ID to start with */
DmiString_t       *className;      /* group class name string
                                     to match*/
DmiAttributeValues_t *keyList;     /* group row keys to match */
```

The *result* parameter is a pointer to a `DmiListComponentsbyClassOUT` structure containing the following members:

```
DmiErrorStatus_t    error_status;
DmiComponentList_t *reply;          /* list of components */
```

The `DmiListLanguages()` function retrieves the set of language mappings installed for the specified component. The *argin* parameter is an instance of a `DmiListLanguagesIN` structure containing the following members:

```
DmiHandle_t      handle;          /* An open session handle */
DmiUnsigned_t   maxCount;        /* maximum number to return,
                                  or 0 for all */
DmiId_t         compId;         /* Component to access */
```

The *result* parameter is a pointer to a `DmiListLanguagesOUT` structure containing the following members:

```
DmiErrorStatus_t error_status;
DmiStringList_t  *reply;        /* List of language strings */
```

The `DmiListClassNames()` function retrieves the class name strings for all groups in a component. This enables the management application to easily determine if a component contains a specific group, or groups. The *argin* parameter is an instance of a `DmiListClassNamesIN` structure containing the following members:

```
DmiHandle_t      handle;          /* An open session handle */
DmiUnsigned_t   maxCount;        /* maximum number to return,
                                  or 0 for all */
DmiId_t         compId;         /* Component to access */
```

The *result* parameter is a pointer to a `DmiListClassNamesOUT` structure containing the following members:

```
DmiErrorStatus_t error_status;
DmiClassNameList_t *reply;       /* List of class names and
                                  group IDs */
```

The `DmiListGroups()` function retrieves a list of groups within a component. With this function you can access a specific group or sequentially access all groups in a component. All enumerations of groups occur within the specified component and do not span components. The *argin* parameter is an instance of a `DmiListGroupsIN` structure containing the following members:

```
DmiHandle_t      handle;          /* An open session handle */
DmiRequestMode_t requestMode;    /* Unique, first or next group */
DmiUnsigned_t   maxCount;        /* Maximum number to return,
                                  or 0 for all */

DmiBoolean_t     getPragma;      /* Get the optional pragma string */
DmiBoolean_t     getDescription; /* Get optional group description */
DmiId_t         compId;         /* Component to access */
DmiId_t         groupId;        /* Group to start with, refer to
                                  requestMode */
```

The *result* parameter is a pointer to a `DmiListGroupsWithOUT` structure containing the following members:

```
DmiErrorStatus_t    error_status;
DmiGroupList_t     *reply;
```

The caller may choose not to retrieve the group description by setting the value `getDescription` to `false`. The caller may choose not to retrieve the pragma string by setting the value of `getPragma` to `false`. The `maxCount`, `requestMode`, and `groupId` parameters allow the caller to control the information returned by the Service Provider. When the `requestMode` is `DMI_UNIQUE`, `groupId` specifies the first group requested (or only group if `maxCount` is one). When the `requestMode` is `DMI_NEXT`, `groupId` specifies the group just before the one requested. When `requestMode` is `DMI_FIRST`, `groupId` is unused. To control the amount of information returned, the caller sets `maxCount` to something other than zero. The service provider must honor this limit on the amount of information returned. When `maxCount` is zero the service provider returns information for all groups, subject to the constraints imposed by `requestMode` and `groupId`.

The `DmiListAttributes()` function retrieves the properties for one or more attributes in a group. All enumerations of attributes occur within the specified group, and do not span groups. The *argIn* parameter is an instance of a `DmiListAttributesIN` structure containing the following members:

```
DmiHandle_t        handle;          /* An open session handle */
DmiRequestMode_t  requestMode;     /* Unique, first or next group */
DmiUnsigned_t     maxCount;        /* Maximum number to return,
                                     or 0 for all */
DmiBoolean_t      getPragma;       /* Get the optional pragma string */
DmiBoolean_t      getDescription; /* Get optional group description */
DmiId_t           compId;          /* Component to access */
DmiId_t           groupId;         /* Group to access */
DmiId_t           attribId;        /* Attribute to start with, refer
                                     to requestMode */
```

The *result* parameter is a pointer to a `DmiListAttributesOUT` structure containing the following members:

```
DmiErrorStatus_t    error_status;
DmiAttributeList_t *reply;         /* List of attributes */
```

You may choose not to retrieve the description string by setting the value of `getDescription` to `false`. Likewise, you may choose not to retrieve the pragma string by setting the value of `getPragma` to `false`. The `maxCount`, `requestMode`, and `attribId` parameters allow you to control the information returned by the Service Provider. When the `requestMode` is `DMI_UNIQUE`,

`attribId` specifies the first attribute requested (or only attribute if `maxCount` is one). When the `requestMode` is `DMI_NEXT`, `attribId` specifies the attribute just before the one requested. When `requestMode` is `DMI_FIRST`, `attribId` is unused. To control the amount of information returned, the caller sets `maxCount` to something other than zero. The Service Provider must honor this limit on the amount of information returned. When `maxCount` is zero the service provider returns information for all attributes, subject to the constraints imposed by `requestMode` and `attribId`.

RETURN VALUES

The `DmiListAttributes()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_ATTRIBUTE_NOT_FOUND
DMIERR_COMPONENT_NOT_FOUND
DMIERR_GROUP_NOT_FOUND
DMIERR_FILE_ERROR
```

The `DmiListClassNames()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_COMPONENT_NOT_FOUND
DMIERR_FILE_ERROR
```

The `DmiListComponents()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_COMPONENT_NOT_FOUND
DMIERR_FILE_ERROR
```

The `DmiListComponentsByClass()` function returns the following possible values:

```
DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_COMPONENT_NOT_FOUND
DMIERR_FILE_ERROR
```

The `DmiListGroup()` function returns the following possible values:

```

DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_COMPONENT_NOT_FOUND
DMIERR_GROUP_NOT_FOUND
DMIERR_FILE_ERROR

```

The `DmiListLanguages()` function returns the following possible values:

```

DMIERR_NO_ERROR
DMIERR_ILLEGAL_RPC_HANDLE
DMIERR_OUT_OF_MEMORY
DMIERR_ILLEGAL_PARAMETER
DMIERR_SP_INACTIVE
DMIERR_COMPONENT_NOT_FOUND
DMIERR_FILE_ERROR

```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | Unsafe |

SEE ALSO

`attributes(5)`

| | |
|--------------------|---|
| NAME | DmiRegisterCi, DmiUnRegisterCi, DmiOriginateEvent – Service Provider functions for components |
| SYNOPSIS | <pre>cc [flag ...] file ... -lci -ldmi -lnsl -lrwtool [library ...] #include <server.h> #include <ciapi.h> extern bool_t DmiRegisterCi(DmiRegisterCiIN <i>argin</i>, DmiRegisterCiOUT *<i>result</i>, DmiRpcHandle *<i>dmi_rpc_handle</i>); bool_t DmiUnregisterCi(DmiUnregisterCiIN <i>argin</i>, DmiUnregisterCiOUT *<i>result</i>, DmiRpcHandle *<i>dmi_rpc_handle</i>); bool_t DmiOriginateEvent(DmiOriginateEventIN <i>argin</i>, DmiOriginateEventOUT *<i>result</i>, DmiRpcHandle *<i>dmi_rpc_handle</i>);</pre> |
| DESCRIPTION | <p>These three functions provide component communication with the DMI through the Component Interface (CI).</p> <p>Component instrumentation code may register with the Service Provider to override its current mechanism for the registered attributes. Instead of manipulating the data in the MIF database or invoking programs, the Service Provider calls the entry points provided in the registration call. Once the component unregisters, the Service Provider returns to a normal method of processing requests for the data as defined in the MIF. Component instrumentation can temporarily interrupt normal processing to perform special functions.</p> <p>Registering attributes through the direct interface overrides attributes that are already being served through the direct interface. RPC is used for communication from the Service Provider to the component instrumentation.</p> <p>For all three functions, <i>argin</i> is the parameter passed to initiate an RPC call, <i>result</i> is the result of the RPC call, and <i>dmi_rpc_handle</i> is an open session RPC handle.</p> <p>The <code>DmiRegisterCi()</code> function registers a callable interface for components that have resident instrumentation code and/or to get the version of the Service Provider.</p> <p>The <code>DmiUnregisterCi()</code> function communicates to the Service Provider to remove a direct component instrumentation interface from the Service Provider table of registered interfaces.</p> <p>The <code>DmiOriginateEvent()</code> function originates an event for filtering and delivery. Any necessary indication filtering is performed by this function (or by subsequent processing) before the event is forwarded to the management applications.</p> <p>A component ID value of zero (0) specifies the event was generated by something that has not been installed as a component, and has no component ID.</p> |

RETURN VALUES

The `DmiRegisterCi()` function returns the following possible values:

- DMIERR_NO_ERROR
- DMIERR_ILLEGAL_HANDLE
- DMIERR_OUT_OF_MEMORY
- DMIERR_INSUFFICIENT_PRIVILEGES
- DMIERR_SP_INACTIVE
- DMIERR_ATTRIBUTE_NOT_FOUND
- DMIERR_COMPONENT_NOT_FOUND
- DMIERR_GROUP_NOT_FOUND
- DMIERR_DATABASE_CORRUPT
- DMIERR_OUT_OF_MEMORY
- DMIERR_ILLEGAL_DMI_LEVEL

The `DmiUnRegisterCi()` function returns the following possible values:

- DMIERR_NO_ERROR
- DMIERR_ILLEGAL_HANDLE
- DMIERR_OUT_OF_MEMORY
- DMIERR_INSUFFICIENT_PRIVILEGES
- DMIERR_SP_INACTIVE
- DMIERR_UNKNOWN_CI_REGISTRY

The `DmiOriginateEvent()` function returns the following possible values:

- DMIERR_NO_ERROR
- DMIERR_ILLEGAL_HANDLE
- DMIERR_OUT_OF_MEMORY
- DMIERR_INSUFFICIENT_PRIVILEGES
- DMIERR_SP_INACTIVE
- DMIERR_UNKNOWN_CI_REGISTRY

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | Unsafe |

SEE ALSO

`attributes(5)`

NAME elf32_checksum, elf64_checksum – return checksum of elf image

SYNOPSIS

```
cc [ flag ... ] file ... -lelf [ library ... ]
#include <libelf.h>
long elf32_checksum(Elf *elf);
long elf64_checksum(Elf *elf);
```

DESCRIPTION The `elf32_checksum()` function returns a simple checksum of selected sections of the image identified by *elf*. The value is typically used as the `.dynamic` tag `DT_CHECKSUM`, recorded in dynamic executables and shared objects.

Selected sections of the image are used to calculate the checksum in order that its value is not affected by utilities such as `strip(1)`.

For the 64-bit class, replace 32 with 64 as appropriate.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO `elf(3ELF)`, `elf_version(3ELF)`, `gelf(3ELF)`, `attributes(5)`

| NAME | elf32_fsize, elf64_fsize – return the size of an object file type | | | | |
|--------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lelf [library ...] #include <libelf.h> size_t elf32_fsize(Elf_Type type, size_t count, unsigned ver); size_t elf64_fsize(Elf_Type type, size_t count, unsigned ver);</pre> | | | | |
| DESCRIPTION | <p>elf32_fsize() gives the size in bytes of the 32-bit file representation of <i>count</i> data objects with the given <i>type</i> . The library uses version <i>ver</i> to calculate the size. See elf(3ELF) and elf_version(3ELF) .</p> <p>Constant values are available for the sizes of fundamental types:</p> <pre>Elf_Type File Size Memory Size ELF_T_ADDR ELF32_FSZ_ADDR sizeof(Elf32_Addr) ELF_T_BYTE 1 sizeof(unsigned char) ELF_T_HALF ELF32_FSZ_HALF sizeof(Elf32_Half) ELF_T_OFF ELF32_FSZ_OFF sizeof(Elf32_Off) ELF_T_SWORD ELF32_FSZ_SWORD sizeof(Elf32_Sword) ELF_T_WORD ELF32_FSZ_WORD sizeof(Elf32_Word)</pre> <p>elf32_fsize() returns 0 if the value of <i>type</i> or <i>ver</i> is unknown. See elf32_xlatetof(3ELF) for a list of the <i>type</i> values.</p> <p>For the 64-bit class, replace 32 with 64 as appropriate.</p> | | | | |
| ATTRIBUTES | <p>See attributes(5) for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | elf(3ELF) , elf32_xlatetof(3ELF) , elf_version(3ELF) , attributes(5) | | | | |

NAME elf32_getehdr, elf32_newehdr, elf64_getehdr, elf64_newehdr – retrieve class-dependent object file header

SYNOPSIS

```
cc [ flag ... ] file ... -lelf [ library ... ]
#include <libelf.h>
Elf32_Ehdr *elf32_getehdr(Elf *elf);

Elf32_Ehdr *elf32_newehdr(Elf *elf);

Elf64_Ehdr *elf64_getehdr(Elf *elf);

Elf64_Ehdr *elf64_newehdr(Elf *elf);
```

DESCRIPTION

For a 32-bit class file, `elf32_getehdr()` returns a pointer to an ELF header, if one is available for the ELF descriptor `elf`. If no header exists for the descriptor, `elf32_newehdr()` allocates a clean one, but it otherwise behaves the same as `elf32_getehdr()`. It does not allocate a new header if one exists already. If no header exists for `elf32_getehdr()`, one cannot be created for `elf32_newehdr()`, a system error occurs, the file is not a 32-bit class file, or `elf` is null, both functions return a null pointer.

For the 64-bit class, replace 32 with 64 as appropriate.

The header includes the following members:

```
unsigned char e_ident[EI_NIDENT];
Elf32_Half e_type;
Elf32_Half e_machine;
Elf32_Word e_version;
Elf32_Addr e_entry;
Elf32_Off e_phoff;
Elf32_Off e_shoff;
Elf32_Word e_flags;
Elf32_Half e_ehsize;
Elf32_Half e_phentsize;
Elf32_Half e_phnum;
Elf32_Half e_shentsize;
Elf32_Half e_shnum;
Elf32_Half e_shstrndx;
```

`elf32_newehdr()` automatically sets the `ELF_F_DIRTY` bit. See `elf_flagdata(3ELF)`. A program may use `elf_getident()` to inspect the identification bytes from a file.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

elf(3ELF), elf_begin(3ELF), elf_flagdata(3ELF), elf_getident(3ELF), attributes(5)

| | |
|--------------------|--|
| NAME | elf32_getphdr, elf32_newphdr, elf64_getphdr, elf64_newphdr – retrieve class-dependent program header table |
| SYNOPSIS | <pre>cc [flag ...] file... -lelf [library ...] #include <libelf.h> Elf32_Phdr *elf32_getphdr(Elf *elf); Elf32_Phdr *elf32_newphdr(Elf *elf, size_t count); Elf64_Phdr *elf64_getphdr(Elf *elf); Elf64_Phdr *elf64_newphdr(Elf *elf, size_t count);</pre> |
| DESCRIPTION | <p>For a 32-bit class file, <code>elf32_getphdr()</code> returns a pointer to the program execution header table, if one is available for the ELF descriptor <code>elf</code>.</p> <p><code>elf32_newphdr()</code> allocates a new table with <code>count</code> entries, regardless of whether one existed previously, and sets the <code>ELF_F_DIRTY</code> bit for the table. See <code>elf_flagdata(3ELF)</code>. Specifying a zero <code>count</code> deletes an existing table. Note this behavior differs from that of <code>elf32_newehdr()</code> allowing a program to replace or delete the program header table, changing its size if necessary. See <code>elf32_getehdr(3ELF)</code>.</p> <p>If no program header table exists, the file is not a 32-bit class file, an error occurs, or <code>elf</code> is <code>NULL</code>, both functions return a null pointer. Additionally, <code>elf32_newphdr()</code> returns a null pointer if <code>count</code> is 0.</p> <p>The table is an array of <code>Elf32_Phdr</code> structures, each of which includes the following members:</p> <pre>Elf32_Word p_type; Elf32_Off p_offset; Elf32_Addr p_vaddr; Elf32_Addr p_paddr; Elf32_Word p_filesz; Elf32_Word p_memsz; Elf32_Word p_flags; Elf32_Word p_align;</pre> <p>The <code>Elf64_Phdr</code> structures include the following members:</p> <pre>Elf64_Word p_type; Elf64_Word p_flags; Elf64_Off p_offset; Elf64_Addr p_vaddr; Elf64_Addr p_paddr; Elf64_Xword p_filesz; Elf64_Xword p_memsz; Elf64_Xword p_align;</pre> |

For the 64-bit class, replace 32 with 64 as appropriate.

The ELF header's `e_phnum` member tells how many entries the program header table has. See `elf32_getehdr(3ELF)`. A program may inspect this value to determine the size of an existing table; `elf32_newphdr()` automatically sets the member's value to `count`. If the program is building a new file, it is responsible for creating the file's ELF header before creating the program header table.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`elf(3ELF)`, `elf32_getehdr(3ELF)`, `elf_begin(3ELF)`,
`elf_flagdata(3ELF)`, `attributes(5)`

NAME elf32_getshdr, elf64_getshdr – retrieve class-dependent section header

SYNOPSIS

```
cc [ flag ... ] file ... -lelf [ library ... ]
#include <libelf.h>
Elf32_Shdr *elf32_getshdr(Elf_Scn *scn);
```

```
Elf64_Shdr *elf64_getshdr(Elf_Scn *scn);
```

DESCRIPTION For a 32-bit class file, `elf32_getshdr()` returns a pointer to a section header for the section descriptor `scn`. Otherwise, the file is not a 32-bit class file, `scn` was `NULL`, or an error occurred; `elf32_getshdr()` then returns `NULL`.

The `elf32_getshdr` header includes the following members:

```
Elf32_Word sh_name;
Elf32_Word sh_type;
Elf32_Word sh_flags;
Elf32_Addr sh_addr;
Elf32_Off sh_offset;
Elf32_Word sh_size;
Elf32_Word sh_link;
Elf32_Word sh_info;
Elf32_Word sh_addralign;
Elf32_Word sh_entsize;
```

while the `elf64_getshdr` header includes the following members:

```
Elf64_Word sh_name;
Elf64_Word sh_type;
Elf64_Xword sh_flags;
Elf64_Addr sh_addr;
Elf64_Off sh_offset;
Elf64_Xword sh_size;
Elf64_Word sh_link;
Elf64_Word sh_info;
Elf64_Xword sh_addralign;
Elf64_Xword sh_entsize;
```

For the 64-bit class, replace 32 with 64 as appropriate.

If the program is building a new file, it is responsible for creating the file's ELF header before creating sections.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

elf(3ELF), elf_flagdata(3ELF), elf_getscn(3ELF), elf_strptr(3ELF),
attributes(5)

| | | | | | | | | | |
|------------------------|---|--------------------|--|---------------------|---|---------------------|---|------------------------|---|
| NAME | elf32_xlatetof, elf32_xlatetom, elf64_xlatetof, elf64_xlatetom – class-dependent data translation | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file... -lelf [library ...] #include <libelf.h> Elf_Data *elf32_xlatetof(Elf_Data *dst, const Elf_Data *src, unsigned encode); Elf_Data *elf32_xlatetom(Elf_Data *dst, const Elf_Data *src, unsigned encode); Elf_Data *elf64_xlatetof(Elf_Data *dst, const Elf_Data *src, unsigned encode); Elf_Data *elf64_xlatetom(Elf_Data *dst, const Elf_Data *src, unsigned encode);</pre> | | | | | | | | |
| DESCRIPTION | <p>elf32_xlatetom() translates various data structures from their 32-bit class file representations to their memory representations; elf32_xlatetof() provides the inverse. This conversion is particularly important for cross development environments. <i>src</i> is a pointer to the source buffer that holds the original data; <i>dst</i> is a pointer to a destination buffer that will hold the translated copy. <i>encode</i> gives the byte encoding in which the file objects are to be represented and must have one of the encoding values defined for the ELF header's <code>e_ident[EI_DATA]</code> entry (see <code>elf_getident(3ELF)</code>). If the data can be translated, the functions return <i>dst</i> . Otherwise, they return <code>NULL</code> because an error occurred, such as incompatible types, destination buffer overflow, etc.</p> <p><code>elf_getdata(3ELF)</code> describes the <code>Elf_Data</code> descriptor, which the translation routines use as follows:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>d_buf</code></td> <td>Both the source and destination must have valid buffer pointers.</td> </tr> <tr> <td style="padding-right: 20px;"><code>d_type</code></td> <td>This member's value specifies the type of the data to which <code>d_buf</code> points and the type of data to be created in the destination. The program supplies a <code>d_type</code> value in the source; the library sets the destination's <code>d_type</code> to the same value. These values are summarized below.</td> </tr> <tr> <td style="padding-right: 20px;"><code>d_size</code></td> <td>This member holds the total size, in bytes, of the memory occupied by the source data and the size allocated for the destination data. If the destination buffer is not large enough, the routines do not change its original contents. The translation routines reset the destination's <code>d_size</code> member to the actual size required, after the translation occurs. The source and destination sizes may differ.</td> </tr> <tr> <td style="padding-right: 20px;"><code>d_version</code></td> <td>This member holds the version number of the objects (desired) in the buffer. The source and destination versions are independent.</td> </tr> </table> | <code>d_buf</code> | Both the source and destination must have valid buffer pointers. | <code>d_type</code> | This member's value specifies the type of the data to which <code>d_buf</code> points and the type of data to be created in the destination. The program supplies a <code>d_type</code> value in the source; the library sets the destination's <code>d_type</code> to the same value. These values are summarized below. | <code>d_size</code> | This member holds the total size, in bytes, of the memory occupied by the source data and the size allocated for the destination data. If the destination buffer is not large enough, the routines do not change its original contents. The translation routines reset the destination's <code>d_size</code> member to the actual size required, after the translation occurs. The source and destination sizes may differ. | <code>d_version</code> | This member holds the version number of the objects (desired) in the buffer. The source and destination versions are independent. |
| <code>d_buf</code> | Both the source and destination must have valid buffer pointers. | | | | | | | | |
| <code>d_type</code> | This member's value specifies the type of the data to which <code>d_buf</code> points and the type of data to be created in the destination. The program supplies a <code>d_type</code> value in the source; the library sets the destination's <code>d_type</code> to the same value. These values are summarized below. | | | | | | | | |
| <code>d_size</code> | This member holds the total size, in bytes, of the memory occupied by the source data and the size allocated for the destination data. If the destination buffer is not large enough, the routines do not change its original contents. The translation routines reset the destination's <code>d_size</code> member to the actual size required, after the translation occurs. The source and destination sizes may differ. | | | | | | | | |
| <code>d_version</code> | This member holds the version number of the objects (desired) in the buffer. The source and destination versions are independent. | | | | | | | | |

Translation routines allow the source and destination buffers to coincide. That is, `dst->d_buf` may equal `src->d_buf`. Other cases where the source and destination buffers overlap give undefined behavior.

```
Elf_Type      32-Bit Memory Type
ELF_T_ADDR   Elf32_Addr
ELF_T_BYTE   unsigned char
ELF_T_DYN    Elf32_Dyn
ELF_T_EHDR   Elf32_Ehdr
ELF_T_HALF   Elf32_Half
ELF_T_OFF    Elf32_Off
ELF_T_PHDR   Elf32_Phdr
ELF_T_REL    Elf32_Rel
ELF_T_RELA   Elf32_Rela
ELF_T_SHDR   Elf32_Shdr
ELF_T_SWORD  Elf32_Sword
ELF_T_SYM    Elf32_Sym
ELF_T_WORD   Elf32_Word
```

Translating buffers of type `ELF_T_BYTE` does not change the byte order.

For the 64-bit class, replace 32 with 64 as appropriate.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`elf(3ELF)`, `elf32_fsize(3ELF)`, `elf_getdata(3ELF)`,
`elf_getident(3ELF)`, `attributes(5)`

NAME | elf – object file access library

SYNOPSIS | `cc [flag ...] file ... -lelf [library ...]`
`#include <libelf.h>`

DESCRIPTION | Functions in the ELF access library let a program manipulate ELF (Executable and Linking Format) object files, archive files, and archive members. The header provides type and function declarations for all library services.

Programs communicate with many of the higher-level routines using an *ELF descriptor*. That is, when the program starts working with a file, `elf_begin(3ELF)` creates an ELF descriptor through which the program manipulates the structures and information in the file. These ELF descriptors can be used both to read and to write files. After the program establishes an ELF descriptor for a file, it may then obtain *section descriptors* to manipulate the sections of the file (see `elf_getscn(3ELF)`). Sections hold the bulk of an object file's real information, such as text, data, the symbol table, and so on. A section descriptor “belongs” to a particular ELF descriptor, just as a section belongs to a file. Finally, *data descriptors* are available through section descriptors, allowing the program to manipulate the information associated with a section. A data descriptor “belongs” to a section descriptor.

Descriptors provide private handles to a file and its pieces. In other words, a data descriptor is associated with one section descriptor, which is associated with one ELF descriptor, which is associated with one file. Although descriptors are private, they give access to data that may be shared. Consider programs that combine input files, using incoming data to create or update another file. Such a program might get data descriptors for an input and an output section. It then could update the output descriptor to reuse the input descriptor's data. That is, the descriptors are distinct, but they could share the associated data bytes. This sharing avoids the space overhead for duplicate buffers and the performance overhead for copying data unnecessarily.

File Classes | ELF provides a framework in which to define a family of object files, supporting multiple processors and architectures. An important distinction among object files is the *class*, or capacity, of the file. The 32-bit class supports architectures in which a 32-bit object can represent addresses, file sizes, and so on, as in the following:

| Name | Purpose |
|--------------------------|-------------------------|
| <code>Elf32_Addr</code> | Unsigned address |
| <code>Elf32_Half</code> | Unsigned medium integer |
| <code>Elf32_Off</code> | Unsigned file offset |
| <code>Elf32_Sword</code> | Signed large integer |

| Name | Purpose |
|---------------|------------------------|
| Elf32_Word | Unsigned large integer |
| unsigned char | Unsigned small integer |

The 64-bit class works the same as the 32-bit class, substituting 64 for 32 as necessary. Other classes will be defined as necessary, to support larger (or smaller) machines. Some library services deal only with data objects for a specific class, while others are class-independent. To make this distinction clear, library function names reflect their status, as described below.

Data Representation

Conceptually, two parallel sets of objects support cross compilation environments. One set corresponds to file contents, while the other set corresponds to the native memory image of the program manipulating the file. Type definitions supplied by the headers work on the native machine, which may have different data encodings (size, byte order, and so on) than the target machine. Although native memory objects should be at least as big as the file objects (to avoid information loss), they may be bigger if that is more natural for the host machine.

Translation facilities exist to convert between file and memory representations. Some library routines convert data automatically, while others leave conversion as the program's responsibility. Either way, programs that create object files must write file-typed objects to those files; programs that read object files must take a similar view. See `elf32_xlatetof(3ELF)` and `elf32_fsize(3ELF)` for more information.

Programs may translate data explicitly, taking full control over the object file layout and semantics. If the program prefers not to have and exercise complete control, the library provides a higher-level interface that hides many object file details. `elf_begin()` and related functions let a program deal with the native memory types, converting between memory objects and their file equivalents automatically when reading or writing an object file.

ELF Versions

Object file versions allow ELF to adapt to new requirements. *Three independent versions* can be important to a program. First, an application program knows about a particular version by virtue of being compiled with certain headers. Second, the access library similarly is compiled with header files that control what versions it understands. Third, an ELF object file holds a value identifying its version, determined by the ELF version known by the file's creator. Ideally, all three versions would be the same, but they may differ.

If a program's version is newer than the access library, the program might use information unknown to the library. Translation routines might not work

properly, leading to undefined behavior. This condition merits installing a new library.

The library's version might be newer than the program's and the file's. The library understands old versions, thus avoiding compatibility problems in this case.

Finally, a file's version might be newer than either the program or the library understands. The program might or might not be able to process the file properly, depending on whether the file has extra information and whether that information can be safely ignored. Again, the safe alternative is to install a new library that understands the file's version.

To accommodate these differences, a program must use `elf_version(3ELF)` to pass its version to the library, thus establishing the *working version* for the process. Using this, the library accepts data from and presents data to the program in the proper representations. When the library reads object files, it uses each file's version to interpret the data. When writing files or converting memory types to the file equivalents, the library uses the program's working version for the file data.

System Services

As mentioned above, `elf_begin()` and related routines provide a higher-level interface to ELF files, performing input and output on behalf of the application program. These routines assume a program can hold entire files in memory, without explicitly using temporary files. When reading a file, the library routines bring the data into memory and perform subsequent operations on the memory copy. Programs that wish to read or write large object files with this model must execute on a machine with a large process virtual address space. If the underlying operating system limits the number of open files, a program can use `elf_cntl(3ELF)` to retrieve all necessary data from the file, allowing the program to close the file descriptor and reuse it.

Although the `elf_begin()` interfaces are convenient and efficient for many programs, they might be inappropriate for some. In those cases, an application may invoke the `elf32_xlatetom(3ELF)` or `elf32_xlatetof(3ELF)` data translation routines directly. These routines perform no input or output, leaving that as the application's responsibility. By assuming a larger share of the job, an application controls its input and output model.

Library Names

Names associated with the library take several forms.

| | |
|-------------------------|---|
| <code>elf_name</code> | These class-independent names perform some service, <i>name</i> , for the program. |
| <code>elf32_name</code> | Service names with an embedded class, 32 here, indicate they work only for the designated class of files. |

| | |
|-----------------------------|--|
| <code>Elf_Type</code> | Data types can be class-independent as well, distinguished by <i>Type</i> . |
| <code>Elf32_Type</code> | Class-dependent data types have an embedded class name, 32 here. |
| <code>ELF_C_CMD</code> | Several functions take commands that control their actions. These values are members of the <code>Elf_Cmd</code> enumeration; they range from zero through <code>ELF_C_NUM-1</code> . |
| <code>ELF_F_FLAG</code> | Several functions take flags that control library status and/or actions. Flags are bits that may be combined. |
| <code>ELF32_FSZ_TYPE</code> | These constants give the file sizes in bytes of the basic ELF types for the 32-bit class of files. See <code>elf32_fsize()</code> for more information. |
| <code>ELF_K_KIND</code> | The function <code>elf_kind()</code> identifies the <i>KIND</i> of file associated with an ELF descriptor. These values are members of the <code>Elf_Kind</code> enumeration; they range from zero through <code>ELF_K_NUM-1</code> . |
| <code>ELF_T_TYPE</code> | When a service function, such as <code>elf32_xlatetom()</code> or <code>elf32_xlatetof()</code> , deals with multiple types, names of this form specify the desired <i>TYPE</i> . Thus, for example, <code>ELF_T_EHDR</code> is directly related to <code>Elf32_Ehdr</code> . These values are members of the <code>Elf_Type</code> enumeration; they range from zero through <code>ELF_T_NUM-1</code> . |

EXAMPLES

EXAMPLE 1 An interpretation of elf file.

The basic interpretation of an ELF file consists of:

- opening an ELF object file
- obtaining an ELF descriptor
- analyzing the file using the descriptor.

The following example opens the file, obtains the ELF descriptor, and prints out the names of each section in the file.

```
#include <fcntl.h>
#include <stdio.h>
#include <libelf.h>
#include <stdlib.h>
#include <string.h>
```



```

static void failure(void);
void
main(int argc, char ** argv)
{
    Elf32_Shdr *   shdr;
    Elf32_Ehdr *   ehdr;
    Elf *          elf;
    Elf_Scn *      scn;
    Elf_Data *     data;
    int            fd;
    unsigned int   cnt;

    /* Open the input file */
    if ((fd = open(argv[1], O_RDONLY)) == -1)
        exit(1);

    /* Obtain the ELF descriptor */
    (void) elf_version(EV_CURRENT);
    if ((elf = elf_begin(fd, ELF_C_READ, NULL)) == NULL)
        failure();

    /* Obtain the .shstrtab data buffer */
    if (((ehdr = elf32_getehdr(elf)) == NULL) ||
        ((scn = elf_getscn(elf, ehdr->e_shstrndx)) == NULL) ||
        ((data = elf_getdata(scn, NULL)) == NULL))
        failure();

    /* Traverse input filename, printing each section */
    for (cnt = 1, scn = NULL; scn = elf_nextscn(elf, scn); cnt++) {
        if ((shdr = elf32_getshdr(scn)) == NULL)
            failure();
        (void) printf("[%d]   %s\n", cnt,
            (char *)data->d_buf + shdr->sh_name);
    }
    /* end main */
}

static void
failure()
{
    (void) fprintf(stderr, "%s\n", elf_errmsg(elf_errno()));
    exit(1);
}

```

ATTRIBUTES

See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

[elf32_checksum\(3ELF\)](#), [elf32_fsize\(3ELF\)](#), [elf32_getshdr\(3ELF\)](#),
[elf32_xlatetof\(3ELF\)](#), [elf_begin\(3ELF\)](#), [elf_cntl\(3ELF\)](#),
[elf_errmsg\(3ELF\)](#), [elf_fill\(3ELF\)](#), [elf_getarhdr\(3ELF\)](#),
[elf_getarsym\(3ELF\)](#), [elf_getbase\(3ELF\)](#), [elf_getdata\(3ELF\)](#),

```
elf_getident(3ELF), elf_getscn(3ELF), elf_hash(3ELF),
elf_kind(3ELF), elf_memory(3ELF), elf_rawfile(3ELF),
elf_strptr(3ELF), elf_update(3ELF), elf_version(3ELF), gelf(3ELF),
ar(3HEAD), attributes(5)
```

ANSI C Programmer's Guide

SPARC only

a.out(4)

NOTES

Information in the ELF headers is separated into common parts and processor-specific parts. A program can make a processor's information available by including the appropriate header: `<sys/elf_NAME.h>` where *NAME* matches the processor name as used in the ELF file header.

| Name | Processor |
|-------|-----------------------------|
| M32 | AT&T WE 32100 |
| SPARC | SPARC |
| 386 | Intel 80386, 80486, Pentium |

Other processors will be added to the table as necessary.

To illustrate, a program could use the following code to “see” the processor-specific information for the SPARC based system.

```
#include <libelf.h>
#include <sys/elf_SPARC.h>
```

Without the `<sys/elf_SPARC.h>` definition, only the common ELF information would be visible.

A program could use the following code to “see” the processor-specific information for the Intel 80386:

```
#include <libelf.h>
#include <sys/elf_386.h>
```

Without the `<sys/elf_386.h>` definition, only the common ELF information would be visible.

Although reading the objects is rather straightforward, writing/updating them can corrupt the shared offsets among sections. Upon creation, relationships are established among the sections that must be maintained even if the object's size is changed.

| | |
|--------------------|---|
| NAME | elf_begin, elf_end, elf_memory, elf_next, elf_rand – process ELF object files |
| SYNOPSIS | <pre>cc [flag...] file ... -lelf [library ...] #include <libelf.h> Elf *elf_begin(int <i>fildev</i>, Elf_Cmd <i>cmd</i>, Elf *<i>ref</i>); int elf_end(Elf *<i>elf</i>); Elf *elf_memory(char *<i>image</i>, size_t <i>tsz</i>); Elf_Cmd elf_next(Elf *<i>elf</i>); size_t elf_rand(Elf *<i>elf</i>, size_t <i>offset</i>);</pre> |
| DESCRIPTION | <p>elf_begin(), elf_end(), elf_memory(), elf_next(), and elf_rand() work together to process Executable and Linking Format (ELF) object files, either individually or as members of archives. After obtaining an ELF descriptor from elf_begin() or elf_memory(), the program may read an existing file, update an existing file, or create a new file. <i>fildev</i> is an open file descriptor that elf_begin() uses for reading or writing. <i>elf</i> is an ELF descriptor previously returned from elf_begin(). The initial file offset (see lseek(2)) is unconstrained, and the resulting file offset is undefined.</p> <p><i>cmd</i> may have the following values:</p> <p>ELF_C_NULL When a program sets <i>cmd</i> to this value, elf_begin() returns a null pointer, without opening a new descriptor. <i>ref</i> is ignored for this command. See the examples below for more information.</p> <p>ELF_C_READ When a program wishes to examine the contents of an existing file, it should set <i>cmd</i> to this value. Depending on the value of <i>ref</i>, this command examines archive members or entire files. Three cases can occur.</p> <p>First, if <i>ref</i> is a null pointer, elf_begin() allocates a new ELF descriptor and prepares to process the entire file. If the file being read is an archive, elf_begin() also prepares the resulting descriptor to examine the initial archive member on the next call to elf_begin(), as if the program had used elf_next() or elf_rand() to “move” to the initial member.</p> <p>Second, if <i>ref</i> is a non-null descriptor associated with an archive file, elf_begin() lets a program obtain a separate ELF descriptor associated with an individual member. The program should have used elf_next() or elf_rand() to position <i>ref</i> appropriately (except for the initial member, which elf_begin() prepares; see the example below). In</p> |

this case, *fildev* should be the same file descriptor used for the parent archive.

Finally, if *ref* is a non-null ELF descriptor that is not an archive, `elf_begin()` increments the number of activations for the descriptor and returns *ref*, without allocating a new descriptor and without changing the descriptor's read/write permissions. To terminate the descriptor for *ref*, the program must call `elf_end()` once for each activation. See the examples below for more information.

- `ELF_C_RDWR` This command duplicates the actions of `ELF_C_READ` and additionally allows the program to update the file image (see `elf_update(3ELF)`). That is, using `ELF_C_READ` gives a read-only view of the file, while `ELF_C_RDWR` lets the program read *and* write the file. `ELF_C_RDWR` is not valid for archive members. If *ref* is non-null, it must have been created with the `ELF_C_RDWR` command.
- `ELF_C_WRITE` If the program wishes to ignore previous file contents, presumably to create a new file, it should set *cmd* to this value. *ref* is ignored for this command.

`elf_begin()` “works” on all files (including files with zero bytes), providing it can allocate memory for its internal structures and read any necessary information from the file. Programs reading object files thus may call `elf_kind(3ELF)` or `elf32_getehdr(3ELF)` to determine the file type (only object files have an ELF header). If the file is an archive with no more members to process, or an error occurs, `elf_begin()` returns a null pointer. Otherwise, the return value is a non-null ELF descriptor.

Before the first call to `elf_begin()`, a program must call `elf_version()` to coordinate versions.

`elf_end()` is used to terminate an ELF descriptor, *elf*, and to deallocate data associated with the descriptor. Until the program terminates a descriptor, the data remain allocated. A null pointer is allowed as an argument, to simplify error handling. If the program wishes to write data associated with the ELF descriptor to the file, it must use `elf_update()` before calling `elf_end()`.

Calling `elf_end()` removes one activation and returns the remaining activation count. The library does not terminate the descriptor until the activation count reaches 0. Consequently, a 0 return value indicates the ELF descriptor is no longer valid.

`elf_memory()` returns a pointer to an ELF descriptor, the ELF image has read operations enabled (`ELF_C_READ`). *image* is a pointer to an image of the Elf file mapped into memory, *sz* is the size of the ELF image. An ELF image that

is mapped in with `elf_memory()` may be read and modified, but the ELF image size may not be changed.

`elf_next()` provides sequential access to the next archive member. That is, having an ELF descriptor, *elf*, associated with an archive member, `elf_next()` prepares the containing archive to access the following member when the program calls `elf_begin()`. After successfully positioning an archive for the next member, `elf_next()` returns the value `ELF_C_READ`. Otherwise, the open file was not an archive, *elf* was `NULL`, or an error occurred, and the return value is `ELF_C_NULL`. In either case, the return value may be passed as an argument to `elf_begin()`, specifying the appropriate action.

`elf_rand()` provides random archive processing, preparing *elf* to access an arbitrary archive member. *elf* must be a descriptor for the archive itself, not a member within the archive. *offset* gives the byte offset from the beginning of the archive to the archive header of the desired member. See `elf_getarsym(3ELF)` for more information about archive member offsets. When `elf_rand()` works, it returns *offset*. Otherwise, it returns 0, because an error occurred, *elf* was `NULL`, or the file was not an archive (no archive member can have a zero offset). A program may mix random and sequential archive processing.

System Services

When processing a file, the library decides when to read or write the file, depending on the program's requests. Normally, the library assumes the file descriptor remains usable for the life of the ELF descriptor. If, however, a program must process many files simultaneously and the underlying operating system limits the number of open files, the program can use `elf_cntl()` to let it reuse file descriptors. After calling `elf_cntl()` with appropriate arguments, the program may close the file descriptor without interfering with the library.

All data associated with an ELF descriptor remain allocated until `elf_end()` terminates the descriptor's last activation. After the descriptors have been terminated, the storage is released; attempting to reference such data gives undefined behavior. Consequently, a program that deals with multiple input (or output) files must keep the ELF descriptors active until it finishes with them.

EXAMPLES

EXAMPLE 1 A sample program of calling the `elf_begin()` function.

A prototype for reading a file appears on the next page. If the file is a simple object file, the program executes the loop one time, receiving a null descriptor in the second iteration. In this case, both *elf* and *arf* will have the same value, the activation count will be 2, and the program calls `elf_end()` twice to terminate the descriptor. If the file is an archive, the loop processes each archive member in turn, ignoring those that are not object files.

```
if (elf_version(EV_CURRENT) == EV_NONE)
{
    /* library out of date */
    /* recover from error */
}
```

```

cmd = ELF_C_READ;
arf = elf_begin(fildes, cmd, (Elf *)0);
while ((elf = elf_begin(fildes, cmd, arf)) != 0)
{
    if ((ehdr = elf32_getehdr(elf)) != 0)
    {
        /* process the file ... */
    }
    cmd = elf_next(elf);
    elf_end(elf);
}
elf_end(arf);

```

Alternatively, the next example illustrates random archive processing. After identifying the file as an archive, the program repeatedly processes archive members of interest. For clarity, this example omits error checking and ignores simple object files. Additionally, this fragment preserves the ELF descriptors for all archive members, because it does not call `elf_end()` to terminate them.

```

elf_version(EV_CURRENT);
arf = elf_begin(fildes, ELF_C_READ, (Elf *)0);
if (elf_kind(arf) != ELF_K_AR)
{
    /* not an archive */
}
/* initial processing */
/* set offset = ... for desired member header */
while (elf_rand(arf, offset) == offset)
{
    if ((elf = elf_begin(fildes, ELF_C_READ, arf)) == 0)
        break;
    if ((ehdr = elf32_getehdr(elf)) != 0)
    {
        /* process archive member ... */
    }
    /* set offset = ... for desired member header */
}

```

An archive starts with a “magic string” that has SARMAG bytes; the initial archive member follows immediately. An application could thus provide the following function to rewind an archive (the function returns `-1` for errors and `0` otherwise).

```

#include <ar.h>
#include <libelf.h>
int
rewindelf(Elf *elf)
{
    if (elf_rand(elf, (size_t)SARMAG) == SARMAG)
        return 0;
    return -1;
}

```

The following outline shows how one might create a new ELF file. This example is simplified to show the overall flow.

```
elf_version(EV_CURRENT);
fildes = open("path/name", O_RDWR|O_TRUNC|O_CREAT, 0666);
if ((elf = elf_begin(fildes, ELF_C_WRITE, (Elf *)0)) == 0)
    return;
ehdr = elf32_newehdr(elf);
phdr = elf32_newphdr(elf, count);
scn = elf_newscn(elf);
shdr = elf32_getshdr(scn);
data = elf_newdata(scn);
elf_update(elf, ELF_C_WRITE);
elf_end(elf);
```

Finally, the following outline shows how one might update an existing ELF file. Again, this example is simplified to show the overall flow.

```
elf_version(EV_CURRENT);
fildes = open("path/name", O_RDWR);
elf = elf_begin(fildes, ELF_C_RDWR, (Elf *)0);
/* add new or delete old information */
...
/* ensure that the memory image of the file is complete */
elf_update(elf, ELF_C_NULL);
elf_update(elf, ELF_C_WRITE); /* update file */
elf_end(elf);
```

Notice that both file creation examples open the file with write *and* read permissions. On systems that support `mmap(2)`, the library uses it to enhance performance, and `mmap(2)` requires a readable file descriptor. Although the library can use a write-only file descriptor, the application will not obtain the performance advantages of `mmap(2)`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`creat(2)`, `lseek(2)`, `mmap(2)`, `open(2)`, `elf(3ELF)`, `elf32_getehdr(3ELF)`, `elf_cntl(3ELF)`, `elf_getarhdr(3ELF)`, `elf_getarsym(3ELF)`, `elf_getbase(3ELF)`, `elf_getdata(3ELF)`, `elf_getscn(3ELF)`, `elf_kind(3ELF)`, `elf_rawfile(3ELF)`, `elf_update(3ELF)`, `elf_version(3ELF)`, `ar(3HEAD)`, `attributes(5)`

| NAME | elf_cntl – control an elf file descriptor | | | | |
|--------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lelf [library ...] #include <libelf.h> int elf_cntl(Elf *elf, Elf_Cmd cmd);</pre> | | | | |
| DESCRIPTION | <p>elf_cntl() instructs the library to modify its behavior with respect to an ELF descriptor, <i>elf</i>. As elf_begin(3ELF) describes, an ELF descriptor can have multiple activations, and multiple ELF descriptors may share a single file descriptor. Generally, elf_cntl() commands apply to all activations of <i>elf</i>. Moreover, if the ELF descriptor is associated with an archive file, descriptors for members within the archive will also be affected as described below. Unless stated otherwise, operations on archive members do not affect the descriptor for the containing archive.</p> <p>The <i>cmd</i> argument tells what actions to take and may have the following values:</p> <p>ELF_C_FDDONE This value tells the library not to use the file descriptor associated with <i>elf</i>. A program should use this command when it has requested all the information it cares to use and wishes to avoid the overhead of reading the rest of the file. The memory for all completed operations remains valid, but later file operations, such as the initial elf_getdata() for a section, will fail if the data are not in memory already.</p> <p>ELF_C_FDREAD This command is similar to ELF_C_FDDONE, except it forces the library to read the rest of the file. A program should use this command when it must close the file descriptor but has not yet read everything it needs from the file. After elf_cntl() completes the ELF_C_FDREAD command, future operations, such as elf_getdata(), will use the memory version of the file without needing to use the file descriptor.</p> <p>If elf_cntl() succeeds, it returns 0. Otherwise <i>elf</i> was NULL or an error occurred, and the function returns -1.</p> | | | | |
| ATTRIBUTES | <p>See attributes(5) for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | elf(3ELF), elf_begin(3ELF), elf_getdata(3ELF), elf_rawfile(3ELF), attributes(5) | | | | |
| NOTES | If the program wishes to use the “raw” operations (see elf_rawdata(), which elf_getdata(3ELF) describes, and elf_rawfile(3ELF)) after disabling the | | | | |

file descriptor with `ELF_C_FDDONE` or `ELF_C_FDREAD`, it must execute the raw operations explicitly beforehand. Otherwise, the raw file operations will fail. Calling `elf_rawfile()` makes the entire image available, thus supporting subsequent `elf_rawdata()` calls.

| NAME | elf_errmsg, elf_errno – error handling | | | | |
|--------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lelf [library ...] #include <libelf.h> const char *elf_errmsg(int err); int elf_errno(void);</pre> | | | | |
| DESCRIPTION | <p>If an ELF library function fails, a program may call <code>elf_errno()</code> to retrieve the library's internal error number. As a side effect, this function resets the internal error number to 0, which indicates no error.</p> <p><code>elf_errmsg()</code> takes an error number, <i>err</i>, and returns a null-terminated error message (with no trailing new-line) that describes the problem. A zero <i>err</i> retrieves a message for the most recent error. If no error has occurred, the return value is a null pointer (not a pointer to the null string). Using <i>err</i> of -1 also retrieves the most recent error, except it guarantees a non-null return value, even when no error has occurred. If no message is available for the given number, <code>elf_errmsg()</code> returns a pointer to an appropriate message. This function does not have the side effect of clearing the internal error number.</p> | | | | |
| EXAMPLES | <p>EXAMPLE 1 A sample program of calling the <code>elf_errmsg()</code> function.</p> <p>The following fragment clears the internal error number and checks it later for errors. Unless an error occurs after the first call to <code>elf_errno()</code>, the next call will return 0.</p> <pre>(void)elf_errno(); /* processing ... */ while (more_to_do) { if ((err = elf_errno()) != 0) { /* print msg */ msg = elf_errmsg(err); } }</pre> | | | | |
| ATTRIBUTES | <p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>elf(3ELF)</code> , <code>attributes(5)</code> | | | | |

NAME | elf_fill – set fill byte

SYNOPSIS | cc [*flag ...*] *file ...* -lelf [*library ...*]
 #include <libelf.h>
 void elf_fill(int *fill*);

DESCRIPTION | Alignment constraints for ELF files sometimes require the presence of “holes.” For example, if the data for one section are required to begin on an eight-byte boundary, but the preceding section is too “short,” the library must fill the intervening bytes. These bytes are set to the *fill* character. The library uses zero bytes unless the application supplies a value. See elf_getdata(3ELF) for more information about these holes.

ATTRIBUTES | See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | elf(3ELF), elf_flagdata(3ELF), elf_getdata(3ELF), elf_update(3ELF), attributes(5)

NOTES | An application can assume control of the object file organization by setting the ELF_F_LAYOUT bit (see elf_flagdata(3ELF)). When this is done, the library does *not* fill holes.

| | |
|--------------------|--|
| NAME | elf_flagdata, elf_flagehdr, elf_flagelf, elf_flagphdr, elf_flagscn, elf_flagshdr – manipulate flags |
| SYNOPSIS | <pre>cc [flag ...] file ... -lelf [library ...] #include <libelf.h> unsigned elf_flagdata(Elf_Data *data, Elf_Cmd cmd, unsigned flags); unsigned elf_flagehdr(Elf *elf, Elf_Cmd cmd, unsigned flags); unsigned elf_flagelf(Elf *elf, Elf_Cmd cmd, unsigned flags); unsigned elf_flagphdr(Elf *elf, Elf_Cmd cmd, unsigned flags); unsigned elf_flagscn(Elf_Scn *scn, Elf_Cmd cmd, unsigned flags); unsigned elf_flagshdr(Elf_Scn *scn, Elf_Cmd cmd, unsigned flags);</pre> |
| DESCRIPTION | <p>These functions manipulate the flags associated with various structures of an ELF file. Given an ELF descriptor (<i>elf</i>), a data descriptor (<i>data</i>), or a section descriptor (<i>scn</i>), the functions may set or clear the associated status bits, returning the updated bits. A null descriptor is allowed, to simplify error handling; all functions return 0 for this degenerate case.</p> <p><i>cmd</i> may have the following values:</p> <p>ELF_C_CLR The functions clear the bits that are asserted in <i>flags</i>. Only the non-zero bits in <i>flags</i> are cleared; zero bits do not change the status of the descriptor.</p> <p>ELF_C_SET The functions set the bits that are asserted in <i>flags</i>. Only the non-zero bits in <i>flags</i> are set; zero bits do not change the status of the descriptor.</p> <p>Descriptions of the defined <i>flags</i> bits appear below:</p> <p>ELF_F_DIRTY When the program intends to write an ELF file, this flag asserts the associated information needs to be written to the file. Thus, for example, a program that wished to update the ELF header of an existing file would call <code>elf_flagehdr()</code> with this bit set in <i>flags</i> and <i>cmd</i> equal to <code>ELF_C_SET</code>. A later call to <code>elf_update()</code> would write the marked header to the file.</p> <p>ELF_F_LAYOUT Normally, the library decides how to arrange an output file. That is, it automatically decides where to place sections, how to align them in the file, etc. If this bit is set for an ELF descriptor, the program assumes responsibility for determining all file positions. This bit is meaningful only for <code>elf_flagelf()</code> and applies to the entire file associated with the descriptor.</p> |

When a flag bit is set for an item, it affects all the subitems as well. Thus, for example, if the program sets the `ELF_F_DIRTY` bit with `elf_flagelf()`, the entire logical file is “dirty.”

EXAMPLES

EXAMPLE 1 A sample display of calling the `elf_flagdata()` function.

The following fragment shows how one might mark the ELF header to be written to the output file:

```
/* dirty ehdr ... */
ehdr = elf32_getehdr(elf);
elf_flagehdr(elf, ELF_C_SET, ELF_F_DIRTY);
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`elf(3ELF)`, `elf32_getehdr(3ELF)`, `elf_getdata(3ELF)`,
`elf_update(3ELF)`, `attributes(5)`

| | |
|--------------------|--|
| NAME | elf_getarhdr – retrieve archive member header |
| SYNOPSIS | <pre>cc [flag ...] file ... -l<code>elf</code> [library...] #include <libelf.h> Elf_Arhdr *elf_getarhdr(Elf *elf);</pre> |
| DESCRIPTION | <p><code>elf_getarhdr()</code> returns a pointer to an archive member header, if one is available for the ELF descriptor <code>elf</code>. Otherwise, no archive member header exists, an error occurred, or <code>elf</code> was null; <code>elf_getarhdr()</code> then returns a null value. The header includes the following members.</p> <pre>char *ar_name; time_t ar_date; uid_t ar_uid; gid_t ar_gid; mode_t ar_mode; off_t ar_size; char *ar_rawname;</pre> <p>An archive member name, available through <code>ar_name</code>, is a null-terminated string, with the <code>ar</code> format control characters removed. The <code>ar_rawname</code> member holds a null-terminated string that represents the original name bytes in the file, including the terminating slash and trailing blanks as specified in the archive format.</p> <p>In addition to “regular” archive members, the archive format defines some special members. All special member names begin with a slash (/), distinguishing them from regular members (whose names may not contain a slash). These special members have the names (<code>ar_name</code>) defined below.</p> <p>/ This is the archive symbol table. If present, it will be the first archive member. A program may access the archive symbol table through <code>elf_getarsym()</code>. The information in the symbol table is useful for random archive processing (see <code>elf_rand()</code> on <code>elf_begin(3ELF)</code>).</p> <p>// This member, if present, holds a string table for long archive member names. An archive member’s header contains a 16-byte area for the name, which may be exceeded in some file systems. The library automatically retrieves long member names from the string table, setting <code>ar_name</code> to the appropriate value.</p> <p>Under some error conditions, a member’s name might not be available. Although this causes the library to set <code>ar_name</code> to a null pointer, the <code>ar_rawname</code> member will be set as usual.</p> |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

elf(3ELF), elf_begin(3ELF), elf_getarsym(3ELF), ar(3HEAD),
attributes(5)

| NAME | elf_getarsym – retrieve archive symbol table | | | | |
|--------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>cc [flag ...] file ... -l<code>elf</code> [library ...] #include <libelf.h> Elf_Arsym *elf_getarsym(Elf *elf, size_t *ptr);</pre> | | | | |
| DESCRIPTION | <p>elf_getarsym() returns a pointer to the archive symbol table, if one is available for the ELF descriptor <i>elf</i>. Otherwise, the archive doesn't have a symbol table, an error occurred, or <i>elf</i> was null; elf_getarsym() then returns a null value. The symbol table is an array of structures that include the following members.</p> <pre>char *as_name; size_t as_off; unsigned long as_hash;</pre> <p>These members have the following semantics:</p> <p>as_name A pointer to a null-terminated symbol name resides here.</p> <p>as_off This value is a byte offset from the beginning of the archive to the member's header. The archive member residing at the given offset defines the associated symbol. Values in <i>as_off</i> may be passed as arguments to elf_rand(). See elf_begin(3ELF) to access the desired archive member.</p> <p>as_hash This is a hash value for the name, as computed by elf_hash().</p> <p>If <i>ptr</i> is non-null, the library stores the number of table entries in the location to which <i>ptr</i> points. This value is set to 0 when the return value is NULL. The table's last entry, which is included in the count, has a null <i>as_name</i>, a zero value for <i>as_off</i>, and ~0UL for <i>as_hash</i>.</p> <p>The hash value returned is guaranteed not to be the bit pattern of all ones (~0UL).</p> | | | | |
| ATTRIBUTES | <p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | elf(3ELF), elf_begin(3ELF), elf_getarhdr(3ELF), elf_hash(3ELF), ar(3HEAD), attributes(5) | | | | |

NAME | elf_getbase – get the base offset for an object file

SYNOPSIS | `cc [flag ...] file ... -lelf [library ...]`
 | `#include <libelf.h>`
 | `off_t elf_getbase(Elf *elf);`

DESCRIPTION | `elf_getbase()` returns the file offset of the first byte of the file or archive member associated with *elf*, if it is known or obtainable, and `-1` otherwise. A null *elf* is allowed, to simplify error handling; the return value in this case is `-1`. The base offset of an archive member is the beginning of the member's information, *not* the beginning of the archive member header.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `elf(3ELF)`, `elf_begin(3ELF)`, `ar(3HEAD)`, `attributes(5)`

| | |
|--------------------|--|
| NAME | elf_getdata, elf_newdata, elf_rawdata – get section data |
| SYNOPSIS | <pre>cc [flag ...] file ... -lelf [library ...] #include <libelf.h> Elf_Data *elf_getdata(Elf_Scn *scn, Elf_Data *data); Elf_Data *elf_newdata(Elf_Scn *scn); Elf_Data *elf_rawdata(Elf_Scn *scn, Elf_Data *data);</pre> |
| DESCRIPTION | <p>These functions access and manipulate the data associated with a section descriptor, <i>scn</i>. When reading an existing file, a section will have a single data buffer associated with it. A program may build a new section in pieces, however, composing the new data from multiple data buffers. For this reason, the data for a section should be viewed as a list of buffers, each of which is available through a data descriptor.</p> <p><code>elf_getdata()</code> lets a program step through a section's data list. If the incoming data descriptor, <i>data</i>, is null, the function returns the first buffer associated with the section. Otherwise, <i>data</i> should be a data descriptor associated with <i>scn</i>, and the function gives the program access to the next data element for the section. If <i>scn</i> is null or an error occurs, <code>elf_getdata()</code> returns a null pointer.</p> <p><code>elf_getdata()</code> translates the data from file representations into memory representations (see <code>elf32_xlatetof(3ELF)</code>) and presents objects with memory data types to the program, based on the file's <i>class</i> (see <code>elf(3ELF)</code>). The working library version (see <code>elf_version(3ELF)</code>) specifies what version of the memory structures the program wishes <code>elf_getdata()</code> to present.</p> <p><code>elf_newdata()</code> creates a new data descriptor for a section, appending it to any data elements already associated with the section. As described below, the new data descriptor appears empty, indicating the element holds no data. For convenience, the descriptor's type (<code>d_type</code> below) is set to <code>ELF_T_BYTE</code>, and the version (<code>d_version</code> below) is set to the working version. The program is responsible for setting (or changing) the descriptor members as needed. This function implicitly sets the <code>ELF_F_DIRTY</code> bit for the section's data (see <code>elf_flagdata(3ELF)</code>). If <i>scn</i> is null or an error occurs, <code>elf_newdata()</code> returns a null pointer.</p> <p><code>elf_rawdata()</code> differs from <code>elf_getdata()</code> by returning only uninterpreted bytes, regardless of the section type. This function typically should be used only to retrieve a section image from a file being read, and then only when a program must avoid the automatic data translation described below. Moreover, a program may not close or disable (see <code>elf_cntl(3ELF)</code>) the file descriptor associated with <i>elf</i> before the initial raw operation, because <code>elf_rawdata()</code> might read the data from the file to ensure it doesn't interfere with <code>elf_getdata()</code>. See</p> |

`elf_rawfile(3ELF)` for a related facility that applies to the entire file. When `elf_getdata()` provides the right translation, its use is recommended over `elf_rawdata()`. If `scn` is null or an error occurs, `elf_rawdata()` returns a null pointer.

The `Elf_Data` structure includes the following members:

```
void *d_buf;
Elf_Type d_type;
size_t d_size;
off_t d_off;
size_t d_align;
unsigned d_version;
```

These members are available for direct manipulation by the program. Descriptions appear below.

| | |
|------------------------|--|
| <code>d_buf</code> | A pointer to the data buffer resides here. A data element with no data has a null pointer. |
| <code>d_type</code> | This member's value specifies the type of the data to which <code>d_buf</code> points. A section's type determines how to interpret the section contents, as summarized below. |
| <code>d_size</code> | This member holds the total size, in bytes, of the memory occupied by the data. This may differ from the size as represented in the file. The size will be zero if no data exist. (See the discussion of <code>SHT_NOBITS</code> below for more information.) |
| <code>d_off</code> | This member gives the offset, within the section, at which the buffer resides. This offset is relative to the file's section, not the memory object's. |
| <code>d_align</code> | This member holds the buffer's required alignment, from the beginning of the section. That is, <code>d_off</code> will be a multiple of this member's value. For example, if this member's value is 4, the beginning of the buffer will be four-byte aligned within the section. Moreover, the entire section will be aligned to the maximum of its constituents, thus ensuring appropriate alignment for a buffer within the section and within the file. |
| <code>d_version</code> | This member holds the version number of the objects in the buffer. When the library originally read the data from the object file, it used the working version to control the translation to memory objects. |

Data Alignment

As mentioned above, data buffers within a section have explicit alignment constraints. Consequently, adjacent buffers sometimes will not abut, causing “holes” within a section. Programs that create output files have two ways of dealing with these holes.

First, the program can use `elf_fill()` to tell the library how to set the intervening bytes. When the library must generate gaps in the file, it uses the fill byte to initialize the data there. The library’s initial fill value is 0, and `elf_fill()` lets the application change that.

Second, the application can generate its own data buffers to occupy the gaps, filling the gaps with values appropriate for the section being created. A program might even use different fill values for different sections. For example, it could set text sections’ bytes to no-operation instructions, while filling data section holes with zero. Using this technique, the library finds no holes to fill, because the application eliminated them.

Section and Memory Types

`elf_getdata()` interprets sections’ data according to the section type, as noted in the section header available through `elf32_getshdr()`. The following table shows the section types and how the library represents them with memory data types for the 32-bit file class. Other classes would have similar tables. By implication, the memory data types control translation by `elf32_xlatetof(3ELF)`

```

Section Type Elf_Type 32-Bit Type
SHT_DYNAMIC ELF_T_DYN Elf32_Dyn
SHT_DYNSYM ELF_T_SYM Elf32_Sym
SHT_HASH ELF_T_WORD Elf32_Word
SHT_NOBITS ELF_T_BYTE unsigned char
SHT_NOTE ELF_T_BYTE unsigned char
SHT_NULL none none
SHT_PROGBITS ELF_T_BYTE unsigned char
SHT_REL ELF_T_REL Elf32_Rel
SHT_RELA ELF_T_RELA Elf32_Rela
SHT_STRTAB ELF_T_BYTE unsigned char
SHT_SYMTAB ELF_T_SYM Elf32_Sym
SHT_SUNW_verdef ELF_T_VDEF Elf32_Verdef
SHT_SUNW_verneed ELF_T_VNEED Elf32_Verneed
SHT_SUNW_versym ELF_T_HALF Elf32_Versym
other ELF_T_BYTE unsigned char

```

`elf_rawdata()` creates a buffer with type `ELF_T_BYTE`.

As mentioned above, the program’s working version controls what structures the library creates for the application. The library similarly interprets section types according to the versions. If a section type belongs to a version newer than the application’s working version, the library does not translate the section data. Because the application cannot know the data format in this case, the library

presents an untranslated buffer of type `ELF_T_BYTE`, just as it would for an unrecognized section type.

A section with a special type, `SHT_NOBITS`, occupies no space in an object file, even when the section header indicates a non-zero size. `elf_getdata()` and `elf_rawdata()` work on such a section, setting the `data` structure to have a null buffer pointer and the type indicated above. Although no data are present, the `d_size` value is set to the size from the section header. When a program is creating a new section of type `SHT_NOBITS`, it should use `elf_newdata()` to add data buffers to the section. These empty data buffers should have the `d_size` members set to the desired size and the `d_buf` members set to `NULL`.

EXAMPLES

EXAMPLE 1 A sample program of calling `elf_getdata()`.

The following fragment obtains the string table that holds section names (ignoring error checking). See `elf_strptr(3ELF)` for a variation of string table handling.

```
ehdr = elf32_getehdr(elf);
scn = elf_getscn(elf, (size_t)ehdr->e_shstrndx);
shdr = elf32_getshdr(scn);
if (shdr->sh_type != SHT_STRTAB)
{
    /* not a string table */
}
data = 0;
if ((data = elf_getdata(scn, data)) == 0 || data->d_size == 0)
{
    /* error or no data */
}
```

The `e_shstrndx` member in an ELF header holds the section table index of the string table. The program gets a section descriptor for that section, verifies it is a string table, and then retrieves the data. When this fragment finishes, `data=>d_buf` points at the first byte of the string table, and `data=>d_size` holds the string table's size in bytes.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`elf(3ELF)`, `elf32_getehdr(3ELF)`, `elf32_getshdr(3ELF)`, `elf32_xlatetof(3ELF)`, `elf_cntl(3ELF)`, `elf_fill(3ELF)`, `elf_flagdata(3ELF)`, `elf_getscn(3ELF)`, `elf_rawfile(3ELF)`, `elf_strptr(3ELF)`, `elf_version(3ELF)`, `attributes(5)`

NAME elf_getident – retrieve file identification data

SYNOPSIS `cc [flag ...] file ... -lelf [library ...]`
`#include <libelf.h>`
`char *elf_getident(Elf *elf, size_t *ptr);`

DESCRIPTION As `elf(3ELF)` explains, ELF provides a framework for various classes of files, where basic objects may have 32 bits, 64 bits, etc. To accommodate these differences, without forcing the larger sizes on smaller machines, the initial bytes in an ELF file hold identification information common to all file classes. Every ELF header's `e_ident` has `EI_NIDENT` bytes with the following interpretation:

| e_ident Index | Value | Purpose |
|---------------|--------------|---------------------|
| EI_MAG0 | ELFMAG0 | File identification |
| EI_MAG1 | ELFMAG1 | |
| EI_MAG2 | ELFMAG2 | |
| EI_MAG3 | ELFMAG3 | |
| EI_CLASS | ELFCLASSNONE | File class |
| | ELFCLASS32 | |
| | ELFCLASS64 | |
| EI_DATA | ELFDATANONE | Data encoding |
| | ELFDATA2LSB | |
| | ELFDATA2MSB | |
| EI_VERSION | EV_CURRENT | File version |
| 7-15 | 0 | Unused, set to zero |

Other kinds of files (see `elf_kind(3ELF)`) also may have identification data, though they would not conform to `e_ident`.

`elf_getident()` returns a pointer to the file's "initial bytes." If the library recognizes the file, a conversion from the file image to the memory image may occur. In any case, the identification bytes are guaranteed not to have been modified, though the size of the unmodified area depends on the file type. If `ptr`

is non-null, the library stores the number of identification bytes in the location to which *ptr* points. If no data are present, *elf* is null, or an error occurs, the return value is a null pointer, with 0 stored through *ptr*, if *ptr* is non-null.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`elf(3ELF)`, `elf32_getehdr(3ELF)`, `elf_begin(3ELF)`, `elf_kind(3ELF)`, `elf_rawfile(3ELF)`, `attributes(5)`

| | |
|--------------------|--|
| NAME | elf_getscn, elf_ndxscn, elf_newscn, elf_nextscn – get section information |
| SYNOPSIS | <pre>cc [flag ...] file ... -lelf [library ...] #include <libelf.h> Elf_Scn *elf_getscn(Elf *elf, size_t index); size_t elf_ndxscn(Elf_Scn *scn); Elf_Scn *elf_newscn(Elf *elf); Elf_Scn *elf_nextscn(Elf *elf, Elf_Scn *scn);</pre> |
| DESCRIPTION | <p>These functions provide indexed and sequential access to the sections associated with the ELF descriptor <i>elf</i>. If the program is building a new file, it is responsible for creating the file's ELF header before creating sections; see <code>elf32_getehdr(3ELF)</code>.</p> <p><code>elf_getscn()</code> returns a section descriptor, given an <i>index</i> into the file's section header table. Note that the first "real" section has an index of 1. Although a program can get a section descriptor for the section whose <i>index</i> is 0 (<code>SHN_UNDEF</code>, the undefined section), the section has no data and the section header is "empty" (though present). If the specified section does not exist, an error occurs, or <i>elf</i> is null, <code>elf_getscn()</code> returns a null pointer.</p> <p><code>elf_newscn()</code> creates a new section and appends it to the list for <i>elf</i>. Because the <code>SHN_UNDEF</code> section is required and not "interesting" to applications, the library creates it automatically. Thus the first call to <code>elf_newscn()</code> for an ELF descriptor with no existing sections returns a descriptor for section 1. If an error occurs or <i>elf</i> is null, <code>elf_newscn()</code> returns a null pointer.</p> <p>After creating a new section descriptor, the program can use <code>elf32_getshdr()</code> to retrieve the newly created, "clean" section header. The new section descriptor will have no associated data (see <code>elf_getdata(3ELF)</code>). When creating a new section in this way, the library updates the <code>e_shnum</code> member of the ELF header and sets the <code>ELF_F_DIRTY</code> bit for the section (see <code>elf_flagdata(3ELF)</code>). If the program is building a new file, it is responsible for creating the file's ELF header (see <code>elf32_getehdr(3ELF)</code>) before creating new sections.</p> <p><code>elf_nextscn()</code> takes an existing section descriptor, <i>scn</i>, and returns a section descriptor for the next higher section. One may use a null <i>scn</i> to obtain a section descriptor for the section whose index is 1 (skipping the section whose index is <code>SHN_UNDEF</code>). If no further sections are present or an error occurs, <code>elf_nextscn()</code> returns a null pointer.</p> <p><code>elf_ndxscn()</code> takes an existing section descriptor, <i>scn</i>, and returns its section table index. If <i>scn</i> is null or an error occurs, <code>elf_ndxscn()</code> returns <code>SHN_UNDEF</code>.</p> |

EXAMPLES

EXAMPLE 1 A sample of calling `elf_getscn()` function.

An example of sequential access appears below. Each pass through the loop processes the next section in the file; the loop terminates when all sections have been processed.

```

scn = 0;
while ((scn = elf_nextscn(elf, scn)) != 0)
{
    /* process section */
}
    
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`elf(3ELF)`, `elf32_getehdr(3ELF)`, `elf32_getshdr(3ELF)`, `elf_begin(3ELF)`, `elf_flagdata(3ELF)`, `elf_getdata(3ELF)`, `attributes(5)`

| NAME | elf_hash – compute hash value | | | | |
|--------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lelf [<i>library ...</i>] #include <libelf.h> unsigned long elf_hash (const char * <i>name</i>); | | | | |
| DESCRIPTION | <p>elf_hash() computes a hash value, given a null terminated string, <i>name</i>. The returned hash value, <i>h</i>, can be used as a bucket index, typically after computing $h \bmod x$ to ensure appropriate bounds.</p> <p>Hash tables may be built on one machine and used on another because elf_hash() uses unsigned arithmetic to avoid possible differences in various machines' signed arithmetic. Although <i>name</i> is shown as char* above, elf_hash() treats it as unsigned char* to avoid sign extension differences. Using char* eliminates type conflicts with expressions such as elf_hash(<i>name</i>) .</p> <p>ELF files' symbol hash tables are computed using this function (see elf_getdata(3ELF) and elf32_xlatetof(3ELF)). The hash value returned is guaranteed not to be the bit pattern of all ones (~0UL).</p> | | | | |
| ATTRIBUTES | <p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | elf(3ELF), elf32_xlatetof(3ELF), elf_getdata(3ELF), attributes(5) | | | | |

NAME elf_kind – determine file type

SYNOPSIS cc [*flag ...*] *file ...* -lelf [*library ...*]
 #include <libelf.h>
 Elf_Kind elf_kind(Elf *elf);

DESCRIPTION This function returns a value identifying the kind of file associated with an ELF descriptor (*elf*). Defined values are below:

ELF_K_AR The file is an archive [see ar(3HEAD)]. An ELF descriptor may also be associated with an archive *member*, not the archive itself, and then elf_kind() identifies the member's type.

ELF_K_COFF The file is a COFF object file. elf_begin(3ELF) describes the library's handling for COFF files.

ELF_K_ELF The file is an ELF file. The program may use elf_getident() to determine the class. Other functions, such as elf32_getehdr(), are available to retrieve other file information.

ELF_K_NONE This indicates a kind of file unknown to the library.

Other values are reserved, to be assigned as needed to new kinds of files. *elf* should be a value previously returned by elf_begin(). A null pointer is allowed, to simplify error handling, and causes elf_kind() to return ELF_K_NONE.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO elf(3ELF), elf32_getehdr(3ELF), elf_begin(3ELF), elf_getident(3ELF), ar(3HEAD), attributes(5)

| NAME | elf_rawfile – retrieve uninterpreted file contents | | | | |
|--------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag...</i>] <i>file ...</i> -lelf [<i>library ...</i>] #include <libelf.h> char *elf_rawfile(Elf *elf, size_t *ptr); | | | | |
| DESCRIPTION | <p>elf_rawfile() returns a pointer to an uninterpreted byte image of the file. This function should be used only to retrieve a file being read. For example, a program might use elf_rawfile() to retrieve the bytes for an archive member.</p> <p>A program may not close or disable (see elf_cntl(3ELF)) the file descriptor associated with <i>elf</i> before the initial call to elf_rawfile() , because elf_rawfile() might have to read the data from the file if it does not already have the original bytes in memory. Generally, this function is more efficient for unknown file types than for object files. The library implicitly translates object files in memory, while it leaves unknown files unmodified. Thus, asking for the uninterpreted image of an object file may create a duplicate copy in memory.</p> <p>elf_rawdata() is a related function, providing access to sections within a file. See elf_getdata(3ELF).</p> <p>If <i>ptr</i> is non-null, the library also stores the file's size, in bytes, in the location to which <i>ptr</i> points. If no data are present, <i>elf</i> is null, or an error occurs, the return value is a null pointer, with 0 stored through <i>ptr</i>, if <i>ptr</i> is non-null.</p> | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | elf(3ELF), elf32_getehdr(3ELF), elf_begin(3ELF), elf_cntl(3ELF), elf_getdata(3ELF), elf_getident(3ELF), elf_kind(3ELF), attributes(5) | | | | |
| NOTES | <p>A program that uses elf_rawfile() and that also interprets the same file as an object file potentially has two copies of the bytes in memory. If such a program requests the raw image first, before it asks for translated information (through such functions as elf32_getehdr(), elf_getdata(), and so on), the library “freezes” its original memory copy for the raw image. It then uses this frozen copy as the source for creating translated objects, without reading the file again. Consequently, the application should view the raw file image returned by elf_rawfile() as a read-only buffer, unless it wants to alter its own view of data subsequently translated. In any case, the application may alter the translated objects without changing bytes visible in the raw image.</p> <p>Multiple calls to elf_rawfile() with the same ELF descriptor return the same value; the library does not create duplicate copies of the file.</p> | | | | |

NAME elf_strptr – make a string pointer

SYNOPSIS cc [*flag ...*] *file ...* -lelf [*library ...*]
 #include <libelf.h>
 char *elf_strptr(Elf *elf, size_t section, size_t offset);

DESCRIPTION This function converts a string section *offset* to a string pointer. *elf* identifies the file in which the string section resides, and *section* identifies the section table index for the strings. elf_strptr() normally returns a pointer to a string, but it returns a null pointer when *elf* is null, *section* is invalid or is not a section of type SHT_STRTAB, the section data cannot be obtained, *offset* is invalid, or an error occurs.

EXAMPLES **EXAMPLE 1** A sample program of calling elf_strptr() function.

A prototype for retrieving section names appears below. The file header specifies the section name string table in the e_shstrndx member. The following code loops through the sections, printing their names.

```
/* handle the error */
if ((ehdr = elf32_getehdr(elf)) == 0) {
    return;
}
ndx = ehdr->e_shstrndx;
scn = 0;
while ((scn = elf_nextscn(elf, scn)) != 0) {
    char *name = 0;
    if ((shdr = elf32_getshdr(scn)) != 0)
        name = elf_strptr(elf, ndx, (size_t)shdr->sh_name);
    printf("%s\n", name? name: "(null)");
}
```

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO elf(3ELF), elf32_getshdr(3ELF), elf32_xlatetof(3ELF), elf_getdata(3ELF), attributes(5)

NOTES A program may call elf_getdata() to retrieve an entire string table section. For some applications, that would be both more efficient and more convenient than using elf_strptr().

| | |
|--------------------|--|
| NAME | elf_update – update an ELF descriptor |
| SYNOPSIS | <pre>cc [flag ...] file ... -lelf [library ...] #include <libelf.h> off_t elf_update(Elf *elf, Elf_Cmd cmd);</pre> |
| DESCRIPTION | <p>elf_update() causes the library to examine the information associated with an ELF descriptor, <i>elf</i>, and to recalculate the structural data needed to generate the file's image.</p> <p><i>cmd</i> may have the following values:</p> <p>ELF_C_NULL This value tells elf_update() to recalculate various values, updating only the ELF descriptor's memory structures. Any modified structures are flagged with the ELF_F_DIRTY bit. A program thus can update the structural information and then reexamine them without changing the file associated with the ELF descriptor. Because this does not change the file, the ELF descriptor may allow reading, writing, or both reading and writing (see elf_begin (3ELF)).</p> <p>ELF_C_WRITE If <i>cmd</i> has this value, elf_update() duplicates its ELF_C_NULL actions and also writes any "dirty" information associated with the ELF descriptor to the file. That is, when a program has used elf_getdata(3ELF) or the elf_flagdata(3ELF) facilities to supply new (or update existing) information for an ELF descriptor, those data will be examined, coordinated, translated if necessary (see elf32_xlatetof(3ELF)), and written to the file. When portions of the file are written, any ELF_F_DIRTY bits are reset, indicating those items no longer need to be written to the file (see elf_flagdata(3ELF)). The sections' data are written in the order of their section header entries, and the section header table is written to the end of the file. When the ELF descriptor was created with elf_begin(), it must have allowed writing the file. That is, the elf_begin() command must have been either ELF_C_RDWR or ELF_C_WRITE.</p> <p>If elf_update() succeeds, it returns the total size of the file image (not the memory image), in bytes. Otherwise an error occurred, and the function returns -1.</p> <p>When updating the internal structures, elf_update() sets some members itself. Members listed below are the application's responsibility and retain the values given by the program.</p> <p>The following table shows ELF Header members:</p> |

| Member | Notes |
|------------------|---------------------------------------|
| e_ident[EI_DATA] | Library controls other e_ident values |
| e_type | |
| e_machine | |
| e_version | |
| e_entry | |
| e_phoff | Only when ELF_F_LAYOUT asserted |
| e_shoff | Only when ELF_F_LAYOUT asserted |
| e_flags | |
| e_shstrndx | |

The following table shows the Program Header members:

| Member | Notes |
|----------|------------------------------|
| p_type | The application controls all |
| p_offset | program header entries |
| p_vaddr | |
| p_paddr | |
| p_filesz | |
| p_memsz | |
| p_flags | |
| p_align | |

The following table shows the Section Header members:

| Member | Notes |
|----------|-------|
| sh_name | |
| sh_type | |
| sh_flags | |
| sh_addr | |

| | |
|--------------|--|
| sh_offset | Only when <code>ELF_F_LAYOUT</code> asserted |
| sh_size | Only when <code>ELF_F_LAYOUT</code> asserted |
| sh_link | |
| sh_info | |
| sh_addralign | Only when <code>ELF_F_LAYOUT</code> asserted |
| sh_entsize | |

The following table shows the Data Descriptor members:

| Member | Notes |
|-----------|--|
| d_buf | |
| d_type | |
| d_size | |
| d_off | Only when <code>ELF_F_LAYOUT</code> asserted |
| d_align | |
| d_version | |

Note that the program is responsible for two particularly important members (among others) in the ELF header. The `e_version` member controls the version of data structures written to the file. If the version is `EV_NONE`, the library uses its own internal version. The `e_ident[EI_DATA]` entry controls the data encoding used in the file. As a special case, the value may be `ELFDATANONE` to request the native data encoding for the host machine. An error occurs in this case if the native encoding doesn't match a file encoding known by the library.

Further note that the program is responsible for the `sh_entsize` section header member. Although the library sets it for sections with known types, it cannot reliably know the correct value for all sections. Consequently, the library relies on the program to provide the values for unknown section types. If the entry size is unknown or not applicable, the value should be set to 0.

When deciding how to build the output file, `elf_update()` obeys the alignments of individual data buffers to create output sections. A section's most strictly aligned data buffer controls the section's alignment. The library also inserts padding between buffers, as necessary, to ensure the proper alignment of each buffer.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

elf(3ELF), elf32_fsize(3ELF), elf32_getehdr(3ELF),
elf32_getshdr(3ELF), elf32_xlatetof(3ELF), elf_begin(3ELF),
elf_flagdata(3ELF), elf_getdata(3ELF), attributes(5)

NOTES

As mentioned above, the `ELF_C_WRITE` command translates data as necessary, before writing them to the file. This translation is *not* always transparent to the application program. If a program has obtained pointers to data associated with a file (for example, see `elf32_getehdr(3ELF)` and `elf_getdata(3ELF)`), the program should reestablish the pointers after calling `elf_update()`.

| NAME | elf_version – coordinate ELF library and application versions | | | | |
|--------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lelf [library ...] #include <libelf.h> unsigned elf_version(unsigned ver);</pre> | | | | |
| DESCRIPTION | <p>As elf(3ELF) explains, the program, the library, and an object file have independent notions of the latest ELF version. elf_version() lets a program query the ELF library's <i>internal version</i>. It further lets the program specify what memory types it uses by giving its own <i>working version</i>, <i>ver</i>, to the library. Every program that uses the ELF library must coordinate versions as described below.</p> <p>The header <libelf.h> supplies the version to the program with the macro EV_CURRENT. If the library's internal version (the highest version known to the library) is lower than that known by the program itself, the library may lack semantic knowledge assumed by the program. Accordingly, elf_version() will not accept a working version unknown to the library.</p> <p>Passing <i>ver</i> equal to EV_NONE causes elf_version() to return the library's internal version, without altering the working version. If <i>ver</i> is a version known to the library, elf_version() returns the previous (or initial) working version number. Otherwise, the working version remains unchanged and elf_version() returns EV_NONE.</p> | | | | |
| EXAMPLES | <p>EXAMPLE 1 A sample display of using the elf_version() function.</p> <p>The following excerpt from an application program protects itself from using an older library:</p> <pre>if (elf_version(EV_CURRENT) == EV_NONE) { /* library out of date */ /* recover from error */ }</pre> | | | | |
| ATTRIBUTES | <p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | elf(3ELF), elf32_xlatetof(3ELF), elf_begin(3ELF), attributes(5) | | | | |
| NOTES | The working version should be the same for all operations on a particular ELF descriptor. Changing the version between operations on a descriptor will probably not give the expected results. | | | | |

| NAME | erf, erfc – error and complementary error functions | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double erf (double <i>x</i>); double erfc (double <i>x</i>); | | | | |
| DESCRIPTION | The <code>erf()</code> function computes the error function of x , defined as: $\frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$ The <code>erfc()</code> function computes $1.0 - \text{erf}(x)$. | | | | |
| RETURN VALUES | Upon successful completion, <code>erf()</code> and <code>erfc()</code> return the value of the error function and complementary error function, respectively. If x is NaN, NaN is returned. | | | | |
| ERRORS | No errors will occur. | | | | |
| USAGE | The <code>erfc()</code> function is provided because of the extreme loss of relative accuracy if <code>erf(x)</code> is called for large x and the result subtracted from 1.0. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>isnan(3M)</code> , <code>attributes(5)</code> | | | | |

| NAME | exp – exponential function | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double exp (double <i>x</i>); | | | | |
| DESCRIPTION | The <code>exp()</code> function computes the exponential of <i>x</i> , defined as e^x . | | | | |
| RETURN VALUES | Upon successful completion, <code>exp()</code> returns the exponential of <i>x</i> . If the correct value would cause overflow, <code>exp()</code> returns HUGE_VAL and sets <code>errno</code> to ERANGE. If the correct value would cause underflow to zero, <code>exp()</code> returns 0 and may set <code>errno</code> to ERANGE. If <i>x</i> is NaN, NaN is returned. For exceptional cases, <code>matherr(3M)</code> tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The <code>exp()</code> function will fail if: ERANGE The result overflows. The <code>exp()</code> function may fail if: ERANGE The result underflows. | | | | |
| USAGE | An application wishing to check for error situations should set <code>errno</code> to 0 before calling <code>exp()</code> . If <code>errno</code> is non-zero on return, or the return value is NaN an error has occurred. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>isnan(3M)</code> , <code>log(3M)</code> , <code>matherr(3M)</code> , <code>mp(3MP)</code> , <code>attributes(5)</code> , <code>standards(5)</code> | | | | |
| NOTES | Prior to Solaris 2.6, there was a conflict between the <code>pow</code> function in this library and the <code>pow</code> function in the <code>libmp</code> library. This conflict was resolved by prepending <code>mp_</code> to all functions in the <code>libmp</code> library. See <code>mp(3MP)</code> for details. | | | | |

NAME | expm1 – computes exponential functions

SYNOPSIS | `cc [flag ...] file ... -lm [library ...]`
`#include <math.h>`
`double expm1(double x);`

DESCRIPTION | The `expm1()` function computes $e^x-1.0$.

RETURN VALUES | If x is NaN, then the function returns NaN.
 If x is positive infinity, `expm1()` returns positive infinity.
 If x is negative infinity, `expm1()` returns -1.0 .
 If the value overflows, `expm1()` returns `HUGE_VAL`.

ERRORS | No errors will occur.

USAGE | The value of `expm1(x)` may be more accurate than `exp(x)-1.0` for small values of x .
 The `expm1()` and `log1p(3M)` functions are useful for financial calculations of $((1+x)^n-1)/x$, namely:

$$\text{expm1}(n * \text{log1p}(x)) / x$$
 when x is very small (for example, when performing calculations with a small daily interest rate). These functions also simplify writing accurate inverse hyperbolic functions.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `exp(3M)`, `ilogb(3M)`, `log1p(3M)`, `attributes(5)`

NAME fabs – absolute value function

SYNOPSIS cc [*flag ...*] *file ...* -lm [*library ...*]
#include <math.h>
double **fabs**(double *x*);

DESCRIPTION The `fabs()` function computes the absolute value of *x*, $|x|$.

RETURN VALUES Upon successful completion, `fabs()` returns the absolute value of *x*.
If *x* is NaN, NaN is returned.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO `isnan(3M)`, `attributes(5)`

NAME | floor – floor function

SYNOPSIS | `cc [flag ...] file ... -lm [library ...]`
`#include <math.h>`
`double floor(double x);`

DESCRIPTION | The `floor()` function computes the largest integral value not greater than x .

RETURN VALUES | Upon successful completion, `floor()` returns the largest integral value not greater than x , expressed as a `double`.
 If x is NaN, NaN is returned.
 If x is $\pm\text{Inf}$ or ± 0 , x is returned.

ERRORS | No errors will occur.

USAGE | The integral value returned by `floor()` as a `double` might not be expressible as an `int` or `long int`. The return value should be tested before assigning it to an integer type to avoid the undefined results of an integer overflow.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `ceil(3M)`, `isnan(3M)`, `attributes(5)`

NAME fmod – floating-point remainder value function

SYNOPSIS cc [*flag ...*] *file ...* -lm [*library ...*]
 #include <math.h>
 double **fmod**(double *x*, double *y*);

DESCRIPTION The `fmod()` function returns the floating-point remainder of the division of *x* by *y*.

RETURN VALUES The `fmod()` function returns the value $x - i * y$, for some integer *i* such that, if *y* is non-zero, the result has the same sign as *x* and magnitude less than the magnitude of *y*.

If *x* or *y* is NaN, NaN is returned. If *y* is 0, NaN is returned and `errno` is set to EDOM. If *x* is ±Inf, NaN is returned. If *y* is non-zero, `fmod(±0, y)` returns the value of *x*. If *x* is not ±Inf, `fmod(x, ±Inf)` returns the value of *x*.

ERRORS The `fmod()` function may fail if:
 EDOM *y* is 0.
 No other errors will occur.

USAGE Portable applications should not call `fmod()` with *y* equal to 0, because the result is implementation-dependent. The application should verify *y* is non-zero before calling `fmod()`.

An application wishing to check for error situations should set `errno` to 0 before calling `fmod()`. If `errno` is non-zero on return, or the return value is NaN, an error has occurred.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO `isnan(3M)`, `attributes(5)`

NAME | freeDmiString – free dynamic memory allocated for input DmiString structure

SYNOPSIS | `cc [flag ...] file ... -ldmi -lnsl -lrwtool [library ...]`
`#include <dmi/util.hh>`
`void freeDmiString(DmiString_t *dstr);`

DESCRIPTION | The `freeDmiString()` function frees dynamic memory allocated for the input DmiString structure.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | MT-Safe |

SEE ALSO | `newDmiString(3DMI)`, `libdmi(3LIB)`, `attributes(5)`

NAME | gelf, gelf_checksum, gelf_fsize, gelf_getclass, gelf_getdyn, gelf_getehdr, gelf_getphdr, gelf_getrel, gelf_getrela, gelf_getshdr, gelf_getsym, gelf_getsyminfo, gelf_newehdr, gelf_newphdr, gelf_update_dyn, gelf_update_ehdr, gelf_update_phdr, gelf_update_rel, gelf_update_rela, gelf_update_shdr, gelf_update_sym, gelf_update_syminfo, gelf_xlatetof, gelf_xlatetom – generic class-independent ELF interface

SYNOPSIS

```
cc [flag ...] file ... -lelf [library ...]

#include <gelf.h>

long gelf_checksum(Elf *elf);
int gelf_getclass(Elf *elf);
size_t gelf_fsize(Elf *elf, Elf_Type type, size_t cnt, unsigned ver);
GElf_Ehdr *gelf_getehdr(Elf *elf, GElf_Ehdr *dst);
int gelf_update_ehdr(Elf *elf, GElf_Ehdr *src);
unsigned long gelf_newehdr(Elf *elf, int class);
GElf_Phdr *gelf_getphdr(Elf *elf, int ndx, GElf_Phdr *src);
int gelf_update_phdr(Elf *elf, int ndx, GElf_Phdr *src);
unsigned long gelf_newphdr(Elf *elf, size_t phnum);
GElf_Shdr *gelf_getshdr(Elf_Scn *scn, Elf_Data *dst);
int gelf_update_shdr(Elf_Scn *scn, GElf_Shdr *src);
Elf_Data *gelf_xlatetof(Elf *elf, Elf_Data *dst, const Elf_Data *src, unsigned encode);
Elf_Data *gelf_xlatetom(Elf *elf, Elf_Data *dst, const Elf_Data *src, unsigned encode);
GElf_Sym *gelf_getsym(Elf_Data *data, int ndx, GElf_Sym *dst);
int gelf_update_sym(Elf_Data *dest, int ndx, GElf_Sym *src);
GElf_Dyn *gelf_getdyn(Elf_Data *src, int ndx, GElf_Dyn *src);
int gelf_update_dyn(Elf_Data *src, int ndx, GElf_Dyn *src);
GElf_Rela *gelf_getrela(Elf_Data *src, int ndx, GElf_Rela *dst);
int gelf_update_rela(Elf_Data *dst, int ndx, GElf_Rela *src);
GElf_Rel *gelf_getrel(Elf_Data *src, int ndx, GElf_Rel *dst);
int gelf_update_rel(Elf_Data *dst, int ndx, GElf_Rel *src);
GElf_Syminfo *gelf_getsyminfo(Elf_Data *src, int ndx, GElf_Syminfo *dst);
```

```
int gelf_update_syminfo(Elf_Data *dst, int ndx, GElf_Syminfo *src);
```

```
GElf_Move *gelf_getmove(Elf_Data *src, int ndx, GElf_Move *dst);
```

```
int gelf_update_move(Elf_Data *dst, int ndx, GElf_Move *src);
```

DESCRIPTION

GELF is a generic, ELF class-independent API, for manipulating ELF object files. GELF provides a single, common interface for handling 32-bit and 64-bit ELF format object files. GELF is a translation layer between the application and the class-dependent parts of the ELF library. Thus, the application can use GELF, which in turn, will call the corresponding elf32_ or elf64_ functions on behalf of the application. The data structures returned are all large enough to hold 32-bit and 64-bit data.

GELF provides a simple, class-independent layer of indirection over the class-dependent ELF32 and ELF64 APIs. GELF is stateless, and may be used along side the ELF32 and ELF64 APIs.

GELF always returns a copy of the underlying ELF32 or ELF64 structure, and therefore the programming practice of using the address of an ELF header as the base offset for the ELF's mapping into memory should be avoided. Also, data accessed by type-casting the Elf_Data buffer to a class-dependent type and treating it like an array, for example, a symbol table, will not work under GELF, and the gelf_get functions must be used instead. See the EXAMPLE section.

Programs which create or modify ELF files using libelf(3LIB) need to perform an extra step when using GELF. Modifications to GELF values must be explicitly flushed to the underlying ELF32 or ELF64 structures by way of the the gelf_update_ interfaces. Use of elf_update or elf_flagelf and the like remains the same.

The sizes of versioning structures remains the same between ELF32 and ELF64. The GELF API only defines types for versioning, rather than a functional API. The processing of versioning information will stay the same in the GELF environment as it was in the class-dependent ELF environment.

List of Functions

| | |
|------------------------------|---|
| <code>gelf_checksum()</code> | An analog to <code>elf32_checksum(3ELF)</code> and <code>elf64_checksum(3ELF)</code> . |
| <code>gelf_getclass()</code> | Returns one of the constants <code>ELFCLASS32</code> , <code>ELFCLASS64</code> or <code>ELFCLASSNONE</code> . |
| <code>gelf_fsize()</code> | An analog to <code>elf32_fsize(3ELF)</code> and <code>elf64_fsize(3ELF)</code> . |
| <code>gelf_getehdr()</code> | An analog to <code>elf32_getehdr(3ELF)</code> and <code>elf64_getehdr(3ELF)</code> . |

| | |
|----------------------------------|--|
| <code>gelf_update_ehdr()</code> | Copies the contents of the <code>GElf_Ehdr</code> ELF header to the underlying <code>Elf32_Ehdr</code> or <code>Elf64_Ehdr</code> structure. |
| <code>gelf_newehdr()</code> | An analog to <code>elf32_newehdr(3ELF)</code> and <code>elf64_newehdr(3ELF)</code> . |
| <code>gelf_getphdr()</code> | An analog to <code>elf32_getphdr(3ELF)</code> and <code>elf64_getphdr(3ELF)</code> . |
| <code>gelf_update_phdr()</code> | Copies of the contents of <code>GElf_Phdr</code> program header to underlying the <code>Elf32_Phdr</code> or <code>Elf64_Phdr</code> structure. |
| <code>gelf_newphdr()</code> | An analog to <code>elf32_newphdr(3ELF)</code> and <code>elf64_newphdr(3ELF)</code> . |
| <code>gelf_getshdr()</code> | An analog to <code>elf32_getshdr(3ELF)</code> and <code>elf64_getshdr(3ELF)</code> . |
| <code>gelf_update_shdr()</code> | Copies of the contents of <code>GElf_Shdr</code> section header to underlying the <code>Elf32_Shdr</code> or <code>Elf64_Shdr</code> structure. |
| <code>gelf_xlatetof()</code> | An analog to <code>elf32_xlatetof(3ELF)</code> and <code>elf64_xlatetof(3ELF)</code> . |
| <code>gelf_xlatetom()</code> | An analog to <code>elf32_xlatetom(3ELF)</code> and <code>elf64_xlatetom(3ELF)</code> . |
| <code>gelf_getsym()</code> | Retrieves the <code>Elf32_Sym</code> or <code>Elf64_Sym</code> information from the symbol table at the given index. |
| <code>gelf_update_sym()</code> | Copies the <code>GElf_Sym</code> information back into the underlying <code>Elf32_Sym</code> or <code>Elf64_Sym</code> structure at the given index. |
| <code>gelf_getdyn()</code> | Retrieves the <code>Elf32_Dyn</code> or <code>Elf64_Dyn</code> information from the dynamic table at the given index. |
| <code>gelf_update_dyn()</code> | Copies the <code>GElf_Dyn</code> information back into the underlying <code>Elf32_Dyn</code> or <code>Elf64_Dyn</code> structure at the given index. |
| <code>gelf_getrela()</code> | Retrieves the <code>Elf32_Rela</code> or <code>Elf64_Rela</code> information from the relocation table at the given index. |

| | |
|------------------------------------|--|
| <code>gelf_update_rela()</code> | Copies the <code>GElf_Refa</code> information back into the underlying <code>Elf32_Refa</code> or <code>Elf64_Refa</code> structure at the given index. |
| <code>gelf_getrel()</code> | Retrieves the <code>Elf32_Rel</code> or <code>Elf64_Rel</code> information from the relocation table at the given index. |
| <code>gelf_update_rel()</code> | Copies the <code>GElf_Rel</code> information back into the underlying <code>Elf32_Rel</code> or <code>Elf64_Rel</code> structure at the given index. |
| <code>gelf_getsyminfo()</code> | Retrieves the <code>Elf32_Syminfo</code> or <code>Elf64_Syminfo</code> information from the relocation table at the given index. |
| <code>gelf_update_syminfo()</code> | Copies the <code>GElf_Syminfo</code> information back into the underlying <code>Elf32_Syminfo</code> or <code>Elf64_Syminfo</code> structure at the given index. |
| <code>gelf_getmove()</code> | Retrieves the <code>Elf32_Move</code> or <code>Elf64_Move</code> information from the move table at the given index. |
| <code>gelf_update_move()</code> | Copies the <code>GElf_Move</code> information back into the underlying <code>Elf32_Move</code> or <code>Elf64_Move</code> structure at the given index. |

RETURN VALUES

Upon failure, all `GElf` functions return 0 and set `elf_errno`. See `elf_errno(3ELF)`.

EXAMPLES**EXAMPLE 1** Printing the ELF Symbol Table

```
#include <stdio.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <libelf.h>
#include <gelf.h>

void
main(int argc, char **argv)
{
    Elf          *elf;
    Elf_Scn      *scn = NULL;
    GElf_Shdr    shdr;
    Elf_Data     *data;
    int          fd, ii, count;

    elf_version(EV_CURRENT);

    fd = open(argv[1], O_RDONLY);
```

```

elf = elf_begin(fd, ELF_C_READ, NULL);

while ((scn = elf_nextscn(elf, scn)) != NULL) {
    gelf_getshdr(scn, &shdr);
    if (shdr.sh_type == SHT_SYMTAB) {
        /* found a symbol table, go print it. */
        break;
    }
}

data      = elf_getdata(scn, NULL);
count     = shdr.sh_size / shdr.sh_entsize;

/* print the symbol names */
for (ii=0; ii < count; ++ii) {
    GElf_Sym sym;
    gelf_getsym(data, ii, &sym);
    printf("%s\
", elf_strptr(elf, shdr.sh_link, sym.st_name));
}
elf_end(elf);
close(fd);
}
    
```

- FILES**
- /usr/lib/libelf.so.1 Shared object.
 - /usr/lib/sparcv9/libelf.so.1 64-bit shared object.
 - /usr/lib//libelf.a Archive library.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT Level | Safe |

SEE ALSO elf(3ELF), elf32_checksum(3ELF), elf32_fsize(3ELF),
 elf32_getehdr(3ELF), elf32_newehdr(3ELF), elf32_getphdr(3ELF),
 elf32_newphdr(3ELF), elf32_getshdr(3ELF), elf32_xlatetof(3ELF),
 elf32_xlatetom(3ELF), elf_errno(3ELF), libelf(3LIB), attributes(5)

| | |
|----------------------|--|
| NAME | getacinfo, getacdir, getacflg, getacmin, getacna, setac, endac – get audit control file information |
| SYNOPSIS | <pre>cc [flag ...] file ... -lbsm -lsocket -lnsl -lintl [library ...] #include <bsm/libbsm.h> int getacdir(char *dir, int len); int getacmin(int *min_val); int getacflg(char *auditstring, int len); int getacna(char *auditstring, int len); void setac(void); void endac(void);</pre> |
| DESCRIPTION | <p>When first called, <code>getacdir()</code> provides information about the first audit directory in the <code>audit_control</code> file; thereafter, it returns the next directory in the file. Successive calls list all the directories listed in <code>audit_control(4)</code>. The parameter <code>len</code> specifies the length of the buffer <code>dir</code>. On return, <code>dir</code> points to the directory entry.</p> <p><code>getacmin()</code> reads the minimum value from the <code>audit_control</code> file and returns the value in <code>min_val</code>. The minimum value specifies how full the file system to which the audit files are being written can get before the script <code>audit_warn(1M)</code> is invoked.</p> <p><code>getacflg()</code> reads the system audit value from the <code>audit_control</code> file and returns the value in <code>auditstring</code>. The parameter <code>len</code> specifies the length of the buffer <code>auditstring</code>.</p> <p><code>getacna()</code> reads the system audit value for non-attributable audit events from the <code>audit_control</code> file and returns the value in <code>auditstring</code>. The parameter <code>len</code> specifies the length of the buffer <code>auditstring</code>. Non-attributable events are events that cannot be attributed to an individual user. <code>inetd(1M)</code> and several other daemons record non-attributable events.</p> <p>Calling <code>setac</code> rewinds the <code>audit_control</code> file to allow repeated searches.</p> <p>Calling <code>endac</code> closes the <code>audit_control</code> file when processing is complete.</p> |
| FILES | <pre>/etc/security/audit_control contains default parameters read by the audit daemon, auditd(1M)</pre> |
| RETURN VALUES | <p><code>getacdir()</code>, <code>getacflg()</code>, <code>getacna()</code> and <code>getacmin()</code> return:</p> <p>0 on success.</p> <p>-2 on failure and set <code>errno</code> to indicate the error.</p> <p><code>getacmin()</code> and <code>getacflg()</code> return:</p> |

1 on EOF.

getacdir() returns:

-1 on EOF.

2 if the directory search had to start from the beginning because one of the other functions was called between calls to getacdir().

These functions return:

-3 if the directory entry format in the audit_control file is incorrect.

getacdir(), getacflg() and getacna() return:

-3 if the input buffer is too short to accommodate the record.

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Safe. |

SEE ALSO

audit_warn(1M), bsmconv(1M), inetd(1M), audit_control(4), attributes(5)

NOTES

The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See bsmconv(1M) for more information.

| | |
|--------------------|--|
| NAME | getauclassent, getauclassnam, setauclass, endauclass, getauclassnam_r, getauclassent_r – get audit_class entry |
| SYNOPSIS | <pre>cc [flag ...] file ... -lbsm -lsocket -lnsl -lintl [library ...] #include <sys/param.h> #include <bsm/libbsm.h> struct au_class_ent *getauclassnam(const char *name); struct au_class_ent *getauclassnam_r(au_class_ent_t *class_int, const char *name); struct au_class_ent *getauclassent(void); struct au_class_ent *getauclassent_r(au_class_ent_t *class_int); void setauclass(void); void endauclass(void);</pre> |
| DESCRIPTION | <p>getauclassent() and getauclassnam() each return an audit_class entry.</p> <p>getauclassnam() searches for an audit_class entry with a given class name <i>name</i>.</p> <p>getauclassent() enumerates audit_class entries: successive calls to getauclassent() will return either successive audit_class entries or NULL.</p> <p>setauclass() “rewinds” to the beginning of the enumeration of audit_class entries. Calls to getauclassnam() may leave the enumeration in an indeterminate state, so setauclass() should be called before the first getauclassent() .</p> <p>endauclass() may be called to indicate that audit_class processing is complete; the system may then close any open audit_class file, deallocate storage, and so forth.</p> <p>getauclassent_r() and getauclassnam_r() both return a pointer to an audit_class entry as do their similarly named counterparts. They each take an additional argument, a pointer to pre-allocated space for an au_class_ent_t, which is returned if the call is successful. To assure there is enough space for the information returned, the applications programmer should be sure to allocate AU_CLASS_NAME_MAX and AU_CLASS_DESC_MAX bytes for the ac_name and ac_desc elements of the au_class_ent_t data structure.</p> <p>The internal representation of an audit_user entry is an au_class_ent structure defined in <bsm/libbsm.h> with the following members:</p> <pre>char *ac_name; au_class_t ac_class; char *ac_desc;</pre> |

RETURN VALUES

getauclassnam() and getauclassnam_r() return a pointer to a struct au_class_ent if they successfully locate the requested entry; otherwise they return NULL.

getauclassent() and getauclassent_r() return a pointer to a struct au_class_ent if they successfully enumerate an entry; otherwise they return NULL, indicating the end of the enumeration.

FILES

/etc/security/audit_class Maps audit class numbers to audit class names

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|--------------------------|
| MT-Level | MT-Safe with exceptions. |

All of the functions described in this man-page are MT-Safe except getauclassent() and getauclassnam . The two functions, getauclassent_r() and getauclassnam_r() have the same functionality as the unsafe functions, but have a slightly different function call interface in order to make them MT-Safe.

SEE ALSO

bsmconv(1M) , audit_class(4) , audit_event(4) , attributes(5)

NOTES

All information is contained in a static area, so it must be copied if it is to be saved.

The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See bsmconv(1M) for more information.

| NAME | getauditflags, getauditflagsbin, getauditflagschar – convert audit flag specifications | | | | |
|----------------------|---|----------------|-----------------|----------|----------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lbsm -lsocket -lnsl -lintl [library ...] #include <sys/param.h> #include <bsm/libbsm.h> int getauditflagsbin(char *auditstring, au_mask_t *masks); int getauditflagschar(char *auditstring, au_mask_t *masks, int verbose);</pre> | | | | |
| DESCRIPTION | <p>getauditflagsbin() converts the character representation of audit values pointed to by <i>auditstring</i> into <i>au_mask_t</i> fields pointed to by <i>masks</i> . These fields indicate which events are to be audited when they succeed and which are to be audited when they fail. The character string syntax is described in <i>audit_control(4)</i> .</p> <p>getauditflagschar() converts the <i>au_mask_t</i> fields pointed to by <i>masks</i> into a string pointed to by <i>auditstring</i> . If <i>verbose</i> is zero, the short (2-character) flag names are used. If <i>verbose</i> is non-zero, the long flag names are used. <i>auditstring</i> should be large enough to contain the ASCII representation of the events.</p> <p><i>auditstring</i> contains a series of event names, each one identifying a single audit class, separated by commas. The <i>au_mask_t</i> fields pointed to by <i>masks</i> correspond to binary values defined in <i><bsm/audit.h></i> , which is read by <i><bsm/libbsm.h></i> .</p> | | | | |
| RETURN VALUES | getauditflagsbin() and getauditflagschar() : -1 is returned on error and 0 on success. | | | | |
| ATTRIBUTES | See <i>attributes(5)</i> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe.</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe. |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe. | | | | |
| SEE ALSO | <i>bsmconv(1M)</i> , <i>audit.log(4)</i> , <i>audit_control(4)</i> , <i>attributes(5)</i> | | | | |
| BUGS | This is not a very extensible interface. | | | | |
| NOTES | The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See <i>bsmconv(1M)</i> for more information. | | | | |

| | |
|--------------------|---|
| NAME | getauevent, getauevnam, getauevnum, getauevnonam, setauevent, endaeuevent, getauevent_r, getauevnam_r, getauevnum_r – get audit_event entry |
| SYNOPSIS | <pre>cc [flag ...] file ... -lbsm -lsocket -lnsl -lintl [library ...] #include <sys/param.h> #include <bsm/libbsm.h> struct au_event_ent *getauevent(void); struct au_event_ent *getauevnam(char *name); struct au_event_ent *getauevnum(au_event_t event_number); au_event_t *getauevnonam(char *event_name); void setauevent(void); void endaeuevent(void); struct au_event_ent *getauevent_r(au_event_ent_t *e); struct au_event_ent *getauevnam_r(au_event_ent_t *e, char *name); struct au_event_ent *getauevnum_r(au_event_ent_t *e, au_event_t event_number);</pre> |
| DESCRIPTION | <p>These interfaces document the programming interface for obtaining entries from the audit_event(4) file. getauevent(), getauevnam(), getauevnum(), getauevent_r(), getauevnam_r(), and getauevnum_r() each return a pointer to an audit_event structure.</p> <p>getauevent() and getauevent_r() enumerate audit_event entries; successive calls to these functions will return either successive audit_event entries or NULL.</p> <p>getauevnam() and getauevnam_r() search for an audit_event entry with a given event_name.</p> <p>getauevnum() and getauevnum_r() search for an audit_event entry with a given event_number.</p> <p>getauevnonam() searches for an audit_event entry with a given event_name and returns the corresponding event number.</p> <p>setauevent() “rewinds” to the beginning of the enumeration of audit_event entries. Calls to getauevnam(), getauevnum(), getauevnonum(), getauevnam_r(), or getauevnum_r() may leave the enumeration in an indeterminate state; setauevent() should be called before the first getauevent() or getauevent_r().</p> <p>endaeuevent() may be called to indicate that audit_event processing is complete; the system may then close any open audit_event file, deallocate storage, and so forth.</p> |

The three functions `getauevent_r()`, `getauevnam_r()`, and `getauevnum_r()` each take an argument `e` which is a pointer to an `au_event_ent_t`. This pointer is returned on a successful function call. To assure there is enough space for the information returned, the applications programmer should be sure to allocate `AU_EVENT_NAME_MAX` and `AU_EVENT_DESC_MAX` bytes for the `ae_name` and `ae_desc` elements of the `au_event_ent_t` data structure.

The internal representation of an `audit_event` entry is an `struct au_event_ent` structure defined in `<bsm/libbsm.h>` with the following members:

```

au_event_t      ae_number
char            *ae_name;
char            *ae_desc*;
au_class_t      ae_class;
    
```

RETURN VALUES

`getauevent()`, `getauevnam()`, `getauevnum()`, `getauevent_r()`, `getauevnam_r()`, and `getauevnum_r()` return a pointer to a `struct au_event_ent` if the requested entry is successfully located; otherwise it returns `NULL`.

`getauevnonam()` returns an event number of type `au_event_t` if it successfully enumerates an entry; otherwise it returns `NULL`, indicating it could not find the requested event name.

FILES

| | |
|--|--|
| <code>/etc/security/audit_event</code> | Maps audit event numbers to audit event names. |
| <code>/etc/passwd</code> | Stores user-ID to username mappings. |

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|--------------------------|
| MT-Level | MT-Safe with exceptions. |

The functions `getauevent()`, `getauevnam()`, and `getauevnum()` are not MT-Safe; however, there are equivalent functions: `getauevent_r()`, `getauevnam_r()`, and `getauevnum_r()` - all of which provide the same functionality and a MT-Safe function call interface.

SEE ALSO

`bsmconv(1M)`, `getauevnam(3BSM)`, `getpwnam(3C)`, `audit_class(4)`, `audit_event(4)`, `passwd(4)`, `attributes(5)`

NOTES

All information for the functions `getauevent()`, `getauevnam()`, and `getauevnum()` is contained in a static area, so it must be copied if it is to be saved.

The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See `bsmconv(1M)` for more information.

| | |
|--------------------|--|
| NAME | getauthattr, getauthnam, free_authattr, setauthattr, endauthattr, chkauthattr – get authorization entry |
| SYNOPSIS | <pre>cc [flag...] file... -lsecdb -lsocket -lnsl -lintl [library...] #include <auth_attr.h> #include <secdb.h> authattr_t *getauthattr(void); authattr_t *getauthnam(const char *name); void free_authattr(authattr_t *auth); void setauthattr(void); void endauthattr(void); int chkauthattr(const char *authname, const char *username);</pre> |
| DESCRIPTION | <p>The <code>getauthattr()</code> and <code>getauthnam()</code> functions each return an <code>auth_attr(4)</code> entry. Entries can come from any of the sources specified in the <code>nsswitch.conf(4)</code> file.</p> <p>The <code>getauthattr()</code> function enumerates <code>auth_attr</code> entries. The <code>getauthnam()</code> function searches for an <code>auth_attr</code> entry with a given authorization name <code>name</code>. Successive calls to these functions return either successive <code>auth_attr</code> entries or <code>NULL</code>.</p> <p>The internal representation of an <code>auth_attr</code> entry is an <code>authattr_t</code> structure defined in <code><auth_attr.h></code> with the following members:</p> <pre>char *name; /* name of the authorization */ char *res1; /* reserved for future use */ char *res2; /* reserved for future use */ char *short_desc; /* short description */ char *long_desc; /* long description */ kva_t *attr; /* array of key-value pair attributes */</pre> <p>The <code>setauthattr()</code> function "rewinds" to the beginning of the enumeration of <code>auth_attr</code> entries. Calls to <code>getauthnam()</code> can leave the enumeration in an indeterminate state. Therefore, <code>setauthattr()</code> should be called before the first call to <code>getauthattr()</code>.</p> <p>The <code>endauthattr()</code> function may be called to indicate that <code>auth_attr</code> processing is complete; the system may then close any open <code>auth_attr</code> file, deallocate storage, and so forth.</p> <p>The <code>chkauthattr()</code> function verifies whether or not a user has a given authorization. It first reads the <code>AUTHS_GRANTED</code> key in the <code>/etc/security/policy.conf</code> file and returns 1 if it finds a match for the given authorization. If <code>chkauthattr()</code> does not find a match, it reads the <code>user_attr(4)</code> database. If it does not find a match in <code>user_attr</code>,</p> |

`chkauthattr()` reads the `prof_attr(4)` database, using the list of profiles assigned to the user, and checks if any of the profiles assigned to the user has the given authorization. The `chkauthattr()` function returns 0 if it does not find a match in any of the three sources.

A user is considered to have been assigned an authorization if either of the following are true:

- The authorization name matches exactly any authorization assigned in the `user_attr` or `prof_attr` databases (authorization names are case-sensitive).
- The authorization name suffix is not the key word `grant` and the authorization name matches any authorization up to the asterisk (*) character assigned in the `user_attr` or `prof_attr` databases.

The examples in the following table illustrate the conditions under which a user is assigned an authorization.

| | <code>/etc/security/policy.conf</code> or | Is user |
|---|--|--------------------|
| Authorization name | <code>user_attr</code> or <code>prof_attr</code> entry | authorized? |
| <code>com.sun.printer.postscript</code> | <code>com.sun.printer.postscript</code> | Yes |
| <code>com.sun.printer.postscript</code> | <code>com.sun.printer.*</code> | Yes |
| <code>com.sun.printer.grant</code> | <code>com.sun.printer.*</code> | No |

The `free_authattr()` function releases memory allocated by the `getauthnam()` and `getauthattr()` functions.

RETURN VALUES

The `getauthattr()` function returns a pointer to an `authattr_t` if it successfully enumerates an entry; otherwise it returns `NULL`, indicating the end of the enumeration.

The `getauthnam()` function returns a pointer to an `authattr_t` if it successfully locates the requested entry; otherwise it returns `NULL`.

The `chkauthattr()` function returns 1 if the user is authorized and 0 otherwise.

USAGE

The `getauthattr()` and `getauthnam()` functions both allocate memory for the pointers they return. This memory should be de-allocated with the `free_authattr()` call.

Applications that use the interfaces described in this manual page cannot be linked statically, since the implementations of these functions employ dynamic

loading and linking of shared objects at run time. Note that these interfaces are reentrant even though they do not use the `_r` suffix naming convention.

Individual attributes in the `attr` structure can be referred to by calling the `kva_match(3SECDB)` function.

WARNINGS

Because the list of legal keys is likely to expand, code must be written to ignore unknown key-value pairs without error.

FILES

| | |
|--|--|
| <code>/etc/nsswitch.conf</code> | configuration file lookup information for the name server switch |
| <code>/etc/user_attr</code> | extended user attributes |
| <code>/etc/security/auth_attr</code> | authorization attributes |
| <code>/etc/security/policy.conf</code> | policy definitions |
| <code>/etc/security/prof_attr</code> | profile information |

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`getexecattr(3SECDB)`, `getprofattr(3SECDB)`, `getuserattr(3SECDB)`, `auth_attr(4)`, `nsswitch.conf(4)`, `prof_attr(4)`, `user_attr(4)`, `attributes(5)`, `rbac(5)`

| | |
|--------------------|--|
| NAME | getausernam, getauserent, setauser, endauser – get audit_user entry |
| SYNOPSIS | <pre>cc [flag ...] file ... -lbsm -lsocket -lnsl -lintl [library ...] #include <sys/param.h> #include <bsm/libbsm.h> struct au_user_ent *getausernam(const char *name); struct au_user_ent *getauserent(void); void setauser(void); void endauser(void); struct au_user_ent *getausernam_r(au_user_ent_t *u, const char *name); struct au_user_ent *getauserent_r(au_user_ent_t *u);</pre> |
| DESCRIPTION | <p>The <code>getauserent()</code>, <code>getausernam()</code>, <code>getauserent_r()</code>, and <code>getausernam_r()</code> functions each return an <code>audit_user</code> entry. Entries can come from any of the sources specified in the <code>/etc/nsswitch.conf</code> file (see <code>nsswitch.conf(4)</code>).</p> <p>The <code>getausernam()</code> and <code>getausernam_r()</code> functions search for an <code>audit_user</code> entry with a given login name <code>name</code>.</p> <p>The <code>getauserent()</code> and <code>getauserent_r()</code> functions enumerate <code>audit_user</code> entries; successive calls to these functions will return either successive <code>audit_user</code> entries or <code>NULL</code>.</p> <p>The <code>setauser()</code> function "rewinds" to the beginning of the enumeration of <code>audit_user</code> entries. Calls to <code>getausernam()</code> and <code>getausernam_r()</code> may leave the enumeration in an indeterminate state, so <code>setauser()</code> should be called before the first call to <code>getauserent()</code> or <code>getauserent_r()</code>.</p> <p>The <code>endauser()</code> function may be called to indicate that <code>audit_user</code> processing is complete; the system may then close any open <code>audit_user</code> file, deallocate storage, and so forth.</p> <p>The <code>getauserent_r()</code> and <code>getausernam_r()</code> functions both take an argument <code>u</code>, which is a pointer to an <code>au_user_ent</code>. This is the pointer that is returned on successful function calls.</p> <p>The internal representation of an <code>audit_user</code> entry is an <code>au_user_ent</code> structure defined in <code><bsm/libbsm.h></code> with the following members:</p> <pre>char *au_name; au_mask_t au_always; au_mask_t au_never;</pre> |

RETURN VALUES

The `getauusernam()` function returns a pointer to a `struct au_user_ent` if it successfully locates the requested entry; otherwise it returns `NULL`.

The `getauuserent()` function returns a pointer to a `struct au_user_ent` if it successfully enumerates an entry; otherwise it returns `NULL`, indicating the end of the enumeration.

USAGE

The functionality described in this manual page is available only if the Basic Security Module (BSM) has been enabled. See `bsmconv(1M)` for more information.

FILES

| | |
|---------------------------------------|-------------------------------------|
| <code>/etc/security/audit_user</code> | stores per-user audit event mask |
| <code>/etc/passwd</code> | stores user-id to username mappings |
| <code>/etc/security/audit_user</code> | stores per-user audit event mask |

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|--------------------------|
| MT-Level | MT-Safe with exceptions. |

SEE ALSO

`bsmconv(1M)`, `getpwnam(3C)`, `audit_user(4)`, `nsswitch.conf(4)`, `passwd(4)`, `attributes(5)`

NOTES

All information for the `getauuserent()` and `getauusernam()` functions is contained in a static area, so it must be copied if it is to be saved.

The `getauusernam()` and `getauuserent()` functions are not MT-safe. The `getauusernam_r()` and `getauuserent_r()` functions provide the same functionality with interfaces that are MT-Safe.

| | |
|--------------------|---|
| NAME | getexecattr, free_execattr, setexecattr, endexecattr, getexecuser, getexecprof, match_execattr – get execution profile entry |
| SYNOPSIS | <pre>cc [flag...] file... -lsecdb -lsocket -lnsl -lintl [library...] #include <exec_attr.h> #include <secdb.h> execattr_t *getexecattr(void); void free_execattr(execattr_t *ep); void setexecattr(void); void endexecattr(void); execattr_t *getexecuser(const char *username, const char *type, const char *id, int search_flag); execattr_t *getexecprof(const char *profname, const char *type, const char *id, int search_flag); execattr_t *match_execattr(execattr_t *ep, char *profname, char *type, char *id);</pre> |
| DESCRIPTION | <p>The <code>getexecattr()</code> function returns a single <code>exec_attr</code> entry. Entries can come from any of the sources specified in the <code>nsswitch.conf(4)</code> file.</p> <p>Successive calls to <code>getexecattr()</code> return either successive <code>exec_attr</code> entries or <code>NULL</code>. Because <code>getexecattr()</code> always returns a single entry, the next pointer in the <code>execattr_t</code> data structure points to <code>NULL</code>.</p> <p>The internal representation of an <code>exec_attr</code> entry is an <code>execattr_t</code> structure defined in <code><exec_attr.h></code> with the following members:</p> <pre>char *name; /* name of the profile */ char *type; /* type of profile */ char *policy; /* policy under which the attributes are */ /* relevant*/ char *res1; /* reserved for future use */ char *res2; /* reserved for future use */ char *id; /* unique identifier */ kva_t *attr; /* attributes */ struct execattr_s *next; /* optional pointer to next profile */</pre> <p>The <code>free_execattr()</code> function releases memory. It follows the <code>next</code> pointers in the <code>execattr_t</code> structure so that the entire linked list is released.</p> <p>The <code>setexecattr()</code> function "rewinds" to the beginning of the enumeration of <code>exec_attr</code> entries. Calls to <code>getexecuser()</code> can leave the enumeration in an indeterminate state. Therefore, <code>setexecattr()</code> should be called before the first call to <code>getexecattr()</code>.</p> |

The `endexecattr()` function can be called to indicate that `exec_attr` processing is complete; the library can then close any open `exec_attr` file, deallocate any internal storage, and so forth.

The `getexecuser()` function returns a linked list of entries filtered by the function's arguments. Only entries assigned to the specified `username`, as described in the `passwd(4)` database, and containing the specified `type` and `id`, as described in the `exec_attr(4)` database, are placed in the list. The `getexecuser()` function is different from the other functions in its family because it spans two databases. It first looks up the list of profiles assigned to a user in the `user_attr` database, then looks up each profile in the `exec_attr` database.

The `getexecprof()` function returns a linked list of entries that have components matching the function's arguments. Only entries in the database matching the argument `profname`, as described in `exec_attr`, and containing the `type` and `id`, also described in `exec_attr`, are placed in the list.

Using `getexecuser()` and `getexecprof()`, programmers can search for any `type` argument, such as the manifest constant `KV_COMMAND`. The arguments are logically AND-ed together so that only entries exactly matching all of the arguments are returned. Wildcard matching applies if there is no exact match for an ID. Any argument can be assigned the `NULL` value to indicate that it is not used as part of the matching criteria. The `search_flag` controls whether the function returns the first match (`GET_ONE`), setting the `next` pointer to `NULL` or all matching entries (`GET_ALL`), using the `next` pointer to create a linked list of all entries that meet the search criteria. See `EXAMPLES`.

Once a list of entries is returned by `getexecuser()` or `getexecprof()`, the convenience function `match_execattr()` can be used to identify an individual entry. It returns a pointer to the individual element with the same profile name (`profname`), type name (`type`), and `id`. Function parameters set to `NULL` are not used as part of the matching criteria. In the event that multiple entries meet the matching criteria, only a pointer to the first entry is returned. The `kva_match(3SECDB)` function can be used to look up a key in a key-value array.

RETURN VALUES

Those functions returning data only return data related to the active policy. The `getexecattr()` function returns a pointer to a `execattr_t` if it successfully enumerates an entry; otherwise it returns `NULL`, indicating the end of the enumeration.

USAGE

The `getexecattr()`, `getexecuser()`, and `getexecprof()` functions all allocate memory for the pointers they return. This memory should be deallocated with the `free_execattr()` call. The `match_execattr()` function does not allocate any memory. Therefore, pointers returned by this function should not be deallocated.

Applications that use the interfaces described in this manual page cannot be linked statically, since the implementations of these functions employ dynamic loading and linking of shared objects at run time. Note that these interfaces are reentrant even though they do not use the `_r` suffix naming convention.

Individual attributes may be referenced in the `attr` structure by calling the `kva_match(3SECDB)` function.

EXAMPLES

EXAMPLE 1 The following finds all profiles that have the `ping` command.

```
if ((execprof=getexecprof(NULL, KV_COMMAND, "/usr/sbin/ping",
    GET_ONE)) == NULL) {
    /* do error */
}
```

EXAMPLE 2 The following finds the entry for the `ping` command in the Network Administration Profile.

```
if ((execprof=getexecprof("Network Administration", KV_COMMAND,
    "/usr/sbin/ping", GET_ALL))==NULL) {
    /* do error */
}
```

EXAMPLE 3 The following tells everything that can be done in the Filesystem Security profile.

```
if ((execprof=getexecprof("Filesystem Security", KV_NULL, NULL,
    GET_ALL))==NULL) {
    /* do error */
}
```

EXAMPLE 4 The following tells if the `tar` command is in a profile assigned to user `wetmore`. If there is no exact profile entry, the wildcard (`*`), if defined, is returned.

```
if ((execprof=getexecuser("wetmore", KV_COMMAND, "/usr/bin/tar",
    GET_ONE))==NULL) {
    /* do error */
}
```

FILES

- `/etc/nsswitch.conf` configuration file lookup information for the name server switch
- `/etc/user_attr` extended user attributes
- `/etc/security/exec_attr` execution profiles

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`getauthattr(3SECDB)`, `getuserattr(3SECDB)`, `kva_match(3SECDB)`, `exec_attr(4)`, `user_attr(4)`, `attributes(5)`

| NAME | getfauditflags – generates the process audit state | | | | |
|----------------------|---|----------------|-----------------|----------|----------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lbsm -lsocket -lnsl -lintl [library ...] #include <sys/param.h> #include <bsm/libbsm.h> int getfauditflags(au_mask_t *usremasks, au_mask_t *usrdmasks, au_mask_t *lastmasks);</pre> | | | | |
| DESCRIPTION | <p>getfauditflags() generates a process audit state by combining the audit masks passed as parameters with the system audit masks specified in the audit_control(4) file. getfauditflags() obtains the system audit value by calling getacflg() (see getacinfo(3BSM)).</p> <p>usremasks points to au_mask_t fields which contains two values. The first value defines which events are <i>always</i> to be audited when they succeed. The second value defines which events are always to be audited when they fail.</p> <p>usrdmasks also points to au_mask_t fields which contains two values. The first value defines which events are <i>never</i> to be audited when they succeed. The second value defines which events are never to be audited when they fail.</p> <p>The structures pointed to by usremasks and usrdmasks may be obtained from the audit_user(4) file by calling getauusernam() which returns a pointer to a structure containing all audit_user(4) fields for a user.</p> <p>The output of this function is stored in lastmasks which is a pointer of type au_mask_t as well. The first value defines which events are to be audited when they succeed and the second defines which events are to be audited when they fail.</p> <p>Both usremasks and usrdmasks override the values in the system audit values.</p> | | | | |
| RETURN VALUES | -1 is returned on error and 0 on success. | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe.</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe. |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe. | | | | |
| SEE ALSO | bsmconv(1M), getacinfo(3BSM), getauditflags(3BSM), getauusernam(3BSM), audit.log(4), audit_control(4), audit_user(4), attributes(5) | | | | |
| NOTES | The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See bsmconv(1M) for more information. | | | | |

| | |
|----------------------|---|
| NAME | getprofattr, getprofnam, free_profattr, setprofattr, endprofattr – get profile description and attributes |
| SYNOPSIS | <pre>cc [flag...] file... -lsecdb -lsocket -lnsl -lintl [library...] #include <prof.h> profattr_t *getprofattr(void); profattr_t *getprofnam(const char * name); void free_profattr(profattr_t *pd); void setprofattr(void); void endprofattr(void);</pre> |
| DESCRIPTION | <p>The <code>getprofattr()</code> and <code>getprofnam()</code> functions each return a <code>prof_attr</code> entry. Entries can come from any of the sources specified in the <code>nsswitch.conf(4)</code> file.</p> <p>The <code>getprofattr()</code> function enumerates <code>prof_attr</code> entries. The <code>getprofnam()</code> function searches for a <code>prof_attr</code> entry with a given <code>name</code>. Successive calls to these functions return either successive <code>prof_attr</code> entries or <code>NULL</code>.</p> <p>The internal representation of a <code>prof_attr</code> entry is a <code>profattr_t</code> structure defined in <code><prof_attr.h></code> with the following members:</p> <pre>char *name; /* Name of the profile */ char *res1; /* Reserved for future use */ char *res2; /* Reserved for future use */ char *desc; /* Description/Purpose of the profile */ kva_t *attr; /* Profile attributes */</pre> <p>The <code>free_profattr()</code> function releases memory allocated by the <code>getprofattr()</code> and <code>getprofnam()</code> functions.</p> <p>The <code>setprofattr()</code> function "rewinds" to the beginning of the enumeration of <code>prof_attr</code> entries. Calls to <code>getprofnam()</code> can leave the enumeration in an indeterminate state. Therefore, <code>setprofattr()</code> should be called before the first call to <code>getprofattr()</code>.</p> <p>The <code>endprofattr()</code> function may be called to indicate that <code>prof_attr</code> processing is complete; the system may then close any open <code>prof_attr</code> file, deallocate storage, and so forth.</p> |
| RETURN VALUES | <p>The <code>getprofattr()</code> function returns a pointer to a <code>profattr_t</code> if it successfully enumerates an entry; otherwise it returns <code>NULL</code>, indicating the end of the enumeration.</p> <p>The <code>getprofnam()</code> function returns a pointer to a <code>profattr_t</code> if it successfully locates the requested entry; otherwise it returns <code>NULL</code>.</p> |

USAGE Individual attributes in the `prof_attr_t` structure can be referred to by calling the `kva_match(3SECDB)` function.

Because the list of legal keys is likely to expand, any code must be written to ignore unknown key-value pairs without error.

The `getprofattr()` and `getprofnam()` functions both allocate memory for the pointers they return. This memory should be deallocated with the `free_profatt()` function.

Applications that use the interfaces described in this manual page cannot be linked statically, since the implementations of these functions employ dynamic loading and linking of shared objects at run time. Note that these interfaces are reentrant even though they do not use the `_r` suffix naming convention.

FILES `/etc/security/prof_attr` profiles and their descriptions

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO `auths(1)`, `profiles(1)`, `getexecattr(3SECDB)`, `getauthattr(3SECDB)`, `prof_attr(4)`

| | |
|----------------------|--|
| NAME | getuserattr, getusernam, getuseruid, free_userattr, setuserattr, enduserattr – get user_attr entry |
| SYNOPSIS | <pre>cc [flag...] file...- lsecdb - lsocket - lnsl - lintl [library...] #include <user_attr.h> userattr_t *getuserattr(void); userattr_t *getusernam(const char * name); userattr_t *getuseruid(uid_t uid); void free_userattr(userattr_t *userattr); void setuserattr(void); void enduserattr(void);</pre> |
| DESCRIPTION | <p>The <code>getuserattr()</code>, <code>getusernam()</code>, and <code>getuseruid()</code> functions each return a <code>user_attr(4)</code> entry. Entries can come from any of the sources specified in the <code>nsswitch.conf(4)</code> file. The <code>getuserattr()</code> function enumerates <code>user_attr</code> entries. The <code>getusernam()</code> function searches for a <code>user_attr</code> entry with a given user name <code>name</code>. The <code>getuseruid()</code> function searches for a <code>user_attr</code> entry with a given user id <code>uid</code>. Successive calls to these functions return either successive <code>user_attr</code> entries or <code>NULL</code>.</p> <p>The <code>free_userattr()</code> function releases memory allocated by the <code>getusernam()</code> and <code>getuserattr()</code> functions.</p> <p>The internal representation of a <code>user_attr</code> entry is a <code>userattr_t</code> structure defined in <code><user_attr.h></code> with the following members:</p> <pre>char *name; /* name of the user */ char *qualifier; /* reserved for future use */ char *res1; /* reserved for future use */ char *res2; /* reserved for future use */ kva_t *attr; /* list of attributes */</pre> <p>The <code>setuserattr()</code> function "rewinds" to the beginning of the enumeration of <code>user_attr</code> entries. Calls to <code>getusernam()</code> may leave the enumeration in an indeterminate state, so <code>setuserattr()</code> should be called before the first call to <code>getuserattr()</code>.</p> <p>The <code>enduserattr()</code> function may be called to indicate that <code>user_attr</code> processing is complete; the library may then close any open <code>user_attr</code> file, deallocate any internal storage, and so forth.</p> |
| RETURN VALUES | <p>The <code>getuserattr()</code> function returns a pointer to a <code>userattr_t</code> if it successfully enumerates an entry; otherwise it returns <code>NULL</code>, indicating the end of the enumeration.</p> <p>The <code>getusernam()</code> function returns a pointer to a <code>userattr_t</code> if it successfully locates the requested entry; otherwise it returns <code>NULL</code>.</p> |

USAGE The `getuserattr()` and `getusernam()` functions both allocate memory for the pointers they return. This memory should be deallocated with the `free_userattr()` function.

Applications that use the interfaces described in this manual page cannot be linked statically, since the implementations of these functions employ dynamic loading and linking of shared objects at run time. Note that these interfaces are reentrant even though they do not use the `_r` suffix naming convention.

Individual attributes may be referenced in the `attr` structure by calling the `kva_match(3SECDB)` function.

WARNINGS Because the list of legal keys is likely to expand, code must be written to ignore unknown key-value pairs without error.

FILES

| | |
|---------------------------------|--|
| <code>/etc/user_attr</code> | extended user attributes |
| <code>/etc/nsswitch.conf</code> | configuration file lookup information for the name server switch |

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO `getauthattr(3SECDB)`, `getexecattr(3SECDB)`, `getprofattr(3SECDB)`, `user_attr(4)`, `attributes(5)`

| NAME | gmatch – shell global pattern matching | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lgen [<i>library ...</i>] #include <libgen.h> int gmatch (const char * <i>str</i> , const char * <i>pattern</i>); | | | | |
| DESCRIPTION | gmatch () checks whether the null-terminated string <i>str</i> matches the null-terminated pattern string <i>pattern</i> . See the sh (1), section File Name Generation, for a discussion of pattern matching. A backslash (\) is used as an escape character in pattern strings. | | | | |
| RETURN VALUES | gmatch () returns non-zero if the pattern matches the string, zero if the pattern does not. | | | | |
| EXAMPLES | EXAMPLE 1 Examples of gmatch () function. In the following example, gmatch () returns non-zero (true) for all strings with “a” or “-” as their last character. <pre>char *s; gmatch (s, "[a\-")</pre> | | | | |
| ATTRIBUTES | See attributes (5) for descriptions of the following attributes: <table border="1" data-bbox="402 827 1299 911"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | sh (1), attributes (5) | | | | |
| NOTES | When compiling multithreaded applications, the _REENTRANT flag must be defined on the compile line. This flag should only be used in multithreaded applications. | | | | |

NAME | hypot – Euclidean distance function

SYNOPSIS | cc [*flag ...*] *file ...* -lm [*library ...*]
 #include <math.h>
 double hypot(double x, double y);

DESCRIPTION | The hypot() function computes the length of the hypotenuse of a right-angled triangle:



RETURN VALUES | Upon successful completion, hypot() returns the length of the hypotenuse of a right angled triangle with sides of length *x* and *y*.

If the result would cause overflow, HUGE_VAL is returned and errno may be set to ERANGE.

If *x* or *y* is NaN, NaN is returned.

ERRORS | The hypot() function may fail if:

ERANGE The result overflows.

USAGE | The hypot() function takes precautions against underflow and overflow during intermediate steps of the computation.

An application wishing to check for error situations should set errno to 0 before calling hypot(). If errno is non-zero on return, or the return value is HUGE_VAL or NaN, an error has occurred.

ATTRIBUTES | See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | isnan(3M), sqrt(3M), attributes(5)

NAME | ilogb – returns an unbiased exponent

SYNOPSIS | cc [*flag ...*] *file ...* -lm [*library ...*]
 | #include <math.h>
 | int ilogb(double x);

DESCRIPTION | The ilogb() function returns the exponent part of x . Formally, the return value is the integral part of $\log_r |x|$ as a signed integral value, for non-zero finite x , where r is the radix of the machine's floating point arithmetic.

RETURN VALUES | Upon successful completion, ilogb() returns the exponent part of x .
 | If x is 0, ilogb() returns `-INT_MAX`.
 | If x is NaN or $\pm\text{Inf}$, ilogb() returns `INT_MAX`.

ATTRIBUTES | See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | logb(3M), attributes(5)

| NAME | isencrypt – determine whether a buffer of characters is encrypted | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [flag...] [file...] -lgen [library...] | | | | |
| DESCRIPTION | <pre>#include<libgen.h> int isencrypt(const char *fbuf, size_t ninbuf);</pre> <p>isencrypt() uses heuristics to determine whether a buffer of characters is encrypted. It requires two arguments: a pointer to an array of characters and the number of characters in the buffer.</p> <p>isencrypt() assumes that the file is not encrypted if all the characters in the first block are ASCII characters. If there are non-ASCII characters in the first <i>ninbuf</i> characters, and if the setlocale() LC_CTYPE category is set to C or ascii, isencrypt() assumes that the buffer is encrypted</p> <p>If the LC_CTYPE category is set to a value other than C or ascii, then isencrypt() uses a combination of heuristics to determine if the buffer is encrypted. If <i>ninbuf</i> has at least 64 characters, a chi-square test is used to determine if the bytes in the buffer have a uniform distribution; if it does, then isencrypt() assumes the buffer is encrypted. If the buffer has less than 64 characters, a check is made for null characters and a terminating new-line to determine whether the buffer is encrypted.</p> | | | | |
| RETURN VALUES | If the buffer is encrypted, 1 is returned; otherwise, zero is returned. | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | setlocale(3C), attributes(5) | | | | |
| NOTES | When compiling multithreaded applications, the <code>_REENTRANT</code> flag must be defined on the compile line. This flag should only be used in multithreaded applications. | | | | |

NAME | isnan – test for NaN

SYNOPSIS | cc [*flag ...*] *file ...* -lm [*library ...*]
 | #include <math.h>
 | int **isnan**(double *x*);

DESCRIPTION | The `isnan()` function tests whether *x* is NaN.

RETURN VALUES | The `isnan()` function returns non-zero if *x* is NaN. Otherwise, 0 is returned.

USAGE | On systems not supporting NaN, `isnan()` always returns 0.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `attributes(5)`

| NAME | j0, j1, jn – Bessel functions of the first kind | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double j0(double x); double j1(double x); double jn(int n, double x); | | | | |
| DESCRIPTION | The j0(), j1() and jn() functions compute Bessel functions of x of the first kind of orders 0, 1 and n respectively. | | | | |
| RETURN VALUES | Upon successful completion, j0(), j1() and jn() return the relevant Bessel value of x of the first kind. If the x argument is too large in magnitude, 0 is returned and errno may be set to ERANGE. If x is NaN, NaN is returned. For exceptional cases, matherr(3M) tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The j0(), j1() and jn() functions may fail if: ERANGE The value of x was too large in magnitude. | | | | |
| USAGE | An application wishing to check for error situations should set errno to 0 before calling j0(), j1() or jn(). If errno is non-zero on return, or the return value is NaN, an error has occurred. | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | isnan(3M), matherr(3M), y0(3M), attributes(5), standards(5) | | | | |

| | |
|--------------------|--|
| NAME | kstat – tied hash interface to the kstat facility |
| SYNOPSIS | <pre>use Sun::Solaris::Kstat; Sun::Solaris::Kstat->new(); Sun::Solaris::Kstat->update(); Sun::Solaris::Kstat->{module}{instance}{name}{statistic}</pre> |
| DESCRIPTION | <p>Kernel statistics are categorized using a 3-part key consisting of the module, the instance, and the statistic name. For example, CPU information can be found under <code>cpu_stat:0:cpu_stat0</code>, as in the above example. The method <code>Sun::Solaris::Kstat->new()</code> creates a new 3-layer tree of Perl hashes with the same structure; that is, the statistic for CPU 0 can be accessed as <code>\$ks->{cpu_stat}{0}{cpu_stat0}</code>. The fourth and lowest layer is a tied hash used to hold the individual statistics values for a particular system resource.</p> <p>For performance reasons, the creation of a <code>Sun::Solaris::Kstat</code> object is not accompanied by a following read of all possible statistics. Instead, the 3-layer structure described above is created, but reads of a statistic's values are done only when referenced. For example, accessing <code>\$ks->{cpu_stat}{0}{cpu_stat0}{syscall}</code> will read in all the statistics for CPU 0, including user, system, and wait times, and the other CPU statistics, as well as the number of system call entries. Once you have accessed a lowest level statistics value, calling <code>\$ks->update</code> will automatically update all the individual values of any statistics you have accessed.</p> <p>Note that there are two values of the lowest-level hash that can be read without causing the full set of statistics to be read from the kernel. These are "class", which is the <code>kstat</code> class of the statistics, and "ctime", which is the time that the <code>kstat</code> was created. See <code>kstat(3KSTAT)</code> for full details of these fields.</p> |
| Methods | <pre>new() Create a new kstat statistics hierarchy and return a reference to the top-level hash. Use it like any normal hash to access the statistics. update() Update all the statistics that have been accessed so far. In scalar context, <code>update()</code> returns 1 if the <code>kstat</code> structure has changed, and 0 otherwise. In list context, <code>update()</code> returns references to two arrays: the first holds the keys of any kstats that have been added, and the second holds the keys of any kstats that have been deleted. Each key will be returned in the form "module:instance:name".</pre> |
| EXAMPLES | <p>EXAMPLE 1 Sun::Solaris::Kstat example</p> <pre>use Sun::Solaris::Kstat;</pre> |

```

my $kstat = Sun::Solaris::Kstat->new();
my ($usr1, $sys1, $wio1, $idle1) =
    @{$kstat->{cpu_stat}{0}{cpu_stat0}}{qw(user kernel wait idle)};
print("usr sys wio idle\n");
while (1) {
    sleep 5;
    if ($kstat->update()) {
        print("Configuration changed\n");
    }
    my ($usr2, $sys2, $wio2, $idle2) =
        @{$kstat->{cpu_stat}{0}{cpu_stat0}}{qw(user kernel wait idle)};
    printf(" %.2d  %.2d  %.2d  %.2d\n",
        ($usr2 - $usr1) / 5, ($sys2 - $sys1) / 5,
        ($wio2 - $wio1) / 5, ($idle2 - $idle1) / 5);
    $usr1 = $usr2;
    $sys1 = $sys2;
    $wio1 = $wio2;
    $idle1 = $idle2;
}

```

SEE ALSO

perl(1), kstat(1M), kstat(3KSTAT), kstat_chain_update(3KSTAT),
kstat_close(3KSTAT), kstat_open(3KSTAT), kstat_read(3KSTAT)

NOTES

As the statistics are stored in a tied hash, taking additional references of members of the hash, such as

```

my $ref = \${ks->{cpu_stat}{0}{cpu_stat0}}{syscall};
print("$ref\n");

```

will be recorded as a hold on that statistic's value, preventing it from being updated by `refresh()`. Copy the values explicitly if persistence is necessary.

Several of the statistics provided by the `kstat` facility are stored as 64-bit integer values. Perl 5 does not yet internally support 64-bit integers, so these values are approximated in this module. There are two classes of 64-bit value to be dealt with:

| | |
|----------------------------|--|
| 64-bit intervals and times | These are the <code>crttime</code> and <code>snaptime</code> fields of all the statistics hashes, and the <code>wtime</code> , <code>wlentime</code> , <code>wlastupdate</code> , <code>rtime</code> , <code>rlentime</code> and <code>rlastupdate</code> fields of the <code>kstat</code> I/O statistics structures. These are measured by the <code>kstat</code> facility in nanoseconds, meaning that a 32-bit value would represent approximately 4 seconds. The alternative is to store the values as floating-point numbers, which offer approximately 53 bits of precision on present hardware. 64-bit intervals and timers as floating point values expressed in seconds, meaning that time-related <code>kstats</code> are being rounded to approximately microsecond resolution. |
|----------------------------|--|

64-bit counters

It is not useful to store these values as 32-bit values. As noted above, floating-point values offer 53 bits of precision. Accordingly, all 64-bit counters are stored as floating-point values.

| | |
|-------------------------|--|
| NAME | kstat – kernel statistics facility |
| DESCRIPTION | The kstat facility is a general-purpose mechanism for providing kernel statistics to users. |
| The kstat model | The kernel maintains a linked list of statistics structures, or kstats. Each kstat has a common header section and a type-specific data section. The header section is defined by the <code>kstat_t</code> structure: |
| kstat header | <pre> typedef int kid_t; /* unique kstat id */ typedef struct kstat { /* * Fields relevant to both kernel and user */ hrttime_t ks_crttime; /* creation time */ struct kstat *ks_next; /* kstat chain linkage */ kid_t ks_kid; /* unique kstat ID */ char ks_module[KSTAT_STRLEN]; /* module name */ uchar_t ks_resv; /* reserved */ int ks_instance; /* module's instance */ char ks_name[KSTAT_STRLEN]; /* kstat name */ uchar_t ks_type; /* kstat data type */ char ks_class[KSTAT_STRLEN]; /* kstat class */ uchar_t ks_flags; /* kstat flags */ void *ks_data; /* kstat type-specific data */ uint_t ks_ndata; /* # of data records */ size_t ks_data_size; /* size of kstat data section */ hrttime_t ks_snaptime; /* time of last data snapshot */ /* * Fields relevant to kernel only */ int (*ks_update)(struct kstat *, int); void *ks_private; int (*ks_snapshot)(struct kstat *, void *, int); void *ks_lock; } kstat_t; </pre> |
| | The fields that are of significance to the user are: |
| <code>ks_crttime</code> | The time the kstat was created. This allows you to compute the rates of various counters since the kstat was created; "rate since boot" is replaced by the more general concept of "rate since kstat creation". All times associated with kstats (such as creation time, last snapshot time, <code>kstat_timer_t</code> and <code>kstat_io_t</code> timestamps, and the like) are 64-bit nanosecond values. The accuracy of kstat timestamps is machine dependent, but the precision (units) is the same across all platforms. See <code>gethrtime(3C)</code> for general information about high-resolution timestamps. |

| | |
|--|---|
| <code>ks_next</code> | kstats are stored as a linked list, or chain. <code>ks_next</code> points to the next kstat in the chain. |
| <code>ks_kid</code> | A unique identifier for the kstat. |
| <code>ks_module,</code> <code>ks_instance</code> | contain the name and instance of the the module that created the kstat. In cases where there can only be one instance, <code>ks_instance</code> is 0. |
| <code>ks_name</code> | gives a meaningful name to a kstat. The full kstat namespace is <code><ks_module,ks_instance,ks_name></code> , so the name only need be unique within a module. |
| <code>ks_type</code> | The type of data in this kstat. kstat data types are discussed below. |
| <code>ks_class</code> | Each kstat can be characterized as belonging to some broad class of statistics, such as disk, tape, net, vm, and streams. This field can be used as a filter to extract related kstats. The following values are currently in use: <code>disk</code> , <code>tape</code> , <code>controller</code> , <code>net</code> , <code>rpc</code> , <code>vm</code> , <code>kvm</code> , <code>hat</code> , <code>streams</code> , <code>kmem</code> , <code>kmem_cache</code> , <code>kstat</code> , and <code>misc</code> . (The kstat class encompasses things like <i>kstat_types</i> .) |
| <code>ks_data,</code> <code>ks_ndata,</code> <code>ks_data_size</code> | <code>ks_data</code> is a pointer to the kstat's data section. The type of data stored there depends on <code>ks_type</code> . <code>ks_ndata</code> indicates the number of data records. Only some kstat types support multiple data records. Currently, <code>KSTAT_TYPE_RAW</code> , <code>KSTAT_TYPE_NAMED</code> and <code>KSTAT_TYPE_TIMER</code> kstats support multiple data records. <code>KSTAT_TYPE_INTR</code> and <code>KSTAT_TYPE_IO</code> kstats support only one data record. <code>ks_data_size</code> is the total size of the data section, in bytes. |
| <code>ks_snaptime</code> | The timestamp for the last data snapshot. This allows you to compute activity rates: $\text{rate} = (\text{new_count} - \text{old_count}) / (\text{new_snaptime} - \text{old_snaptime});$ |

kstat data types

The following types of kstats are currently available:

```
#define KSTAT_TYPE_RAW    0    /* can be anything */
```

```
#define KSTAT_TYPE_NAMED 1 /* name/value pairs */
#define KSTAT_TYPE_INTR 2 /* interrupt statistics */
#define KSTAT_TYPE_IO 3 /* I/O statistics */
#define KSTAT_TYPE_TIMER 4 /* event timers */
```

To get a list of all kstat types currently supported in the system, tools can read out the standard system kstat *kstat_types* (full name spec is <"unix", 0, "kstat_types">). This is a KSTAT_TYPE_NAMED kstat in which the *name* field describes the type of kstat, and the *value* field is the kstat type number (for example, KSTAT_TYPE_IO is type 3 – see above).

Raw kstat

KSTAT_TYPE_RAW raw data

The "raw" kstat type is just treated as an array of bytes. This is generally used to export well-known structures, like *sysinfo*.

Name=value kstat

KSTAT_TYPE_NAMED A list of arbitrary *name=value* statistics.

```
typedef struct kstat_named {
    char name[KSTAT_STRLLEN]; /* name of counter */
    uchar_t data_type; /* data type */
    union {
        char c[16]; /* enough for 128-bit ints */
        int32_t i32;
        uint32_t ui32;
        int64_t i64;
        uint64_t ui64;

        /* These structure members are obsolete */

        int32_t l;
        uint32_t ul;
        int64_t ll;
        uint64_t ull;
    } value; /* value of counter */
} kstat_named_t;
#define KSTAT_DATA_CHAR 0
#define KSTAT_DATA_INT32 1
#define KSTAT_DATA_UINT32 2
#define KSTAT_DATA_INT64 3
#define KSTAT_DATA_UINT64 4

/* These types are obsolete */

#define KSTAT_DATA_LONG 1
#define KSTAT_DATA_ULONG 2
#define KSTAT_DATA_LONGLONG 3
#define KSTAT_DATA_ULONGLONG 4
#define KSTAT_DATA_FLOAT 5
#define KSTAT_DATA_DOUBLE 6
```

Interrupt kstat

KSTAT_TYPE_INTR Interrupt statistics.

An interrupt is a hard interrupt (sourced from the hardware device itself), a soft interrupt (induced by the system via the use of some system interrupt source), a watchdog interrupt (induced by a periodic timer call), spurious (an interrupt entry point was entered but there was no interrupt to service), or multiple service (an interrupt was detected and serviced just prior to returning from any of the other types).

```
#define KSTAT_INTR_HARD      0
#define KSTAT_INTR_SOFT     1
#define KSTAT_INTR_WATCHDOG  2
#define KSTAT_INTR_SPURIOUS  3
#define KSTAT_INTR_MULTSVC   4
#define KSTAT_NUM_INTRS     5

typedef struct kstat_intr {
    uint_t intrs[KSTAT_NUM_INTRS]; /* interrupt counters */
} kstat_intr_t;
```

Event timer kstat

KSTAT_TYPE_TIMER Event timer statistics.

These provide basic counting and timing information for any type of event.

```
typedef struct kstat_timer {
    char name[KSTAT_STRLEN]; /* event name */
    uchar_t resv; /* reserved */
    u_longlong_t num_events; /* number of events */
    hrtime_t elapsed_time; /* cumulative elapsed time */
    hrtime_t min_time; /* shortest event duration */
    hrtime_t max_time; /* longest event duration */
    hrtime_t start_time; /* previous event start time */
    hrtime_t stop_time; /* previous event stop time */
} kstat_timer_t;
```

I/O kstat

KSTAT_TYPE_IO I/O statistics.

```
typedef struct kstat_io {
    /*
     * Basic counters.
     */
    u_longlong_t nread; /* number of bytes read */
    u_longlong_t nwritten; /* number of bytes written */
    uint_t reads; /* number of read operations */
    uint_t writes; /* number of write operations */
    /*
     * Accumulated time and queue length statistics.
     */
}
```

(continued)

(Continuation)

```

* Time statistics are kept as a running sum of "active" time.
* Queue length statistics are kept as a running sum of the
* product of queue length and elapsed time at that length -
* that is, a Riemann sum for queue length integrated against time.

```

```

*   ^
*   |
*   8 |-----| i4 |
*   | | | |
* Queue 6 | | | |
* Length |-----| | |
* 4 | i2 |-----| | |
* | | i3 | |
* 2-----| |
* | i1 | |
* |-----|
* Time-> t1 t2 t3 t4

```

```

* At each change of state (entry or exit from the queue),
* we add the elapsed time (since the previous state change)
* to the active time if the queue length was non-zero during
* that interval; and we add the product of the elapsed time
* times the queue length to the running length*time sum.

```

```

* This method is generalizable to measuring residency
* in any defined system: instead of queue lengths, think
* of "outstanding RPC calls to server X".

```

```

* A large number of I/O subsystems have at least two basic
* "lists" of transactions they manage: one for transactions
* that have been accepted for processing but for which processing
* has yet to begin, and one for transactions which are actively
* being processed (but not done). For this reason, two cumulative
* time statistics are defined here: pre-service (wait) time,
* and service (run) time.

```

```

* The units of cumulative busy time are accumulated nanoseconds.
* The units of cumulative length*time products are elapsed time
* times queue length.

```

```

*/
hrtime_t wtime;      /* cumulative wait (pre-service) time */
hrtime_t wlentime;  /* cumulative wait length*time product*/
hrtime_t wlastupdate; /* last time wait queue changed */
hrtime_t rtime;     /* cumulative run (service) time */
hrtime_t rlentime;  /* cumulative run length*time product */
hrtime_t rlastupdate; /* last time run queue changed */
uint_t wcnt;        /* count of elements in wait state */
uint_t rcnt;        /* count of elements in run state */
} kstat_io_t;

```

Using libkstat

The kstat library, `libkstat`, defines the user interface (API) to the system's kstat facility.

You begin by opening libkstat with `kstat_open(3KSTAT)`, which returns a pointer to a fully initialized kstat control structure. This is your ticket to subsequent libkstat operations:

```
typedef struct kstat_ctl {
    kid_t    kc_chain_id;    /* current kstat chain ID */
    kstat_t  *kc_chain;     /* pointer to kstat chain */
    int      kc_kd;        /* /dev/kstat descriptor */
} kstat_ctl_t;
```

Only the first two fields, `kc_chain_id` and `kc_chain`, are of interest to libkstat clients. (`kc_kd` is the descriptor for `/dev/kstat`, the kernel statistics driver. libkstat functions are built on top of `/dev/kstat ioctl(2)` primitives. Direct interaction with `/dev/kstat` is strongly discouraged, since it is *not* a public interface.)

`kc_chain` points to your copy of the kstat chain. You typically walk the chain to find and process a certain kind of kstat. For example, to display all I/O kstats:

```
kstat_ctl_t    *kc;
kstat_t        *ksp;
kstat_io_t     kio;

kc = kstat_open();
for (ksp = kc->kc_chain; ksp != NULL; ksp = ksp->ks_next) {
    if (ksp->ks_type == KSTAT_TYPE_IO) {
        kstat_read(kc, ksp, &kio);
        my_io_display(kio);
    }
}
```

`kc_chain_id` is the kstat chain ID, or KCID, of your copy of the kstat chain. See `kstat_chain_update(3KSTAT)` for an explanation of KCIDs.

FILES

```
/dev/kstat                kernel statistics driver
/usr/include/kstat.h
/usr/include/sys/kstat.h
```

SEE ALSO

```
ioctl(2), gethrtime(3C), getloadavg(3C),
kstat_chain_update(3KSTAT), kstat_close(3KSTAT),
kstat_data_lookup(3KSTAT), kstat_lookup(3KSTAT),
kstat_open(3KSTAT), kstat_read(3KSTAT), kstat_write(3KSTAT)
```

| | |
|----------------------|--|
| NAME | kstat_chain_update – update the kstat header chain |
| SYNOPSIS | <pre>cc [flag ...] file ... -lkstat [library ...] #include <kstat.h></pre> |
| DESCRIPTION | <pre>kid_t kstat_chain_update(kstat_ctl_t *kc);</pre> <p>The <code>kstat_chain_update()</code> function brings the user's kstat header chain in sync with that of the kernel. The kstat chain is a linked list of kstat headers (<code>kstat_t</code>'s) pointed to by <code>kc->kc_chain</code>, which is initialized by <code>kstat_open(3KSTAT)</code>. This chain constitutes a list of all kstats currently in the system.</p> <p>During normal operation, the kernel creates new kstats and delete old ones as various device instances are added and removed, thereby causing the user's copy of the kstat chain to become out of date. The <code>kstat_chain_update()</code> function detects this condition by comparing the kernel's current kstat chain ID(KCID), which is incremented every time the kstat chain changes, to the user's KCID, <code>kc->kc_chain_id</code>. If the KCIDs match, <code>kstat_chain_update()</code> does nothing. Otherwise, it deletes any invalid kstat headers from the user's kstat chain, adds any new ones, and sets <code>kc->kc_chain_id</code> to the new KCID. All other kstat headers in the user's kstat chain are unmodified.</p> |
| RETURN VALUES | The <code>kstat_chain_update()</code> function returns the new KCID if the kstat chain has changed, 0 if it hasn't, or -1 on failure. |
| FILES | <code>/dev/kstat</code> kernel statistics driver |
| SEE ALSO | <code>kstat(3KSTAT)</code> , <code>kstat_close(3KSTAT)</code> , <code>kstat_data_lookup(3KSTAT)</code> , <code>kstat_lookup(3KSTAT)</code> , <code>kstat_open(3KSTAT)</code> , <code>kstat_read(3KSTAT)</code> , <code>kstat_write(3KSTAT)</code> |

| | |
|----------------------|---|
| NAME | kstat_lookup, kstat_data_lookup – find a kstat by name |
| SYNOPSIS | <pre>cc [flag ...] file ... -lkstat [library ...] #include <kstat.h> kstat_t *kstat_lookup(kstat_ctl_t *kc, char *ks_module, int ks_instance, char *ks_name); void *kstat_data_lookup(kstat_t *ksp, char *name);</pre> |
| DESCRIPTION | <p>The <code>kstat_lookup()</code> function traverses the kstat chain, <code>kc->kc_chain</code>, searching for a kstat with the same <code>ks_module</code>, <code>ks_instance</code>, and <code>ks_name</code> fields; this triplet uniquely identifies a kstat. If <code>ks_module</code> is <code>NULL</code>, <code>ks_instance</code> is <code>-1</code>, or <code>ks_name</code> is <code>NULL</code>, then those fields will be ignored in the search. For example, <code>kstat_lookup(kc, NULL, -1, "foo")</code> will simply find the first kstat with name "foo".</p> <p>The <code>kstat_data_lookup()</code> function searches the kstat's data section for the record with the specified <code>name</code>. This operation is valid only for kstat types which have named data records. Currently, only the <code>KSTAT_TYPE_NAMED</code> and <code>KSTAT_TYPE_TIMER</code> kstats have named data records.</p> |
| RETURN VALUES | <p>The <code>kstat_lookup()</code> function returns a pointer to the requested kstat if it is found, or <code>NULL</code> if it is not.</p> <p>The <code>kstat_data_lookup()</code> function returns a pointer to the requested data record if it is found. If the requested record is not found, or if the kstat type is invalid, <code>kstat_data_lookup()</code> returns <code>NULL</code>.</p> |
| FILES | <code>/dev/kstat</code> kernel statistics driver |
| SEE ALSO | <code>kstat(3KSTAT)</code> , <code>kstat_chain_update(3KSTAT)</code> , <code>kstat_close(3KSTAT)</code> , <code>kstat_open(3KSTAT)</code> , <code>kstat_read(3KSTAT)</code> , <code>kstat_write(3KSTAT)</code> |

| | |
|----------------------|---|
| NAME | kstat_open, kstat_close – initialize kernel statistics facility |
| SYNOPSIS | <pre>cc[<i>flag ...</i>] <i>file ...</i> -lkstat [<i>library ...</i>] #include <kstat.h> kstat_ctl_t *kstat_open(<i>void</i>); int kstat_close(kstat_ctl_t *<i>kc</i>);</pre> |
| DESCRIPTION | <p>kstat_open() initializes a kstat control structure, which provides access to the kernel statistics library. It returns a pointer to this structure, which must be supplied as the <i>kc</i> argument in subsequent libkstat function calls.</p> <p>kstat_close() frees all resources that were associated with <i>kc</i>. This is done automatically on exit(2) and execve() (see exec(2)).</p> |
| RETURN VALUES | <p>kstat_open() returns a pointer to a kstat control structure. On failure, it returns NULL and no resources are allocated.</p> <p>kstat_close() returns 0 on success, -1 on failure.</p> |
| FILES | /dev/kstat kernel statistics driver |
| SEE ALSO | kstat(3KSTAT) , kstat_chain_update(3KSTAT) , kstat_data_lookup(3KSTAT) , kstat_lookup(3KSTAT) , kstat_read(3KSTAT) , kstat_write(3KSTAT) |

| | |
|----------------------|---|
| NAME | kstat_read, kstat_write – read or write kstat data |
| SYNOPSIS | <pre>cc [flag ...] file ... -lkstat [library ...] #include <kstat.h> kid_t kstat_read(kstat_ctl_t *kc, kstat_t *ksp, void *buf); kid_t kstat_write(kstat_ctl_t *kc, kstat_t *ksp, void *buf);</pre> |
| DESCRIPTION | <p>kstat_read() gets data from the kernel for the kstat pointed to by <i>ksp</i>. <i>ksp->ks_data</i> is automatically allocated (or reallocated) to be large enough to hold all of the data. <i>ksp->ks_ndata</i> is set to the number of data fields, <i>ksp->ks_data_size</i> is set to the total size of the data, and <i>ksp->ks_snaptime</i> is set to the high-resolution time at which the data snapshot was taken. If <i>buf</i> is non-NULL, the data is copied from <i>ksp->ks_data</i> into <i>buf</i>.</p> <p>kstat_write() writes data from <i>buf</i>, or from <i>ksp->ks_data</i> if <i>buf</i> is NULL, to the corresponding kstat in the kernel. Only the superuser can use kstat_write().</p> |
| RETURN VALUES | On success, kstat_read() and kstat_write() return the current kstat chain ID (KCID). On failure, they return -1 . |
| FILES | /dev/kstat kernel statistics driver |
| SEE ALSO | kstat(3KSTAT), kstat_chain_update(3KSTAT), kstat_close(3KSTAT), kstat_data_lookup(3KSTAT), kstat_lookup(3KSTAT), kstat_open(3KSTAT) |

| NAME | kva_match – look up a key in a key-value array | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag...</i>] <i>file...</i> -lsecdb [<i>library...</i>] #include <secdb.h> char *kva_match(kva_t *kva, char *, key); | | | | |
| DESCRIPTION | The kva_match() function searches a kva_t structure, which is part of the authattr_t, execattr_t, profattr_t, or userattr_t structures. The function takes two arguments: a pointer to a key value array, and a key. If the key is in the array, the function returns a pointer to the first corresponding value that matches that key. Otherwise, the function returns NULL. | | | | |
| RETURN VALUES | Upon success, the function returns a pointer to the value sought. Otherwise, it returns NULL. | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | getauthattr(3SECDB), getexecattr(3SECDB), getprofattr(3SECDB), getuserattr(3SECDB) | | | | |
| NOTES | The kva_match() function returns a pointer to data that already exists in the key-value array. It does not allocate its own memory for this pointer but obtains it from the key-value array that is passed as its first argument. | | | | |

| | |
|----------------------|---|
| NAME | kvm_getu, kvm_getcmd – get the u-area or invocation arguments for a process |
| SYNOPSIS | <pre>#include <kvm.h> #include <sys/param.h> #include <sys/user.h> #include <sys/proc.h> struct user *kvm_getu(kvm_t *kd, struct proc *proc); int kvm_getcmd(kvm_t *kd, struct proc *proc, struct user *u, char ***arg, char ***env);</pre> |
| DESCRIPTION | |
| kvm_getu() | <p>The <code>kvm_getu()</code> function reads the u-area of the process specified by <code>proc</code> to an area of static storage associated with <code>kd</code> and returns a pointer to it. Subsequent calls to <code>kvm_getu()</code> will overwrite this static area.</p> <p>The <code>kd</code> argument is a pointer to a kernel descriptor returned by <code>kvm_open(3KVM)</code>. The <code>proc</code> argument is a pointer to a copy in the current process' address space of a <code>proc</code> structure, obtained, for instance, by a prior <code>kvm_nextproc(3KVM)</code> call.</p> |
| kvm_getcmd() | <p>The <code>kvm_getcmd()</code> function constructs a list of string pointers that represent the command arguments and environment that were used to initiate the process specified by <code>proc</code>.</p> <p>The <code>kd</code> argument is a pointer to a kernel descriptor returned by <code>kvm_open(3KVM)</code>. The <code>u</code> argument is a pointer to a copy in the current process' address space of a <code>user</code> structure, obtained, for instance, by a prior <code>kvm_getu()</code> call. If <code>arg</code> is not <code>NULL</code>, the command line arguments are formed into a null-terminated array of string pointers. The address of the first such pointer is returned in <code>arg</code>. If <code>env</code> is not <code>NULL</code>, then the environment is formed into a null-terminated array of string pointers. The address of the first of these is returned in <code>env</code>.</p> <p>The pointers returned in <code>arg</code> and <code>env</code> refer to data allocated by <code>malloc(3C)</code> and should be freed by a call to <code>free()</code> when no longer needed. See <code>malloc(3C)</code>. Both the string pointers and the strings themselves are deallocated when freed.</p> <p>Since the environment and command line arguments may have been modified by the user process, there is no guarantee that it will be possible to reconstruct the original command at all. Thus, <code>kvm_getcmd()</code> will make the best attempt possible, returning -1 if the user process data is unrecognizable.</p> |
| RETURN VALUES | <p>On success, <code>kvm_getu()</code> returns a pointer to a copy of the u-area of the process specified by <code>proc</code>. On failure, it returns <code>NULL</code>.</p> <p>The <code>kvm_getcmd()</code> function returns 0 on success and -1 on failure.</p> |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Unsafe |

SEE ALSO

`kvm_nextproc(3KVM)`, `kvm_open(3KVM)`, `kvm_read(3KVM)`, `malloc(3C)`, `libkvm(3LIB)`, `attributes (5)`

NOTES

If `kvm_getcmd()` returns `-1`, the caller still has the option of using the command line fragment that is stored in the `u-area`.

On systems that support both 32-bit and 64-bit processes, the 64-bit implementation of `libkvm` ensures that the `arg` and `env` pointer arrays for `kvm_getcmd()` are translated to the same form as if they were 64-bit processes. Applications that wish to access the raw 32-bit stack directly can use `kvm_uread()`. See `kvm_read(3KVM)`.

| | |
|----------------------|--|
| NAME | kvm_nextproc, kvm_getproc, kvm_setproc – read system process structures |
| SYNOPSIS | <pre>#include <kvm.h> #include <sys/param.h> #include <sys/time.h> #include <sys/proc.h> struct proc *kvm_nextproc(kvm_t *kd); int kvm_setproc(kvm_t *kd); struct proc *kvm_getproc(kvm_t *kd, pid_t pid);</pre> |
| DESCRIPTION | |
| kvm_nextproc() | <p>The <code>kvm_nextproc()</code> function may be used to sequentially read all of the system process structures from the kernel identified by <code>kd</code> (see <code>kvm_open(3KVM)</code>). Each call to <code>kvm_nextproc()</code> returns a pointer to the static memory area that contains a copy of the next valid process table entry. There is no guarantee that the data will remain valid across calls to <code>kvm_nextproc()</code>, <code>kvm_setproc()</code>, or <code>kvm_getproc()</code>. Therefore, if the process structure must be saved, it should be copied to non-volatile storage.</p> <p>For performance reasons, many implementations will cache a set of system process structures. Since the system state is liable to change between calls to <code>kvm_nextproc()</code>, and since the cache may contain obsolete information, there is no guarantee that <i>every</i> process structure returned refers to an active process, nor is it certain that <i>all</i> processes will be reported.</p> |
| kvm_setproc() | <p>The <code>kvm_setproc()</code> function rewinds the process list, enabling <code>kvm_nextproc()</code> to rescan from the beginning of the system process table. This function will always flush the process structure cache, allowing an application to re-scan the process table of a running system.</p> |
| kvm_getproc() | <p>The <code>kvm_getproc()</code> function locates the <code>proc</code> structure of the process specified by <code>pid</code> and returns a pointer to it. This function does not interact with the process table pointer manipulated by <code>kvm_nextproc()</code>; however, the restrictions regarding the validity of the data still apply.</p> |
| RETURN VALUES | <p>On success, <code>kvm_nextproc()</code> returns a pointer to a copy of the next valid process table entry. On failure, it returns <code>NULL</code>.</p> <p>On success, <code>kvm_getproc()</code> returns a pointer to the <code>proc</code> structure of the process specified by <code>pid</code>. On failure, it returns <code>NULL</code>.</p> <p>The <code>kvm_setproc()</code> function returns 0 on success -1 on failure.</p> |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Unsafe |

SEE ALSO

kvm_getu(3KVM) , kvm_open(3KVM) , kvm_read(3KVM) , attributes(5)

| NAME | kvm_nlist – get entries from kernel symbol table | | | | |
|----------------------|--|----------------|-----------------|----------|--------|
| SYNOPSIS | <pre>#include <kvm.h> #include <nlist.h> int kvm_nlist(kvm_t *kd, struct nlist *nl);</pre> | | | | |
| DESCRIPTION | <p><code>kvm_nlist()</code> examines the symbol table from the kernel image identified by <i>kd</i> (see <code>kvm_open(3KVM)</code>) and selectively extracts a list of values and puts them in the array of <code>nlist</code> structures pointed to by <code>nl</code>. The name list pointed to by <code>nl</code> consists of an array of structures containing names, types and values. The <i>n_name</i> field of each such structure is taken to be a pointer to a character string representing a symbol name. The list is terminated by an entry with a NULL pointer (or a pointer to a null string) in the <i>n_name</i> field. For each entry in <code>nl</code>, if the named symbol is present in the kernel symbol table, its value and type are placed in the <i>n_value</i> and <i>n_type</i> fields. If a symbol cannot be located, the corresponding <i>n_type</i> field of <code>nl</code> is set to zero.</p> | | | | |
| RETURN VALUES | <p><code>kvm_nlist()</code> returns the value of <code>nlist(3UCB)</code> or <code>nlist(3ELF)</code>, depending on the library used.</p> | | | | |
| ATTRIBUTES | <p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">Unsafe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | Unsafe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | Unsafe | | | | |
| SEE ALSO | <p><code>nlist(3UCB)</code>, <code>nlist(3ELF)</code>, <code>kvm_open(3KVM)</code>, <code>kvm_read(3KVM)</code>, <code>attributes(5)</code></p> | | | | |

| | |
|--------------------|---|
| NAME | kvm_open, kvm_close – specify a kernel to examine |
| SYNOPSIS | <pre>#include <kvm.h> #include <fcntl.h> kvm_t *kvm_open(char *namelist, char *corefile, char *swapfile, int flag, char *errstr); int kvm_close(kvm_t *kd);</pre> |
| DESCRIPTION | |
| kvm_open() | <p>The <code>kvm_open()</code> function initializes a set of file descriptors to be used in subsequent calls to kernel virtual memory (VM) routines. It returns a pointer to a kernel identifier that must be used as the <code>kd</code> argument in subsequent kernel VM function calls.</p> <p>The <code>namelist</code> argument specifies an unstripped executable file whose symbol table will be used to locate various offsets in <code>corefile</code> . If <code>namelist</code> is <code>NULL</code> , the symbol table of the currently running kernel is used to determine offsets in the core image. In this case, it is up to the implementation to select an appropriate way to resolve symbolic references, for instance, using <code>/dev/ksyms</code> as a default <code>namelist</code> file.</p> <p>The <code>corefile</code> argument specifies a file that contains an image of physical memory, for instance, a kernel crash dump file (see <code>savecore(1M)</code>) or the special device <code>/dev/mem</code> . If <code>corefile</code> is <code>NULL</code> , the currently running kernel is accessed, using <code>/dev/mem</code> and <code>/dev/kmem</code> .</p> <p>The <code>swapfile</code> argument specifies a file that represents the swap device. If both <code>corefile</code> and <code>swapfile</code> are <code>NULL</code> , the swap device of the currently running kernel is accessed. Otherwise, if <code>swapfile</code> is <code>NULL</code> , <code>kvm_open()</code> may succeed but subsequent <code>kvm_getu(3KVM)</code> function calls may fail if the desired information is swapped out.</p> <p>The <code>flag</code> function is used to specify read or write access for <code>corefile</code> and may have one of the following values:</p> <pre>O_RDONLY open for reading O_RDWR open for reading and writing</pre> <p>The <code>errstr</code> argument is used to control error reporting. If it is a null pointer, no error messages will be printed. If it is non-null, it is assumed to be the address of a string that will be used to prefix error messages generated by <code>kvm_open</code> . Errors are printed to <code>stderr</code> . A useful value to supply for <code>errstr</code> would be <code>argv [0]</code> . This has the effect of printing the process name in front of any error messages.</p> <p>Applications using <code>libkvm</code> are dependent on the underlying data model of the kernel image, that is, whether it is a 32-bit or 64-bit kernel.</p> |

The data model of these applications must match the data model of the kernel in order to correctly interpret the size and offsets of kernel data structures. For example, a 32-bit application that uses the 32-bit version of the `libkvm` interfaces will fail to open a 64-bit kernel image. Similarly, a 64-bit application that uses the 64-bit version of the `libkvm` interfaces will fail to open a 32-bit kernel image.

kvm_close()

The `kvm_close()` function closes all file descriptors that were associated with `kd`. These files are also closed on `exit(2)` and `execve()` (see `exec(2)`). `kvm_close()` also resets the `proc` pointer associated with `kvm_nextproc(3KVM)` and flushes any cached kernel data.

RETURN VALUES

The `kvm_open()` function returns a non-null value suitable for use with subsequent kernel VM function calls. On failure, it returns `NULL` and no files are opened.

The `kvm_close()` function returns 0 on success -1 on failure.

FILES

`/dev/kmem`
`/dev/ksyms`
`/dev/mem`

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Unsafe |

SEE ALSO

`savecore(1M)`, `exec(2)`, `exit(2)`, `pathconf(2)`, `getloadavg(3C)`, `kstat(3KSTAT)`, `kvm_getu(3KVM)`, `kvm_nextproc(3KVM)`, `kvm_nlist(3KVM)`, `kvm_read(3KVM)`, `sysconf(3C)`, `libkvm(3LIB)`, `proc(4)`, `attributes(5)`

NOTES

Kernel core dumps should be examined on the platform on which they were created. While a 32-bit application running on a 64-bit kernel can examine a 32-bit core dump, a 64-bit application running on a 64-bit kernel cannot examine a kernel core dump from the 32-bit system.

Applications using `libkvm` are likely to be platform- and release-dependent.

On 32-bit systems, applications that use `libkvm` to access the running kernel must be 32-bit applications. On systems that support both 32-bit and 64-bit applications, applications that use the `libkvm` interfaces to access the running kernel must themselves be 64-bit applications.

Most of the traditional uses of `libkvm` have been superseded by more stable interfaces that allow the same information to be extracted more efficiently, yet independent of the kernel data model. For examples, see `sysconf(3C)`, `proc(4)`, `kstat(3KSTAT)`, `getloadavg(3C)`, and `pathconf(2)`.

| | |
|---------------------------|--|
| NAME | kvm_read, kvm_write, kvm_uread, kvm_uwrite, kvm_kread, kvm_kwrite – copy data to or from a kernel image or running system |
| SYNOPSIS | <pre>#include <kvm.h> ssize_t kvm_read(kvm_t *kd, uintptr_t addr, void *buf, size_t nbytes); ssize_t kvm_write(kvm_t *kd, uintptr_t addr, void *buf, size_t nbytes); ssize_t kvm_kread(kvm_t *kd, uintptr_t addr, void *buf, size_t nbytes); ssize_t kvm_kwrite(kvm_t *kd, uintptr_t addr, void *buf, size_t nbytes); ssize_t kvm_uread(kvm_t *kd, uintptr_t addr, void *buf, size_t nbytes); ssize_t kvm_uwrite(kvm_t *kd, uintptr_t addr, void *buf, size_t nbytes);</pre> |
| DESCRIPTION | |
| <code>kvm_kread()</code> | The <code>kvm_kread()</code> function transfers data from the kernel address space to the address space of the process. <i>nbytes</i> bytes of data are copied from the kernel virtual address given by <i>addr</i> to the buffer pointed to by <i>buf</i> . |
| <code>kvm_kwrite()</code> | The <code>kvm_kwrite()</code> function is like <code>kvm_kread()</code> , except that the direction of the transfer is reversed. To use this function, the <code>kvm_open(3KVM)</code> call that returned <i>kd</i> must have specified write access. |
| <code>kvm_uread()</code> | The <code>kvm_uread()</code> function transfers data from the address space of the processes specified in the most recent <code>kvm_getu(3KVM)</code> call. <i>nbytes</i> bytes of data are copied from the user virtual address given by <i>addr</i> to the buffer pointed to by <i>buf</i> . |
| <code>kvm_uwrite()</code> | The <code>kvm_uwrite()</code> function is like <code>kvm_uread()</code> , except that the direction of the transfer is reversed. To use this function, the <code>kvm_open(3KVM)</code> call that returned <i>kd</i> must have specified write access. The address is resolved in the address space of the process specified in the most recent <code>kvm_getu(3KVM)</code> call. |
| <code>kvm_read()</code> | The <code>kvm_read()</code> function transfers data from the kernel image specified by <i>kd</i> (see <code>kvm_open(3KVM)</code>) to the address space of the process. <i>nbytes</i> bytes of data are copied from the kernel virtual address given by <i>addr</i> to the buffer pointed to by <i>buf</i> . |
| <code>kvm_write()</code> | The <code>kvm_write()</code> function is like <code>kvm_read()</code> , except that the direction of data transfer is reversed. To use this function, the <code>kvm_open(3KVM)</code> call that returned <i>kd</i> must have specified write access. If a user virtual address is given, it is resolved in the address space of the process specified in the most recent <code>kvm_getu(3KVM)</code> call. |
| USAGE | The use of <code>kvm_read()</code> and <code>kvm_write()</code> is strongly discouraged. On some platforms, there is considerable ambiguity over which address space is to be accessed by these functions, possibly leading to unexpected results. |

The `kvm_kread()`, `kvm_kwrite()`, `kvm_uread()`, and `kvm_uwrite()` functions are much more clearly defined in this respect.

RETURN VALUES

On success, these functions return the number of bytes actually transferred. On failure, they return -1.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Unsafe |

SEE ALSO

`kvm_getu(3KVM)`, `kvm_nlist(3KVM)`, `kvm_open(3KVM)`, `attributes(5)`

NAME | lgamma, lgamma_r, gamma, gamma_r – log gamma function

SYNOPSIS | cc [*flag ...*] *file ...* -lm [*library ...*]
 #include <math.h>
 extern int signgam;
 double **lgamma**(double *x*);

double **lgamma_r**(double *x*, int **signgamp*);

DESCRIPTION | Both `lgamma()` and `lgamma_r()` return

$$\ln |\Gamma(x)|$$

where

$$\Gamma(x) = \int_0^{\infty} t^{x-1} e^{-t} dt$$

for $x > 0$ and

$$\Gamma(x) = \pi / (\Gamma(1-x) \sin(\pi x))$$

for $x < 1$.

`lgamma()` uses the external integer `signgam` to return the sign of $|\sim(x)$ while `lgamma_r()` uses the user-allocated space addressed by `signgamp`.

IDIOSYNCRASIES | In the case of `lgamma()`, do *not* use the expression `signgam*exp(lgamma(x))` to compute

$$g = \Gamma(x)$$

Instead compute `lgamma()` first:

```
lg = lgamma(x); g = signgam*exp(lg);
```

only after `lgamma()` has returned can `signgam` be correct. Note that $|\sim(x)$ must overflow when x is large enough, underflow when $-x$ is large enough, and generate a division by 0 exception at the singularities x a nonpositive integer.

RETURN VALUES | For exceptional cases, `matherr(3M)` tabulates the values to be returned as dictated by various Standards.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|------------------|
| MT-Level | See NOTES below. |

SEE ALSO | `matherr(3M)`, `attributes(5)`

NOTES

Although `lgamma_r()` is not mentioned by POSIX.4a Draft 6, it was added to complete the functionality provided by similar thread-safe functions. This interface is subject to change to be compatible with the "spirit" of POSIX.4a when it is approved as a standard.

When compiling multi-thread applications, the `_REENTRANT` flag must be defined on the compile line. This flag should only be used in multi-thread applications.

`lgamma()` is unsafe in multithreaded applications. `lgamma_r()` should be used instead.

| | |
|--------------------|---|
| NAME | libdevinfo – library of device information functions |
| SYNOPSIS | <pre>cc [flag ...] file ...-ldevinfo [library ...] #include <libdevinfo.h></pre> |
| DESCRIPTION | <p>libdevinfo is a set of interfaces used to access device configuration data.</p> <p>Device configuration data is organized as a tree of device nodes, defined as <code>di_node_t</code> in the libdevinfo interfaces. Each <code>di_node_t</code> represents a physical or logical (pseudo) device. Three types of data are associated with device nodes:</p> <ul style="list-style-type: none"> ■ data defined for all device nodes (attributes) ■ properties specific to each device ■ minor node data <p>All device nodes have a set of common attributes, such as a node name, an instance number, and a driver binding name. Common device node attributes are accessed by calling interfaces listed on the <code>di_binding_name(3DEVINFO)</code> man page. Each device node also has a physical path, which is accessed by calling <code>di_devfs_path(3DEVINFO)</code>.</p> <p>Properties provide device specific information for device configuration and usage. Properties may be defined by software (<code>di_prop_t</code>) or by firmware (<code>di_prom_prop_t</code>). One way to access each <code>di_prop_t</code> is to make successive calls to <code>di_prop_next(3DEVINFO)</code> until <code>DI_PROP_NIL</code> is returned. For each <code>di_prop_t</code>, use interfaces on the <code>di_prop_bytes(3DEVINFO)</code> man page to obtain property names and values. Another way to access these properties is to call <code>di_prop_lookup_bytes(3DEVINFO)</code> to find the value of a property with a given name. Accessing a <code>di_prom_prop_t</code> is similar to accessing a <code>di_prop_t</code>, except that the interface names start with <code>di_prom_prop</code> and additional calls to <code>di_prom_init(3DEVINFO)</code> and <code>di_prom_fini(3DEVINFO)</code> are required.</p> <p>Minor nodes contain information exported by the device for creating special files for the device. Each device node has 0 or more minor nodes associated with it. A list minor nodes (<code>di_minor_t</code>) may be obtained by making successive calls to <code>di_minor_next(3DEVINFO)</code> until <code>DI_MINOR_NIL</code> is returned. For each minor node, <code>di_minor_devt(3DEVINFO)</code> and related interfaces are called to get minor node data.</p> <p>Using libdevinfo involves three steps:</p> <ul style="list-style-type: none"> ■ Creating a snapshot of the device tree ■ Traversing the device tree to get information of interest |

- Destroying the snapshot of the device tree

A snapshot of the device tree is created by calling `di_init(3DEVINFO)` and destroyed by calling `di_fini(3DEVINFO)`. An application may specify the data to be included in the snapshot (full or partial tree, include or exclude properties and minor nodes) and get a handle to the root of the device tree. See `di_init(3DEVINFO)` for details. The application then traverses the device tree in the snapshot to obtain device configuration data.

The device tree is normally traversed through parent-child-sibling linkage. Each device node contains references to its parent, its next sibling, and the first of its children. Given the `di_node_t` returned from `di_init(3DEVINFO)`, one can find all children by first calling `di_child_node(3DEVINFO)`, followed by successive calls to `di_sibling_node(3DEVINFO)`, until `DI_NODE_NIL` is returned. By following this procedure recursively, an application can visit all device nodes contained in the snapshot. Two interfaces, `di_walk_node(3DEVINFO)` and `di_walk_minor(3DEVINFO)`, are provided to facilitate device tree traversal. The `di_walk_node(3DEVINFO)` interface visits all device nodes and executes a user-supplied callback function for each node visited. The `di_walk_minor(3DEVINFO)` does the same for each minor node in the device tree.

An alternative way to traverse the device tree is through the per-driver device node linkage. Device nodes contain a reference to the next device node bound to the same driver. Given the `di_node_t` returned from `di_init(3DEVINFO)`, an application can find all device nodes bound to a driver by first calling `di_drv_first_node(3DEVINFO)`, followed by successive calls to `di_drv_next_node(3DEVINFO)` until `DI_NODE_NIL` is returned. Note that traversing the per-driver device node list works only when the snapshot includes all device nodes.

See `libdevinfo(3LIB)` for a complete list of `libdevinfo` interfaces. See `di_init(3DEVINFO)` for examples of `libdevinfo` usage. See *Writing Device Drivers* for details of Solaris device configuration.

EXAMPLES

EXAMPLE 1 Information Accessible Through `libdevinfo` Interfaces

The following example illustrates the kind of information accessible through `libdevinfo` interfaces for a device node representing a hard disk (`sd2`):

```
Attributes
  node name:  sd
  instance:   2
  physical path: /sbus@1f,0/espdma@e,8400000/esp@e,8800000/sd@2,0

Properties
  target=2
  lun=0
```

```

Minor nodes
(disk partition /dev/dsk/c0t2d0s0)
  name:      a
  dev_t:     0x0080010 (32/16)
  spectype:  IF_BLK (block special)
(disk partition /dev/rdisk/c0t2d0s2)
  name:      c,raw
  dev_t:     0x0080012 (32/18)
  spectype:  IF_CHR (character special)
    
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level | Safe |
| Interface Stability | Evolving |

SEE ALSO

`devlinks(1M)`, `prtconf(1M)`, `di_binding_name(3DEVINFO)`,
`di_child_node(3DEVINFO)`, `di_devfs_path(3DEVINFO)`,
`di_drv_first_node(3DEVINFO)`, `di_drv_next_node(3DEVINFO)`,
`di_fini(3DEVINFO)`, `di_init(3DEVINFO)`, `di_minor_devt(3DEVINFO)`,
`di_minor_next(3DEVINFO)`, `di_prom_fini(3DEVINFO)`,
`di_prom_init(3DEVINFO)` `di_prop_bytes(3DEVINFO)`,
`di_prop_lookup_bytes (3DEVINFO)`, `di_prop_next(3DEVINFO)`,
`di_sibling_node (3DEVINFO)`, `di_walk_minor(3DEVINFO)`,
`di_walk_node(3DEVINFO)`, `libdevinfo(3LIB)`, `attributes(5)`

Writing Device Drivers

| | | | | | |
|-------------------------------------|---|-------------------------------------|-----|---------------------------------|---|
| NAME | libtnfctl – library for TNF probe control in a process or the kernel | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -ltnfctl [library ...] #include <tnf/tnfctl.h></pre> | | | | |
| DESCRIPTION | <p>The <code>libtnfctl</code> library provides an API to control TNF ("Trace Normal Form") probes within a process or the kernel. See <code>tracing(3TNF)</code> for an overview of the Solaris tracing architecture. The client of <code>libtnfctl</code> controls probes in one of four modes:</p> <p>internal mode The target is the controlling process itself; that is, the client controls its own probes.</p> <p>direct mode The target is a separate process; a client can either <code>exec(2)</code> a program or attach to a running process for probe control. The <code>libtnfctl</code> library uses <code>proc(4)</code> on the target process for probe and process control in this mode, and additionally provides basic process control features.</p> <p>indirect mode The target is a separate process, but the controlling process is already using <code>proc(4)</code> to control the target, and hence <code>libtnfctl</code> cannot use those interfaces directly. Use this mode to control probes from within a debugger. In this mode, the client must provide a set of functions that <code>libtnfctl</code> can use to query and update the target process.</p> <p>kernel mode The target is the Solaris kernel.</p> <p>A process is controlled "externally" if it is being controlled in either direct mode or indirect mode. Alternatively, a process is controlled "internally" when it uses internal mode to control its own probes.</p> <p>There can be only one client at a time doing probe control on a given process. Therefore, it is not possible for a process to be controlled internally while it is being controlled externally. It is also not possible to have a process controlled by multiple external processes. Similarly, there can be only one process at a time doing kernel probe control. Note, however, that while a given target may only be controlled by one <code>libtnfctl</code> client, a single client may control an arbitrary number of targets. That is, it is possible for a process to simultaneously control its own probes, probes in other processes, and probes in the kernel.</p> <p>The following tables denotes the modes applicable to all <code>libtnfctl</code> interfaces (INT = internal mode; D = direct mode; IND = indirect mode; K = kernel mode).</p> <p>These interfaces create handles in the specified modes:</p> <table border="0"> <tr> <td><code>tnfctl_internal_open()</code></td> <td>INT</td> </tr> <tr> <td><code>tnfctl_exec_open()</code></td> <td>D</td> </tr> </table> | <code>tnfctl_internal_open()</code> | INT | <code>tnfctl_exec_open()</code> | D |
| <code>tnfctl_internal_open()</code> | INT | | | | |
| <code>tnfctl_exec_open()</code> | D | | | | |

| | | | | |
|-------------------------|--|---|-----|---|
| tnfctl_pid_open() | | D | | |
| tnfctl_indirect_open() | | | IND | |
| tnfctl_kernel_open() | | | | K |

These interfaces are used with the specified modes:

| | | | | |
|------------------------------------|-----|---|-----|---|
| tnfctl_continue() | | D | | |
| tnfctl_probe_connect() | INT | D | IND | |
| tnfctl_probe_disconnect_all () | INT | D | IND | |
| tnfctl_trace_attrs_get() | INT | D | IND | K |
| tnfctl_buffer_alloc() | INT | D | IND | K |
| tnfctl_register_funcs() | INT | D | IND | K |
| tnfctl_probe_apply() | INT | D | IND | K |
| tnfctl_probe_apply_ids() | INT | D | IND | K |
| tnfctl_probe_state_get () | INT | D | IND | K |
| tnfctl_probe_enable() | INT | D | IND | K |
| tnfctl_probe_disable() | INT | D | IND | K |
| tnfctl_probe_trace() | INT | D | IND | K |
| tnfctl_probe_untrace() | INT | D | IND | K |
| tnfctl_check_libs() | INT | D | IND | K |
| tnfctl_close() | INT | D | IND | K |
| tnfctl_strerror() | INT | D | IND | K |
| tnfctl_buffer_dealloc() | | | | K |
| tnfctl_trace_state_set() | | | | K |
| tnfctl_filter_state_set() | | | | K |
| tnfctl_filter_list_get() | | | | K |
| tnfctl_filter_list_add() | | | | K |
| tnfctl_filter_list_delete() | | | | K |

When using `libtnfctl`, the first task is to create a handle for controlling probes. The `tnfctl_internal_open()` function creates an internal mode handle for controlling probes in the same process, as described above. The `tnfctl_pid_open()` and `tnfctl_exec_open()` functions create handles

in direct mode. The `tnfctl_indirect_open()` function creates an indirect mode handle, and the `tnfctl_kernel_open()` function creates a kernel mode handle. A handle is required for use in nearly all other `libtnfctl` functions. The `tnfctl_close()` function releases the resources associated with a handle.

The `tnfctl_continue()` function is used in direct mode to resume execution of the target process.

The `tnfctl_buffer_alloc()` function allocates a trace file or, in kernel mode, a trace buffer.

The `tnfctl_probe_apply()` and `tnfctl_probe_apply_ids()` functions call a specified function for each probe or for a designated set of probes.

The `tnfctl_register_funcs()` function registers functions to be called whenever new probes are seen or probes have disappeared, providing an opportunity to do one-time processing for each probe.

The `tnfctl_check_libs()` function is used primarily in indirect mode to check whether any new probes have appeared, that is, they have been made available by `dlopen(3DL)`, or have disappeared, that is, they have disassociated from the process by `dlclose(3DL)`.

The `tnfctl_probe_enable()` and `tnfctl_probe_disable()` functions control whether the probe, when hit, will be ignored.

The `tnfctl_probe_trace()` and `tnfctl_probe_untrace()` functions control whether an enabled probe, when hit, will cause an entry to be made in the trace file.

The `tnfctl_probe_connect()` and `tnfctl_probe_disconnect_all()` functions control which functions, if any, are called when an enabled probe is hit.

The `tnfctl_probe_state_get()` function returns information about the status of a probe, such as whether it is currently enabled.

The `tnfctl_trace_attrs_get()` function returns information about the tracing session, such as the size of the trace buffer or trace file.

The `tnfctl_strerror()` function maps a `tnfctl` error code to a string, for reporting purposes.

The remaining functions apply only to kernel mode.

The `tnfctl_trace_state_set()` function controls the master switch for kernel tracing. See `prex(1)` for more details.

The `tnfctl_filter_state_set()`, `tnfctl_filter_list_get()`, `tnfctl_filter_list_add()`, and `tnfctl_filter_list_delete()` functions allow a set of processes to be specified for which probes will not

be ignored when hit. This prevents kernel activity caused by uninteresting processes from cluttering up the kernel's trace buffer.

The `tnfctl_buffer_dealloc()` function deallocates the kernel's internal trace buffer.

RETURN VALUES

Upon successful completion, these functions return `TNFCTL_ERR_NONE`.

ERRORS

The error codes for `libtnfctl` are:

| | |
|---------------------------------------|---|
| <code>TNFCTL_ERR_ACCES</code> | Permission denied. |
| <code>TNFCTL_ERR_NOTARGET</code> | The target process completed. |
| <code>TNFCTL_ERR_ALLOCFAIL</code> | A memory allocation failure occurred. |
| <code>TNFCTL_ERR_INTERNAL</code> | An internal error occurred. |
| <code>TNFCTL_ERR_SIZETOOSMALL</code> | The requested trace size is too small. |
| <code>TNFCTL_ERR_SIZETOOBIG</code> | The requested trace size is too big. |
| <code>TNFCTL_ERR_BADARG</code> | Bad input argument. |
| <code>TNFCTL_ERR_NOTDYNAMIC</code> | The target is not a dynamic executable. |
| <code>TNFCTL_ERR_NOLIBTNFPROBE</code> | <code>libtnfprobe.so</code> not linked in target. |
| <code>TNFCTL_ERR_BUFBROKEN</code> | Tracing is broken in the target. |
| <code>TNFCTL_ERR_BUFEXISTS</code> | A buffer already exists. |
| <code>TNFCTL_ERR_NOBUF</code> | No buffer exists. |
| <code>TNFCTL_ERR_BADDEALLOC</code> | Cannot deallocate buffer. |
| <code>TNFCTL_ERR_NOPROCESS</code> | No such target process exists. |
| <code>TNFCTL_ERR_FILENOTFOUND</code> | File not found. |
| <code>TNFCTL_ERR_BUSY</code> | Cannot attach to process or kernel because it is already tracing. |
| <code>TNFCTL_ERR_INVALIDPROBE</code> | Probe no longer valid. |
| <code>TNFCTL_ERR_USR1</code> | Error code reserved for user. |
| <code>TNFCTL_ERR_USR2</code> | Error code reserved for user. |
| <code>TNFCTL_ERR_USR3</code> | Error code reserved for user. |
| <code>TNFCTL_ERR_USR4</code> | Error code reserved for user. |
| <code>TNFCTL_ERR_USR5</code> | Error code reserved for user. |

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-------------------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe with exceptions |

SEE ALSO

`prex(1)`, `exec(2)`, `dlclose(3DL)`, `dlopen(3DL)`, `TNF_PROBE(3TNF)`, `tnfctl_buffer_alloc(3TNF)`, `tnfctl_buffer_dealloc(3TNF)`, `tnfctl_check_libs(3TNF)`, `tnfctl_close(3TNF)`, `tnfctl_continue(3TNF)`, `tnfctl_internal_open(3TNF)`, `tnfctl_exec_open(3TNF)`, `tnfctl_filter_list_add(3TNF)`, `tnfctl_filter_list_delete(3TNF)`, `tnfctl_filter_list_get(3TNF)`, `tnfctl_filter_state_set(3TNF)`, `tnfctl_kernel_open(3TNF)`, `tnfctl_pid_open(3TNF)`, `tnfctl_probe_apply(3TNF)`, `tnfctl_probe_apply_ids(3TNF)`, `tnfctl_probe_connect(3TNF)`, `tnfctl_probe_disable(3TNF)`, `tnfctl_probe_enable(3TNF)`, `tnfctl_probe_state_get(3TNF)`, `tnfctl_probe_trace(3TNF)`, `tnfctl_probe_untrace(3TNF)`, `tnfctl_indirect_open(3TNF)`, `tnfctl_register_funcs(3TNF)`, `tnfctl_strerror(3TNF)`, `tnfctl_trace_attrs_get(3TNF)`, `tnfctl_trace_state_set(3TNF)`, `libtnfctl(3LIB)`, `proc(4)`, `attributes(5)`

Linker and Libraries Guide

NOTES

This API is MT-Safe. Multiple threads may concurrently operate on independent `tnfctl` handles, which is the typical behavior expected. The `libtnfctl` library does not support multiple threads operating on the same `tnfctl` handle. If this is desired, it is the client's responsibility to implement locking to ensure that two threads that use the same `tnfctl` handle are not simultaneously in a `libtnfctl` interface.

| NAME | log10 – base 10 logarithm function | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double log10(double <i>x</i>); | | | | |
| DESCRIPTION | The log10() function computes the base 10 logarithm of <i>x</i> , $\log_{10}(x)$. The value of <i>x</i> must be positive. | | | | |
| RETURN VALUES | Upon successful completion, log10() returns the base 10 logarithm of <i>x</i> . If <i>x</i> is NaN, NaN is returned. If <i>x</i> is less than 0, -HUGE_VAL or NaN is returned, and errno is set to EDOM. If <i>x</i> is 0, -HUGE_VAL is returned and errno may be set to ERANGE. For exceptional cases, matherr(3M) tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The log10() function will fail if: EDOM The value of <i>x</i> is negative. The log10() function may fail if: ERANGE The value of <i>x</i> is 0. No other errors will occur. | | | | |
| USAGE | An application wishing to check for error situations should set errno to 0 before calling log10(). If errno is non-zero on return, or the return value is NaN, an error has occurred. | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | isnan(3M), log(3M), matherr(3M), pow(3M), attributes(5), standards(5) | | | | |

| NAME | log1p – compute natural logarithm | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double log1p (double <i>x</i>); | | | | |
| DESCRIPTION | The <code>log1p()</code> function computes $\log_e(1.0 + x)$. The value of <i>x</i> must be greater than -1.0 . | | | | |
| RETURN VALUES | Upon successful completion, <code>log1p()</code> returns the natural logarithm of $1.0 + x$. If <i>x</i> is NaN, <code>log1p()</code> returns NaN. If <i>x</i> is less than -1.0 , <code>log1p()</code> returns <code>-HUGE_VAL</code> or NaN and sets <code>errno</code> to <code>EDOM</code> . If <i>x</i> is -1.0 , <code>log1p()</code> returns <code>-HUGE_VAL</code> and may set <code>errno</code> to <code>ERANGE</code> . For exceptional cases, <code>matherr(3M)</code> tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The <code>log1p()</code> function will fail if: <code>EDOM</code> The value of <i>x</i> is less than -1.0 . The <code>log1p()</code> function may fail and set <code>errno</code> to: <code>ERANGE</code> The value of <i>x</i> is -1.0 . | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>log(3M)</code> , <code>matherr(3M)</code> , <code>attributes(5)</code> , <code>standards(5)</code> | | | | |

| NAME | log – natural logarithm function | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double log(double <i>x</i>); | | | | |
| DESCRIPTION | The log() function computes the natural logarithm of <i>x</i> , $\log_e(x)$. The value of <i>x</i> must be positive. | | | | |
| RETURN VALUES | Upon successful completion, log() returns the natural logarithm of <i>x</i> . If <i>x</i> is NaN, NaN is returned. If <i>x</i> is less than 0, -HUGE_VAL or NaN is returned and <code>errno</code> is set to <code>EDOM</code> . If <i>x</i> is 0, -HUGE_VAL is returned and <code>errno</code> may be set to <code>ERANGE</code> . In IEEE 754 mode (the <code>-xlibmieee</code> cc compilation option), if <i>x</i> is Inf or a quiet NaN, <i>x</i> is returned; if <i>x</i> is a signaling NaN, a quiet NaN is returned and the invalid operation exception is raised; if <i>x</i> is 1, 0 is returned; for all other positive <i>x</i> , a normalized number is returned and the inexact exception is raised. For exceptional cases, <code>matherr(3M)</code> tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The log() function will fail if: EDOM The value of <i>x</i> is negative. The log() function may fail if: ERANGE The value of <i>x</i> is 0. No other errors will occur. | | | | |
| USAGE | An application wishing to check for error situations should set <code>errno</code> to 0 before calling log(). If <code>errno</code> is non-zero on return, or the return value is NaN, an error has occurred. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>exp(3M)</code> , <code>isnan(3M)</code> , <code>log10(3M)</code> , <code>log1p(3M)</code> , <code>matherr(3M)</code> , <code>attributes(5)</code> , <code>standards(5)</code> | | | | |

| NAME | logb – radix-independent exponent | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double logb (double <i>x</i>); | | | | |
| DESCRIPTION | The <code>logb()</code> function computes the exponent of <i>x</i> , which is the integral part of $\log_r x $, as a signed floating point value, for non-zero <i>x</i> , where <i>r</i> is the radix of the machine's floating-point arithmetic. | | | | |
| RETURN VALUES | Upon successful completion, <code>logb()</code> returns the exponent of <i>x</i> . If <i>x</i> is 0.0, <code>logb()</code> returns <code>-HUGE_VAL</code> and sets <code>errno</code> to <code>EDOM</code> . If <i>x</i> is $\pm\text{Inf}$, <code>logb()</code> returns <code>+Inf</code> . If <i>x</i> is NaN, <code>logb()</code> returns NaN. For exceptional cases, <code>matherr(3M)</code> tabulates the values to be returned as dictated by various Standards. | | | | |
| ERRORS | The <code>logb()</code> function will fail if: <code>EDOM</code> The <i>x</i> argument is 0.0. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>ilogb(3M)</code> , <code>matherr(3M)</code> , <code>attributes(5)</code> | | | | |

| NAME | maillock, mailunlock, touchlock – functions to manage lockfile(s) for user's mailbox | | | | |
|----------------------|---|----------------|-----------------|----------|--------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lmail [library ...] #include <maillock.h> int maillock(const char *user, int retrycnt); void mailunlock(void); void touchlock(void);</pre> | | | | |
| DESCRIPTION | <p>The maillock() function attempts to create a lockfile for the user's mailfile. If a lockfile already exists, and it has not been modified in the last 5 minutes, maillock() will remove the lockfile and set its own lockfile.</p> <p>It is crucial that programs locking mail files refresh their locks at least every three minutes to maintain the lock. Refresh the lockfile by calling the touchlock() function with no arguments.</p> <p>The algorithm used to determine the age of the lockfile takes into account clock drift between machines using a network file system. A zero is written into the lockfile so that the lock will be respected by systems running the standard version of System V.</p> <p>If the lockfile has been modified in the last 5 minutes the process will sleep until the lock is available. The sleep algorithm is to sleep for 5 seconds times the attempt number. That is, the first sleep will be for 5 seconds, the next sleep will be for 10 seconds, etc. until the number of attempts reaches <i>retrycnt</i>.</p> <p>When the lockfile is no longer needed, it should be removed by calling mailunlock().</p> <p>The <i>user</i> argument is the login name of the user for whose mailbox the lockfile will be created. maillock() assumes that user's mailfiles are in the "standard" place as defined in <maillock.h>.</p> | | | | |
| RETURN VALUES | Upon successful completion, maillock() returns 0. Otherwise it returns -1. | | | | |
| FILES | <pre>/var/mail/* user mailbox files /var/mail/*.lock user mailbox lockfiles</pre> | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: <table border="1" style="margin-left: 20px; border-collapse: collapse; width: 50%;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">Unsafe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | Unsafe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | Unsafe | | | | |
| SEE ALSO | libmail(3LIB), attributes(5) | | | | |

NOTES

The `mailunlock()` function will only remove the lockfile created from the most previous call to `maillock()`. Calling `maillock()` for different users without intervening calls to `mailunlock()` will cause the initially created lockfile(s) to remain, potentially blocking subsequent message delivery until the current process finally terminates.

| | | | | | | | | | | | | | |
|--------------------|---|--------|---------------------------|------|----------------------|----------|--------------------------|-----------|---------------------------|-------|----------------------------|-------|------------------------------|
| NAME | matherr – math library exception-handling function | | | | | | | | | | | | |
| SYNOPSIS | <pre>#include <math.h> int matherr(struct exception *exc);</pre> | | | | | | | | | | | | |
| DESCRIPTION | <p>The The System V Interface Definition, Third Edition (SVID3) specifies that certain libm functions call <code>matherr()</code> when exceptions are detected. Users may define their own mechanisms for handling exceptions, by including a function named <code>matherr()</code> in their programs. The <code>matherr()</code> function is of the form described above. When an exception occurs, a pointer to the exception structure <code>exc</code> will be passed to the user-supplied <code>matherr()</code> function. This structure, which is defined in the <code><math.h></code> header file, is as follows:</p> <pre>struct exception { int type; char *name; double arg1, arg2, retval; };</pre> <p>The <code>type</code> member is an integer describing the type of exception that has occurred, from the following list of constants (defined in the header file):</p> <table border="0"> <tr> <td>DOMAIN</td> <td>argument domain exception</td> </tr> <tr> <td>SING</td> <td>argument singularity</td> </tr> <tr> <td>OVERFLOW</td> <td>overflow range exception</td> </tr> <tr> <td>UNDERFLOW</td> <td>underflow range exception</td> </tr> <tr> <td>TLOSS</td> <td>total loss of significance</td> </tr> <tr> <td>PLOSS</td> <td>partial loss of significance</td> </tr> </table> <p>Note that both <code>TLOSS</code> and <code>PLOSS</code> reflect limitations of particular algorithms for trigonometric functions that suffer abrupt declines in accuracy at definite boundaries. Since the implementation does not suffer such abrupt declines, <code>PLOSS</code> is never signaled. <code>TLOSS</code> is signaled for Bessel functions <i>only</i> to satisfy SVID3 requirements.</p> <p>The <code>name</code> member points to a string containing the name of the function that incurred the exception. The <code>arg1</code> and <code>arg2</code> members are the arguments with which the function was invoked. <code>retval</code> is set to the default value that will be returned by the function unless the user's <code>matherr()</code> sets it to a different value.</p> <p>If the user's <code>matherr()</code> function returns non-zero, no exception message will be printed, and <code>errno</code> will not be set.</p> | DOMAIN | argument domain exception | SING | argument singularity | OVERFLOW | overflow range exception | UNDERFLOW | underflow range exception | TLOSS | total loss of significance | PLOSS | partial loss of significance |
| DOMAIN | argument domain exception | | | | | | | | | | | | |
| SING | argument singularity | | | | | | | | | | | | |
| OVERFLOW | overflow range exception | | | | | | | | | | | | |
| UNDERFLOW | underflow range exception | | | | | | | | | | | | |
| TLOSS | total loss of significance | | | | | | | | | | | | |
| PLOSS | partial loss of significance | | | | | | | | | | | | |

**SVID3
STANDARD
CONFORMANCE**

When an application is built as a SVID3 conforming application (see standards(5)), if `matherr()` is not supplied by the user, the default `matherr` exception-handling mechanisms, summarized in the table below, will be invoked upon exception:

- DOMAIN 0.0 is usually returned, `errno` is set to `EDOM`, and a message is usually printed on standard error.
- SING The largest finite single-precision number, `HUGE` of appropriate sign is returned, `errno` is set to `EDOM`, and a message is printed on standard error.
- OVERFLOW The largest finite single-precision number, `HUGE` of appropriate sign is usually returned, `errno` is set to `ERANGE`.
- UNDERFLOW 0.0 is returned, and `errno` is set to `ERANGE`.
- TLOSS 0.0 is returned, `errno` is set to `ERANGE`, and a message is printed on standard error.

In general, `errno` is not a reliable error indicator in that it may be unexpectedly set by a function in a handler for an asynchronous signal.

**SVID3 ERROR
HANDLING
PROCEDURES
(compile with `cc \-Xt`)**

| <math.h> type | DOMAIN | SING | OVERFLOW | UNDERFLOW | TLOSS |
|---|-------------------------|--------------------------|--------------------------|---------------------------|---------|
| <code>errno</code> | EDOM | EDOM | ERANGE | ERANGE | ERANGE |
| IEEE Exception | Invalid Operation | Division by Zero | Overflow | Underflow | - |
| <code>fp_exception_type</code> | <code>fp_invalid</code> | <code>fp_division</code> | <code>fp_overflow</code> | <code>fp_underflow</code> | - |
| ACOS, ASIN ($ x > 1$): | Md, 0.0 | - | - | - | - |
| ACOSH ($x < 1$), ATANH ($ x > 1$): | NaN | - | - | - | - |
| ATAN2 (0,0): | Md, 0.0 | - | - | - | - |
| COSH, SINH: | - | - | +HUGE | - | - |
| EXP: | - | - | +HUGE | 0.0 | - |
| FMOD (x,0): | x | - | - | - | - |
| HYPOT: | - | - | +HUGE | - | - |
| J0, J1, JN ($ x > X_TLOSS$): | - | - | - | - | Mt, 0.0 |
| LGAMMA: usual cases | - | - | +HUGE | - | - |

| <math.h> type | DOMAIN | SING | OVERFLOW | UNDERFLOW | TLOSS |
|-------------------------------|--------------|--------------|-----------|-----------|---------|
| (x = 0 or -integer) | - | Ms, +HUGE | - | - | - |
| LOG, LOG10: | | | | | |
| (x < 0) | Md, -HUGE | - | - | - | - |
| (x = 0) | - | Ms, -HUGE | - | - | - |
| POW: | | | | | |
| usual cases | - | - | ±HUGE | ±0.0 | - |
| (x < 0) ** (y not an integer) | Md, 0.0 | - | - | - | - |
| 0 ** 0 | Md, 0.0 | - | - | - | - |
| 0 ** (y < 0) | Md, 0.0 | - | - | - | - |
| REMAINDER (x,0): | NaN | - | - | - | - |
| SCALB: | - | - | ±HUGE_VAL | ±0.0 | - |
| SQRT (x < 0): | Md, 0.0 | - | - | - | - |
| Y0, Y1, YN: | | | | | |
| (x < 0) | Md, -HUGE | - | - | - | - |
| (x = 0) | - | Md, -HUGE | - | - | - |
| (x > X_TLOSS) | - | - | - | - | Mt, 0.0 |

Abbreviations

| | |
|----------|--|
| Md | Message is printed (DOMAIN error). |
| Ms | Message is printed (SING error). |
| Mt | Message is printed (TLOSS error). |
| NaN | IEEE NaN result and invalid operation exception. |
| HUGE | Maximum finite single-precision floating-point number. |
| HUGE_VAL | IEEE ∞ result and division-by-zero exception. |
| X_TLOSS | The value X_TLOSS is defined in <values.h>. |

**X/OPEN
COMMON
APPLICATION
ENVIRONMENT
(CAE)
SPECIFICATIONS
CONFORMANCE
CAE
SPECIFICATION
ERROR HANDLING
PROCEDURES
(compile with cc -Xa)**

The interaction of IEEE arithmetic and `matherr()` is not defined when executing under IEEE rounding modes other than the default round to nearest: `matherr()` is not always called on overflow or underflow, and the `matherr()` may return results that differ from those in this table.

The X/Open System Interfaces and Headers (XSH) Issue 3 and later revisions of that specification no longer sanctions the use of the `matherr()` interface. The following table summarizes the values returned in the exceptional cases. In general, XSH dictates that as long as one of the input argument(s) is a NaN, NaN shall be returned. In particular, `pow(NaN, 0) = NaN`.

| <math.h> type | DOMAIN | SING | OVERFLOW | UNDERFLOW | TLOSS |
|----------------------------------|-----------|-----------|-------------|-----------|--------|
| <code>errno</code> | EDOM | EDOM | ERANGE | ERANGE | ERANGE |
| ACOS, ASIN ($ x > 1$): | 0.0 | - | - | - | - |
| ATAN2(0,0): | 0.0 | - | - | - | - |
| COSH, SINH: | - | - | {±HUGE_VAL} | - | - |
| EXP: | - | - | {+HUGE_VAL} | {0.0} | - |
| FMOD(x,0): | {NaN} | - | - | - | - |
| HYPOT: | - | - | {+HUGE_VAL} | - | - |
| J0, J1, JN ($ x > X_TLOSS$): | - | - | - | - | {0.0} |
| LGAMMA: | | | | | |
| usual cases | - | - | {+HUGE_VAL} | - | - |
| (x = 0 or -integer) | - | +HUGE_VAL | - | - | - |
| LOG, LOG10: | | | | | |
| (x < 0) | -HUGE_VAL | - | - | - | - |
| (x = 0) | - | -HUGE_VAL | - | - | - |
| POW: | | | | | |

| <math.h> type | DOMAIN | SING | OVERFLOW | UNDERFLOW | TLOSS |
|-------------------------------|-------------|-------------|-----------|-----------|-------|
| usual cases | - | - | ±HUGE_VAL | ±0.0 | - |
| (x < 0) ** (y not an integer) | 0.0 | - | - | - | - |
| 0 ** 0 | {1.0} | - | - | - | - |
| 0 ** (y < 0) | {-HUGE_VAL} | - | - | - | - |
| SQRT (x < 0): | 0.0 | - | - | - | - |
| Y0, Y1, YN: | | | | | |
| (x < 0) | {-HUGE_VAL} | - | - | - | - |
| (x = 0) | - | {-HUGE_VAL} | - | - | - |
| (x > X_TLOSS) | - | - | - | - | 0.0 |

Abbreviations

{...} errno is not to be relied upon in all braced cases.
 NaN IEEE NaN result and invalid operation exception.
 HUGE_VAL IEEE ∞ result and division-by-zero exception.
 X_TLOSS The value X_TLOSS is defined in <values.h>.

ANSI/ISO-C STANDARD CONFORMANCE ANSI/ISO-C ERROR HANDLING PROCEDURES (compile with cc -xc)

The ANSI/ISO-C standard covers a small subset of the CAE specification.
 The following table summarizes the values returned in the exceptional cases.

| <math.h> type | DOMAIN | SING | OVERFLOW | UNDERFLOW |
|-----------------------|--------|------|-----------|-----------|
| errno | EDOM | EDOM | ERANGE | ERANGE |
| ACOS, ASIN (x > 1): | 0.0 | - | - | - |
| ATAN2 (0,0): | 0.0 | - | - | - |
| EXP: | - | - | +HUGE_VAL | 0.0 |
| FMOD (x,0): | NaN | - | - | - |
| LOG, LOG10: | | | | |

| <math.h> type | DOMAIN | SING | OVERFLOW | UNDERFLOW |
|-------------------------------|-----------|-----------|-----------|-----------|
| (x < 0) | -HUGE_VAL | - | - | - |
| (x = 0) | - | -HUGE_VAL | - | - |
| POW: | | | | |
| usual cases | - | - | ±HUGE_VAL | ±0.0 |
| (x < 0) ** (y not an integer) | 0.0 | - | - | - |
| 0 ** (y < 0) | -HUGE_VAL | - | - | - |
| SQRT (x < 0): | 0.0 | - | - | - |

ABBREVIATIONS

NaN IEEE NaN result and invalid operation exception.
 HUGE_VAL IEEE ∞ result and division-by-zero

EXAMPLES

EXAMPLE 1 Example of matherr() function

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int
matherr(struct exception *x) {
    switch (x->type) {
        case DOMAIN:
            /* change sqrt to return sqrt(-arg1), not NaN */
            if (!strcmp(x->name, "sqrt")) {
                x->retval = sqrt(-x->arg1);
                return (0); /* print message and set errno */
            } /* FALLTHRU */
        case SING:
            /* all other domain or sing exceptions, print message and */
            /* abort */
            fprintf(stderr, "domain exception in %s\n", x->name);
            abort( );
            break;
    }
    return (0); /* all other exceptions, execute default procedure */
}
```

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

attributes(5), standards(5)

| | |
|----------------------|---|
| NAME | m_create_layout – initialize a layout object |
| SYNOPSIS | cc [<i>flag...</i>] <i>file...</i> -llayout [<i>library...</i>] #include <sys/layout.h> |
| DESCRIPTION | <p>LayoutObject m_create_layout(const AttrObject <i>attrobj</i>, const char* <i>modifier</i>);</p> <p>The m_create_layout() function creates a LayoutObject associated with the locale identified by <i>attrobj</i>.</p> <p>The LayoutObject is an opaque object containing all the data and methods necessary to perform the layout operations on context-dependent or directional characters of the locale identified by the <i>attrobj</i>. The memory for the LayoutObject is allocated by m_create_layout(). The LayoutObject created has default layout values. If the <i>modifier</i> argument is not NULL, the layout values specified by the <i>modifier</i> overwrite the default layout values associated with the locale. Internal states maintained by the layout transformation function across transformations are set to their initial values.</p> <p>The <i>attrobj</i> argument is or may be an amalgam of many opaque objects. A locale object is just one example of the type of object that can be attached to an attribute object. The <i>attrobj</i> argument specifies a name that is usually associated with a locale category. If <i>attrobj</i> is NULL, the created LayoutObject is associated with the current locale as set by the setlocale(3C) function.</p> <p>The <i>modifier</i> argument announces a set of layout values when the LayoutObject is created.</p> |
| RETURN VALUES | Upon successful completion, the m_create_layout () function returns a LayoutObject for use in subsequent calls to m*_layout () functions. Otherwise the m_create_layout () function returns (LayoutObject) 0 and sets errno to indicate the error. |
| ERRORS | <p>The m_create_layout() function may fail if:</p> <p>EBADF The attribute object is invalid or the locale associated with the attribute object is not available.</p> <p>EINVAL The <i>modifier</i> string has a syntax error or it contains unknown layout values.</p> <p>EMFILE There are {OPEN_MAX} file descriptors currently open in the calling process.</p> <p>ENOMEM Insufficient storage space is available.</p> |
| ATTRIBUTES | See attributes (5) for descriptions of the following attributes: |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

setlocale(3C), attributes(5)

| | |
|----------------------|--|
| NAME | md5, MD5Init, MD5Update, MD5Final, md5_calc – MD5 hashing functions |
| SYNOPSIS | <pre>cc [flag ...] file ... -lmd5 [library ...] #include <md5.h> void MD5Init(MD5_CTX *context); void MD5Update(MD5_CTX *context, unsigned char *input, unsigned int inlen); void MD5Final(unsigned char *output, MD5_CTX *context); void md5_calc(unsigned char *output, unsigned char *input, unsigned int inlen);</pre> |
| DESCRIPTION | <p>The MD5Init() initializes an MD5 context structure that is used as an input argument to the MD5Update() and MD5Final() functions. This function must be called every time a new hash needs to be computed.</p> <p>The MD5Update() function updates the MD5 context buffer that will be used in the final hash output.</p> <p>The MD5Final() function generates the final MD5 hash, using the md5 context that was updated in the calls to MD5Update() .</p> <p>These function should be called if scatter/gather buffer support is required. For applications that need to hash a single contiguous buffer, md5_calc() provides a single call to generate the MD5 hash.</p> |
| RETURN VALUES | These functions do not return a value. |
| EXAMPLES | <p>EXAMPLE 1 Authenticate a message found in multiple buffers</p> <p>The following is a sample function that must authenticate a message that is found in multiple buffers. The calling function provides an authentication buffer that will contain the result of the MD5 hash.</p> <pre>int AuthenticateMsg(unsigned char *auth_buffer, struct iovec *messageIov, unsigned int num_buffers) { MD5_CTX md5_context; unsigned int i; MD5Init(&md5_context); for(i=0, i<num_buffers; i++) { MD5Update(&md5_context, messageIov->iiov_base, messageIov->iiov_len); messageIov += sizeof(struct iovec); } MD5Final(auth_buffer, &md5_context); return 0; }</pre> |

CODE EXAMPLE 1 Use `md5_calc()` to generate the MD5 hash

Since the buffer to be computed is contiguous, the `md5_calc()` function can be used to generate the MD5 hash.

```
int AuthenticateMsg(unsigned char *auth_buffer, unsigned
                   char *buffer, unsigned int length)
{
    md5_calc(buffer, auth_buffer, length);

    return (0);
}
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

NAME | m_destroy_layout – destroy a layout object

SYNOPSIS | cc [*flag...*] *file...* -llayout [*library...*]
 | #include <sys/layout.h>

DESCRIPTION | int **m_destroy_layout**(const LayoutObject *layoutobject*);
 | The **m_destroy_layout**() function destroys a LayoutObject by
 | deallocating the layout object and all the associated resources previously
 | allocated by the **m_create_layout**(3LAYOUT) function.

RETURN VALUES | Upon successful completion, 0 is returned. Otherwise -1 is returned and *errno*
 | is set to indicate the error.

ERRORS | The **m_destroy_layout**() function may fail if:
 | EBADF The attribute object is erroneous.
 | EFAULT Errors occurred while processing the request.

ATTRIBUTES | See **attributes**(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | **m_create_layout**(3LAYOUT), **attributes**(5)

| | | | | | | | | | | | |
|-----------------------|---|-----------------|--|-----------------|--|--------------------|---|-----------------------|---|-------------------|---|
| NAME | media_findname – convert a supplied name into an absolute pathname that can be used to access removable media | | | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lvolmgt [library ...] #include <volmgt.h></pre> | | | | | | | | | | |
| DESCRIPTION | <p>char *media_findname(char *start);</p> <p>media_findname() converts the supplied <i>start</i> string into an absolute pathname that can then be used to access a particular piece of media.</p> <p>The <i>start</i> parameter can be one of the following types of specifications:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><i>/dev/...</i></td> <td>An absolute pathname in /dev, such as /dev/rdiskette0, in which case a copy of that string is returned (see NOTES on this page).</td> </tr> <tr> <td style="padding-right: 20px;"><i>/vol/...</i></td> <td>An absolute Volume Management pathname, such as /vol/dev/aliases/floppy0 or /vol/dsk/fred. If this supplied pathname is not a symbolic link, then a copy of that pathname is returned. If the supplied pathname is a symbolic link then it is dereferenced and a copy of that dereferenced pathname is returned.</td> </tr> <tr> <td style="padding-right: 20px;"><i>volume_name</i></td> <td>The Volume Management volume name for a particular volume, such as fred (see fdformat(1) for a description of how to label floppies). In this case a pathname in the Volume Management namespace is returned.</td> </tr> <tr> <td style="padding-right: 20px;"><i>volmgt_symname</i></td> <td>The Volume Management symbolic name for a device, such as floppy0 or cdrom2 (see volfs(7FS) for more information on Volume Management symbolic names), in which case a pathname in the Volume Management namespace is returned.</td> </tr> <tr> <td style="padding-right: 20px;"><i>media_type</i></td> <td>The Volume Management generic media type name. For example, floppy or cdrom. In this case media_findname() looks for the first piece of media that matches that media type, starting at 0 (zero) and continuing on until a match is found (or some fairly large maximum number is reached). In this case, if a match is found, a copy of the pathname to the volume found is returned.</td> </tr> </table> | <i>/dev/...</i> | An absolute pathname in /dev, such as /dev/rdiskette0, in which case a copy of that string is returned (see NOTES on this page). | <i>/vol/...</i> | An absolute Volume Management pathname, such as /vol/dev/aliases/floppy0 or /vol/dsk/fred. If this supplied pathname is not a symbolic link, then a copy of that pathname is returned. If the supplied pathname is a symbolic link then it is dereferenced and a copy of that dereferenced pathname is returned. | <i>volume_name</i> | The Volume Management volume name for a particular volume, such as fred (see fdformat(1) for a description of how to label floppies). In this case a pathname in the Volume Management namespace is returned. | <i>volmgt_symname</i> | The Volume Management symbolic name for a device, such as floppy0 or cdrom2 (see volfs(7FS) for more information on Volume Management symbolic names), in which case a pathname in the Volume Management namespace is returned. | <i>media_type</i> | The Volume Management generic media type name. For example, floppy or cdrom. In this case media_findname() looks for the first piece of media that matches that media type, starting at 0 (zero) and continuing on until a match is found (or some fairly large maximum number is reached). In this case, if a match is found, a copy of the pathname to the volume found is returned. |
| <i>/dev/...</i> | An absolute pathname in /dev, such as /dev/rdiskette0, in which case a copy of that string is returned (see NOTES on this page). | | | | | | | | | | |
| <i>/vol/...</i> | An absolute Volume Management pathname, such as /vol/dev/aliases/floppy0 or /vol/dsk/fred. If this supplied pathname is not a symbolic link, then a copy of that pathname is returned. If the supplied pathname is a symbolic link then it is dereferenced and a copy of that dereferenced pathname is returned. | | | | | | | | | | |
| <i>volume_name</i> | The Volume Management volume name for a particular volume, such as fred (see fdformat(1) for a description of how to label floppies). In this case a pathname in the Volume Management namespace is returned. | | | | | | | | | | |
| <i>volmgt_symname</i> | The Volume Management symbolic name for a device, such as floppy0 or cdrom2 (see volfs(7FS) for more information on Volume Management symbolic names), in which case a pathname in the Volume Management namespace is returned. | | | | | | | | | | |
| <i>media_type</i> | The Volume Management generic media type name. For example, floppy or cdrom. In this case media_findname() looks for the first piece of media that matches that media type, starting at 0 (zero) and continuing on until a match is found (or some fairly large maximum number is reached). In this case, if a match is found, a copy of the pathname to the volume found is returned. | | | | | | | | | | |

RETURN VALUES

Upon successful completion `media_findname()` returns a pointer to the pathname found. In the case of an error a null pointer is returned.

ERRORS

For cases where the supplied *start* parameter is an absolute pathname, `media_findname()` can fail, returning a null string pointer, if an `lstat(2)` of that supplied pathname fails. Also, if the supplied absolute pathname is a symbolic link, `media_findname()` can fail if a `readlink(2)` of that symbolic link fails, or if a `stat(2)` of the pathname pointed to by that symbolic link fails, or if any of the following is true:

ENXIO The specified absolute pathname was not a character special device, and it was not a directory with a character special device in it.

EXAMPLES

EXAMPLE 1 Sample programs of the `media_findname()` function.

The following example attempts to find what the Volume Management pathname is to a piece of media called fred. Notice that a `volmgt_check()` is done first (see the NOTES section on this page).

```
(void) volmgt_check(NULL);
if ((nm = media_findname("fred")) != NULL) {
    (void) printf("media named \"fred\" is at \"%s\"\n", nm);
} else {
    (void) printf("media named \"fred\" not found\n");
}
```

This example looks for whatever volume is in the first floppy drive, letting `media_findname()` call `volmgt_check()` if and only if no floppy is currently known to be the first floppy drive.

```
if ((nm = media_findname("floppy0")) != NULL) {
    (void) printf("path to floppy0 is \"%s\"\n", nm);
} else {
    (void) printf("nothing in floppy0\n");
}
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Unsafe |

SEE ALSO

`cc(1B)`, `fdformat(1)`, `vold(1M)`, `lstat(2)`, `readlink(2)`, `stat(2)`, `free(3C)`, `malloc(3C)`, `volmgt_check(3VOLMGT)`, `volmgt_inuse(3VOLMGT)`, `volmgt_root(3VOLMGT)`, `volmgt_running(3VOLMGT)`, `volmgt_symname(3VOLMGT)`, `attributes(5)`, `volfs(7FS)`

NOTES

If `media_findname()` cannot find a match for the supplied name, it performs a `volmgt_check(3VOLMGT)` and tries again, so it can be more efficient to perform `volmgt_check()` before calling `media_findname()`.

Upon success `media_findname()` returns a pointer to string which has been allocated; this should be freed when no longer in use (see `free(3C)`).

NAME media_getattr, media_setattr – get and set media attributes

SYNOPSIS

```
cc [ flag ... ] file ... -lvolmgt [ library ... ]
#include <volmgt.h>
char *media_getattr(char *vol_path, char *attr);
```

```
int media_setattr(char *vol_path, char *attr, char *value);
```

DESCRIPTION media_setattr() and media_getattr() respectively set and get attribute-value pairs (called properties) on a per-volume basis.

Volume Management supports system properties and user properties. System properties are ones that Volume Management predefines. Some of these system properties are writable, but only by the user that owns the volume being specified, and some system properties are read only:

| Attribute | Writable | Value | Description |
|--------------|----------|---|---|
| s-access | RO | "seq", "rand" | sequential or random access |
| s-density | RO | "low", "medium", "high" | media density |
| s-parts | RO | comma separated list of slice numbers | list of partitions on this volume |
| s-location | RO | <i>pathname</i> | Volume Management pathname to media |
| s-mejectable | RO | "true", "false" | whether or not media is manually ejectable |
| s-rmoneject | R/W | "true", "false" | should media access points be removed from database upon ejection |
| s-enxio | R/W | "true", "false" | if set return ENXIO when media access attempted |

Properties can also be defined by the user. In this case the value can be any string the user wishes.

RETURN VALUES Upon successful completion `media_getattr()` returns a pointer to the value corresponding to the specified attribute. A null pointer is returned if the specified volume doesn't exist, if the specified attribute for that volume doesn't exist, if the specified attribute is boolean and its value is false, or if `malloc(3C)` fails to allocate space for the return value.

`media_setattr()` returns 1 upon success, and 0 upon failure.

ERRORS

Both `media_getattr()` and `media_setattr()` can fail returning a null pointer if an `open(2)` of the specified `vol_path` fails, if an `fstat(2)` of that pathname fails, or if that pathname is not a block or character special device.

`media_getattr()` can also fail if the specified attribute was not found, and `media_setattr()` can also fail if the caller doesn't have permission to set the attribute, either because it's is a system attribute, or because the caller doesn't own the specified volume.

Additionally, either routine can fail returning the following error values:

- ENXIO The Volume Management daemon, `vold`, is not running
- EINTR The routine was interrupted by the user before finishing

EXAMPLES

EXAMPLE 1 Using `media_getattr()`

The following example checks to see if the volume called *fred* that Volume Management is managing can be ejected by means of software, or if it can only be manually ejected:

```

if (media_getattr("/vol/rdisk/fred", "s-mejectable") != NULL) {
    (void) printf("\\"fred\\" must be manually ejected\
");
} else {
    (void) printf("software can eject\\"fred\\"");
}

```

This example shows setting the *s-enxio* property for the floppy volume currently in the first floppy drive:

```

int    res;
if ((res = media_setattr("/vol/dev/aliases/floppy0", "s-enxio",
"true")) == 0) {
    (void) printf("can't set s-enxio flag for floppy0\
");
}

```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`cc(1B)`, `vold(1M)`, `lstat(2)`, `open(2)`, `readlink(2)`, `stat(2)`, `free(3C)`, `malloc(3C)`, `media_findname(3VOLMGT)`, `volmgt_check(3VOLMGT)`, `volmgt_inuse(3VOLMGT)`, `volmgt_root(3VOLMGT)`, `volmgt_running(3VOLMGT)`, `volmgt_symname(3VOLMGT)`, `attributes(5)`

NOTES

Upon success `media_getattr()` returns a pointer to a string which has been allocated, and should be freed when no longer in use (see `free(3C)`).

NAME | media_getid – return the id of a piece of media

SYNOPSIS | cc [flag ...] file ...-lvolgmt [library ...]

| #include <volmgt.h>

| ulonglong_t media_getid(char *vol_path);

DESCRIPTION | media_getid() returns the *id* of a piece of media. Volume Management must be running. See volmgt_running(3VOLMGT).

PARAMETERS | *vol_path* Path to the block or character special device.

RETURN VALUES | media_getid() returns the *id* of the volume. This value is unique for each volume. If media_getid() returns 0, the *path* provided is not valid, for example, it is a block or char device.

EXAMPLES | **EXAMPLE 1** Using media_getid()

| The following example first checks if Volume Management is running, then checks the volume management name space for *path*, and then returns the *id* for the piece of media.

|

```
char *path;
```

| ...

|

```
if (volmgt_running()) {
    if (volmgt_ownspath(path)) {
        (void) printf("id of %s is %lld\n",
            path, media_getid(path));
    }
}
```

| If a program using media_getid() does not check whether or not Volume Management is running, then any NULL return value will be ambiguous, as it could mean that either Volume Management does not have *path* in its name space, or Volume Management is not running.

ATTRIBUTES | See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|------------------|-----------------|
| MT Level | Safe |
| Commitment Level | Public |

SEE ALSO | volmgt_ownspath(3VOLMGT), volmgt_running(3VOLMGT), attributes(5)

| NAME | m_getvalues_layout – query layout values of a LayoutObject | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag...</i>] <i>file...</i> -l <i>library...</i>] #include <sys/layout.h> | | | | |
| DESCRIPTION | <p>int m_getvalues_layout(const LayoutObject <i>layout_object</i>, LayoutValues <i>values</i>, int *<i>index_returned</i>);</p> <p>The m_getvalues_layout() function queries the current setting of layout values within a LayoutObject.</p> <p>The <i>layout_object</i> argument specifies a LayoutObject returned by the m_create_layout(3LAYOUT) function.</p> <p>The <i>values</i> argument specifies the list of layout values that are to be queried. Each value element of a LayoutValueRec must point to a location where the layout value is stored. That is, if the layout value is of type T, the argument must be of type T*. The values are queried from the LayoutObject and represent its current state.</p> <p>It is the user's responsibility to manage the space allocation for the layout values queried. If the layout value name has QueryValueSize OR-ed to it, instead of the value of the layout value, only its size is returned. The caller can use this option to determine the amount of memory needed to be allocated for the layout values queried.</p> | | | | |
| RETURN VALUES | Upon successful completion, the m_getvalues_layout () function returns 0. If any value cannot be queried, the index of the value causing the error is returned in <i>index_returned</i> , -1 is returned and <i>errno</i> is set to indicate the error. | | | | |
| ERRORS | <p>The m_getvalues_layout() function may fail if:</p> <p>EINVAL The layout value specified by <i>index_returned</i> is unknown, its value is invalid, or the <i>layout_object</i> argument is invalid. In the case of an invalid <i>layout_object</i> argument, the value returned in <i>index_returned</i> is -1.</p> | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | m_create_layout(3LAYOUT) , attributes(5) | | | | |

NAME mkdirp, rmdirp – create or remove directories in a path

SYNOPSIS cc [*flag ...*] *file ...* -lgen [*library ...*]
#include <libgen.h>
int mkdirp(const char *path, mode_t mode);
int rmdirp(char *dir, char *dir1);

DESCRIPTION The mkdirp() function creates all the missing directories in *path* with *mode* . See chmod(2) for the values of *mode* .
The rmdirp() function removes directories in path *dir* . This removal begins at the end of the path and moves backward toward the root as far as possible. If an error occurs, the remaining path is stored in *dir1* .

RETURN VALUES If *path* already exists or if a needed directory cannot be created, mkdirp() returns -1 and sets errno to one of the error values listed for mkdir(2) . It returns zero if all the directories are created.
The rmdirp() function returns 0 if it is able to remove every directory in the path. It returns -2 if a “ . ” or “ . . ” is in the path and -3 if an attempt is made to remove the current directory. Otherwise it returns -1 .

EXAMPLES **EXAMPLE 1** Example of creating scratch directories.
The following example creates scratch directories.

```

/* create scratch directories */
if(mkdirp("/tmp/sub1/sub2/sub3", 0755) == -1) {
    fprintf(stderr, "cannot create directory");
    exit(1);
}
chdir("/tmp/sub1/sub2/sub3");
.
.
.
/* cleanup */
chdir("/tmp");
rmdirp("sub1/sub2/sub3");

```

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO chmod(2) , mkdir(2) , rmdir(2) , malloc(3C) , attributes(5)

NOTES mkdirp() uses malloc(3C) to allocate temporary space for the string.
When compiling multithreaded applications, the _REENTRANT flag must be defined on the compile line. This flag should only be used in multithreaded applications.

NAME | mp, mp_madd, mp_msub, mp_mult, mp_mdiv, mp_mcmp, mp_min, mp_mout,
mp_pow, mp_gcd, mp_rpow, mp_itom, mp_xtom, mp_mtox, mp_mfree –
multiple precision integer arithmetic

SYNOPSIS

```
cc [ flag ... ] file ... -lmp [ library ... ]
#include <mp.h>
void mp_madd(MINT *a, MINT *b, MINT *c);
void mp_msub(MINT *a, MINT *b, MINT *c);
void mp_mult(MINT *a, MINT *b, MINT *c);
void mp_mdiv(MINT *a, MINT *b, MINT *q, MINT *r);
int mp_mcmp(MINT *a, MINT *b);
int mp_min(MINT *a);
void mp_mout(MINT *a);
void mp_pow(MINT *a, MINT *b, MINT *c, MINT *d);
void mp_gcd(MINT *a, MINT *b, MINT *c);
void mp_rpow(MINT *a, short n, MINT *b);
int mp_msqrt(MINT *a, MINT *b, MINT *r);
void mp_sdiv(MINT *a, short n, MINT *q, short *r);
MINT * mp_itom(short n);
MINT * mp_xtom(char *a);
char * mp_mtox(MINT *a);
void mp_mfree(MINT *a);
```

DESCRIPTION | These routines perform arithmetic on integers of arbitrary length. The integers are stored using the defined type `MINT`. Pointers to a `MINT` should be initialized using the function `mp_itom(n)`, which sets the initial value to *n*. Alternatively, `mp_xtom(a)` may be used to initialize a `MINT` from a string of hexadecimal digits. `mp_mfree(a)` may be used to release the storage allocated by the `mp_itom(a)` and `mp_xtom(a)` routines.

The `mp_madd(a,b,c)`, `mp_msub(a,b,c)` and `mp_mult(a,b,c)` functions assign to their third arguments the sum, difference, and product, respectively, of their first two arguments. The `mp_mdiv(a,b,q,r)` function assigns the quotient and remainder, respectively, to its third and fourth arguments. The `mp_sdiv(a,n,q,r)` function is similar to `mp_mdiv(a,b,q,r)` except that the divisor is an ordinary integer. The `mp_msqrt(a,b,r)` function produces the square root and remainder of its first argument. The `mp_mcmp(a,b)` function compares the values of its

arguments and returns 0 if the two values are equal, a value greater than 0 if the first argument is greater than the second, and a value less than 0 if the second argument is greater than the first. The `mp_rpow(a,n,b)` function raises `a` to the `n`th power and assigns this value to `b`. The `mp_pow(a,b,c,d)` function raises `a` to the `b`th power, reduces the result modulo `c` and assigns this value to `d`. The `mp_min(a)` and `mp_mout(a)` functions perform decimal input and output. The `mp_gcd(a,b,c)` function finds the greatest common divisor of the first two arguments, returning it in the third argument. The `mp_mtox(a)` function provides the inverse of `mp_xtom(a)`. To release the storage allocated by `mp_mtox(a)` use `free()` (see `malloc(3C)`).

Use the `-lmp` loader option to obtain access to these functions.

FILES

`/usr/lib/libmp.a`
`/usr/lib/libmp.so`

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Unsafe |

SEE ALSO

`exp(3M)`, `malloc(3C)`, `libmp(3LIB)`, `attributes(5)`

DIAGNOSTICS

Illegal operations and running out of memory produce messages and core images.

WARNINGS

The function `pow()` exists in both `libmp` and `libm` with widely differing semantics. This is why `libmp.so.2` exists. `libmp.so.1` exists solely for reasons of backward compatibility, and should not be used otherwise. Use the `mp_*` functions instead. See `libmp(3LIB)`.

| NAME | m_setvalues_layout – set layout values of a LayoutObject | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag...</i>] <i>file...</i> -l <i>library...</i>] #include <sys/layout.h> | | | | |
| DESCRIPTION | <p>int m_setvalues_layout(LayoutObject <i>layout_object</i>, const LayoutValues <i>values</i>, int <i>*index_returned</i>);</p> <p>The <code>m_setvalues_layout()</code> function changes the layout values of a LayoutObject.</p> <p>The <i>layout_object</i> argument specifies a LayoutObject returned by the <code>m_create_layout(3LAYOUT)</code> function.</p> <p>The <i>values</i> argument specifies the list of layout values that are to be changed. The values are written into the LayoutObject and may affect the behavior of subsequent layout functions. Some layout values do alter internal states maintained by a LayoutObject.</p> <p>The <code>m_setvalues_layout()</code> function can be implemented as a macro that evaluates the first argument twice.</p> | | | | |
| RETURN VALUES | Upon successful completion, the requested layout values are set and 0 is returned. Otherwise -1 is returned and <code>errno</code> is set to indicate the error. If any value cannot be set, none of the layout values are changed and the (zero-based) index of the first value causing the error is returned in <i>index_returned</i> . | | | | |
| ERRORS | <p>The <code>m_setvalues_layout()</code> function may fail if:</p> <p><code>EINVAL</code> The layout value specified by <i>index_returned</i> is unknown, its value is invalid, or the <i>layout_object</i> argument is invalid.</p> <p><code>EMFILE</code> There are {<code>OPEN_MAX</code>} file descriptors currently open in the calling process.</p> | | | | |
| USAGE | Do not use expressions with side effects such as auto-increment or auto-decrement within the first argument to the <code>m_setvalues_layout()</code> function. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>m_create_layout(3LAYOUT)</code> , <code>attributes(5)</code> | | | | |

| | | | | | | | | | |
|-----------------------|---|---------------------|--|----------------------|--|-----------------------|---|-----------------------|--|
| NAME | m_transform_layout – layout transformation | | | | | | | | |
| SYNOPSIS | <pre>cc [flag...] file... -llayout [library...] #include <sys/layout.h> int m_transform_layout(LayoutObject layout_object, const char *InpBuf, const size_t InpSize, const void *OutBuf, size_t *Outsize, size_t *InpToOut, size_t *OutToInp, unsigned char *Property, size_t *InpBufIndex);</pre> | | | | | | | | |
| DESCRIPTION | <p>The <code>m_transform_layout()</code> function performs layout transformations (reordering, shaping, cell determination) or provides additional information needed for layout transformation (such as the expected size of the transformed layout, the nesting level of different segments in the text and cross-references between the locations of the corresponding elements before and after the layout transformation). Both the input text and output text are character strings.</p> <p>The <code>m_transform_layout()</code> function transforms the input text in <code>InpBuf</code> according to the current layout values in <code>layout_object</code>. Any layout value whose value type is <code>LayoutTextDescriptor</code> describes the attributes of the <code>InpBuf</code> and <code>OutBuf</code> arguments. If the attributes are the same for both <code>InpBuf</code> and <code>OutBuf</code>, a null transformation is performed with respect to that specific layout value.</p> <p>The <code>InpBuf</code> argument specifies the source text to be processed. The <code>InpBuf</code> may not be <code>NULL</code>, unless there is a need to reset the internal state.</p> <p>The <code>InpSize</code> argument is the number of bytes within <code>InpBuf</code> to be processed by the transformation. Its value will not change after return from the transformation. <code>InpSize</code> set to <code>-1</code> indicates that the text in <code>InpBuf</code> is delimited by a null code element. If <code>InpSize</code> is not set to <code>-1</code>, it is possible to have some null elements in the input buffer. This might be used, for example, for a “one shot” transformation of several strings, separated by nulls.</p> <p>Output of this function may be one or more of the following depending on the setting of the arguments:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><code>OutBuf</code></td> <td>Any transformed data is stored in <code>OutBuf</code>, converted to <code>ShapeCharset</code>.</td> </tr> <tr> <td><code>Outsize</code></td> <td>The number of bytes in <code>OutBuf</code>.</td> </tr> <tr> <td><code>InpToOut</code></td> <td>A cross-reference from each <code>InpBuf</code> code element to the transformed data. The cross-reference relates to the data in <code>InpBuf</code> starting with the first element that <code>InpBufIndex</code> points to (and not necessarily starting from the beginning of the <code>InpBuf</code>).</td> </tr> <tr> <td><code>OutToInp</code></td> <td>A cross-reference to each <code>InpBuf</code> code element from the transformed data. The cross-reference relates to the data in <code>InpBuf</code> starting with the first element that <code>InpBufIndex</code></td> </tr> </table> | <code>OutBuf</code> | Any transformed data is stored in <code>OutBuf</code> , converted to <code>ShapeCharset</code> . | <code>Outsize</code> | The number of bytes in <code>OutBuf</code> . | <code>InpToOut</code> | A cross-reference from each <code>InpBuf</code> code element to the transformed data. The cross-reference relates to the data in <code>InpBuf</code> starting with the first element that <code>InpBufIndex</code> points to (and not necessarily starting from the beginning of the <code>InpBuf</code>). | <code>OutToInp</code> | A cross-reference to each <code>InpBuf</code> code element from the transformed data. The cross-reference relates to the data in <code>InpBuf</code> starting with the first element that <code>InpBufIndex</code> |
| <code>OutBuf</code> | Any transformed data is stored in <code>OutBuf</code> , converted to <code>ShapeCharset</code> . | | | | | | | | |
| <code>Outsize</code> | The number of bytes in <code>OutBuf</code> . | | | | | | | | |
| <code>InpToOut</code> | A cross-reference from each <code>InpBuf</code> code element to the transformed data. The cross-reference relates to the data in <code>InpBuf</code> starting with the first element that <code>InpBufIndex</code> points to (and not necessarily starting from the beginning of the <code>InpBuf</code>). | | | | | | | | |
| <code>OutToInp</code> | A cross-reference to each <code>InpBuf</code> code element from the transformed data. The cross-reference relates to the data in <code>InpBuf</code> starting with the first element that <code>InpBufIndex</code> | | | | | | | | |

Property

points to (and not necessarily starting from the beginning of the *InpBuf*).

A weighted value that represents peculiar input string transformation properties with different connotations as explained below. If this argument is not a null pointer, it represents an array of values with the same number of elements as the source substring text before the transformation. Each byte will contain relevant “property” information of the corresponding element in *InpBuf* starting from the element pointed by *InpBufIndex*. The four rightmost bits of each “property” byte will contain information for bidirectional environments (when `ActiveDirectional` is `True`) and they will mean “`NestingLevels`.” The possible value from 0 to 15 represents the nesting level of the corresponding element in the *InpBuf* starting from the element pointed by *InpBufIndex*. If `ActiveDirectional` is `false` the content of `NestingLevel` bits will be ignored. The leftmost bit of each “property” byte will contain a “new cell indicator” for composed character environments, and will have a value of either 1 (for an element in *InpBuf* that is transformed to the beginning of a new cell) or 0 (for the “zero-length” composing character elements, when these are grouped into the same presentation cell with a non-composing character). Here again, each element of “property” pertains to the elements in the *InpBuf* starting from the element pointed by *InpBufIndex*. (Remember that this is not necessarily the beginning of *InpBuf*). If none of the transformation properties is required, the argument *Property* can be `NULL`. The use of “property” can be enhanced in the future to pertain to other possible usage in other environments.

The *InpBufIndex* argument is an offset value to the location of the transformed text. When `m_transform_layout()` is called, *InpBufIndex* contains the offset to the element in *InpBuf* that will be transformed first. (Note that this is not necessarily the first element in *InpBuf*). At the return from the transformation, *InpBufIndex* contains the offset to the first element in the *InpBuf* that has not been transformed. If the entire substring has been transformed successfully, *InpBufIndex* will be incremented by the amount defined by *InpSize*.

Each of these output arguments may be `NULL` to specify that no output is desired for the specific argument, but at least one of them should be set to a non-null value to perform any significant work.

The layout object maintains a directional state that keeps track of directional changes, based on the last segment transformed. The directional state is maintained across calls to the layout transformation functions and allows stream data to be processed with the layout functions. The directional state is reset to its initial state whenever any of the layout values `TypeOfText`, `Orientation`, or `ImplicitAlg` is modified by means of a call to `m_setvalues_layout()`.

The *layout_object* argument specifies a `LayoutObject` returned by the `m_create_layout()` function.

The *OutBuf* argument contains the transformed data. This argument can be specified as a null pointer to indicate that no transformed data is required.

The encoding of the *OutBuf* argument depends on the `ShapeCharset` layout value defined in *layout_object*. If the `ActiveShapeEditing` layout value is not set (False), the encoding of *OutBuf* is guaranteed to be the same as the codeset of the locale associated with the `LayoutObject` defined by *layout_object*.

On input, the *OutSize* argument specifies the size of the output buffer in number of bytes. The output buffer should be large enough to contain the transformed result; otherwise, only a partial transformation is performed. If the `ActiveShapeEditing` layout value is set (True) the *OutBuf* should be allocated to contain at least the *InpSize* multiplied by `ShapeCharsetSize`.

On return, the *OutSize* argument is modified to the actual number of bytes placed in *OutBuf*.

When the *OutSize* argument is specified as zero, the function calculates the size of an output buffer large enough to contain the transformed text, and the result is returned in this field. The content of the buffers specified by *InpBuf* and *OutBuf*, and the value of *InpBufIndex*, remain unchanged. If *OutSize* = NULL, the `EINVAL` error condition should be returned.

If the *InpToOut* argument is not a null pointer, it points to an array of values with the same number of bytes in *InpBuf* starting with the one pointed by *InpBufIndex* and up to the end of the substring in the buffer. On output, the *nth* value in *InpToOut* corresponds to the *nth* byte in *InpBuf*. This value is the index (in units of bytes) in *OutBuf* that identifies the transformed `ShapeCharset` element of the *nth* byte in *InpBuf*. In the case of multibyte encoding, the index points (for each of the bytes of a code element in the *InpBuf*) to the first byte of the transformed code element in the *OutBuf*.

InpToOut may be specified as NULL if no index array from *InpBuf* to *OutBuf* is desired.

If the *OutToInp* argument is not a null pointer, it points to an array of values with the same number of bytes as contained in *OutBuf*. On output, the *nth* value in *OutToInp* corresponds to the *nth* byte in *OutBuf*. This value is the index in *InpBuf*,

starting with the byte pointed to by *InpBufIndex*, that identifies the logical code element of the *n*th byte in *OutBuf*. In the case of multibyte encoding, the index will point for each of the bytes of a transformed code element in the *OutBuf* to the first byte of the code element in the *InpBuf*.

OutToInp may be specified as NULL if no index array from *OutBuf* to *InpBuf* is desired.

To perform shaping of a text string without reordering of code elements, the *layout_object* should be set with input and output layout value *TypeOfText* set to TEXT_VISUAL and both in and out of Orientation set to the same value.

RETURN VALUES

If successful, the `m_transform_layout()` function returns 0. If unsuccessful, the returned value is -1 and the `errno` is set to indicate the source of error. When the size of *OutBuf* is not large enough to contain the entire transformed text, the input text state at the end of the uncompleted transformation is saved internally and the error condition E2BIG is returned in `errno`.

ERRORS

The `m_transform_layout()` function may fail if:

| | |
|--------|--|
| E2BIG | The output buffer is full and the source text is not entirely processed. |
| EBADF | The layout values are set to a meaningless combination or the layout object is not valid. |
| EILSEQ | Transformation stopped due to an input code element that cannot be shaped or is invalid. The <i>InpBufIndex</i> argument is set to indicate the code element causing the error. The suspect code element is either a valid code element but cannot be shaped into the <i>ShapeCharset</i> layout value, or is an invalid code element not defined by the codeset of the locale of <i>layout_object</i> . The <code>mbtowc()</code> and <code>wctomb()</code> functions, when used in the same locale as the <i>LayoutObject</i> , can be used to determine if the code element is valid. |
| EINVAL | Transformation stopped due to an incomplete composite sequence at the end of the input buffer, or <i>OutSize</i> contains NULL. |
| ERANGE | More than 15 embedding levels are in source text or <i>InpBuf</i> contain unbalanced directional layout information (push/pop) or an incomplete composite sequence has been detected in the input buffer at the beginning of the string pointed to by <i>InpBufIndex</i> . An incomplete composite sequence at the end of the input buffer is not always detectable. Sometimes, the fact that |

the sequence is incomplete will only be detected when additional character elements belonging to the composite sequence are found at the beginning of the next input buffer.

USAGE

A `LayoutObject` will have a meaningful combination of default layout values. Whoever chooses to change the default layout values is responsible for making sure that the combination of layout values is meaningful. Otherwise, the result of `m_transform_layout()` might be unpredictable or implementation-specific with `errno` set to `EBADF`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`attributes(5)`

| | | | | | | | | | |
|--------------------|--|---------------|--|----------------|--|-----------------|---|-----------------|--|
| NAME | m_wtransform_layout – layout transformation for wide character strings | | | | | | | | |
| SYNOPSIS | <pre>cc [flag...] file... -llayout [library...] #include <sys/layout.h></pre> <p>int m_wtransform_layout(LayoutObject <i>layout_object</i>, const wchar_t *<i>InpBuf</i>, const size_t <i>InpSize</i>, const void *<i>OutBuf</i>, size_t *<i>Outsize</i>, size_t *<i>InpToOut</i>, size_t *<i>OutToInp</i>, unsignedchar *<i>Property</i>, size_t *<i>InpBufIndex</i>);</p> | | | | | | | | |
| DESCRIPTION | <p>The <code>m_wtransform_layout()</code> function performs layout transformations (reordering, shaping, cell determination) or provides additional information needed for layout transformation (such as the expected size of the transformed layout, the nesting level of different segments in the text and cross-references between the locations of the corresponding elements before and after the layout transformation). Both the input text and output text are wide character strings.</p> <p>The <code>m_wtransform_layout()</code> function transforms the input text in <i>InpBuf</i> according to the current layout values in <i>layout_object</i>. Any layout value whose value type is <code>LayoutTextDescriptor</code> describes the attributes of the <i>InpBuf</i> and <i>OutBuf</i> arguments. If the attributes are the same for both <i>InpBuf</i> and <i>OutBuf</i>, a null transformation is performed with respect to that specific layout value.</p> <p>The <i>InpBuf</i> argument specifies the source text to be processed. The <i>InpBuf</i> may not be <code>NULL</code>, unless there is a need to reset the internal state.</p> <p>The <i>InpSize</i> argument is the number of bytes within <i>InpBuf</i> to be processed by the transformation. Its value will not change after return from the transformation. <i>InpSize</i> set to <code>-1</code> indicates that the text in <i>InpBuf</i> is delimited by a null code element. If <i>InpSize</i> is not set to <code>-1</code>, it is possible to have some null elements in the input buffer. This might be used, for example, for a “one shot” transformation of several strings, separated by nulls.</p> <p>Output of this function may be one or more of the following depending on the setting of the arguments:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><i>OutBuf</i></td> <td>Any transformed data is stored in <i>OutBuf</i>, converted to <code>ShapeCharset</code>.</td> </tr> <tr> <td><i>Outsize</i></td> <td>The number of wide characters in <i>OutBuf</i>.</td> </tr> <tr> <td><i>InpToOut</i></td> <td>A cross-reference from each <i>InpBuf</i> code element to the transformed data. The cross-reference relates to the data in <i>InpBuf</i> starting with the first element that <i>InpBufIndex</i> points to (and not necessarily starting from the beginning of the <i>InpBuf</i>).</td> </tr> <tr> <td><i>OutToInp</i></td> <td>A cross-reference to each <i>InpBuf</i> code element from the transformed data. The cross-reference relates to the data in <i>InpBuf</i> starting with the first element that <i>InpBufIndex</i></td> </tr> </table> | <i>OutBuf</i> | Any transformed data is stored in <i>OutBuf</i> , converted to <code>ShapeCharset</code> . | <i>Outsize</i> | The number of wide characters in <i>OutBuf</i> . | <i>InpToOut</i> | A cross-reference from each <i>InpBuf</i> code element to the transformed data. The cross-reference relates to the data in <i>InpBuf</i> starting with the first element that <i>InpBufIndex</i> points to (and not necessarily starting from the beginning of the <i>InpBuf</i>). | <i>OutToInp</i> | A cross-reference to each <i>InpBuf</i> code element from the transformed data. The cross-reference relates to the data in <i>InpBuf</i> starting with the first element that <i>InpBufIndex</i> |
| <i>OutBuf</i> | Any transformed data is stored in <i>OutBuf</i> , converted to <code>ShapeCharset</code> . | | | | | | | | |
| <i>Outsize</i> | The number of wide characters in <i>OutBuf</i> . | | | | | | | | |
| <i>InpToOut</i> | A cross-reference from each <i>InpBuf</i> code element to the transformed data. The cross-reference relates to the data in <i>InpBuf</i> starting with the first element that <i>InpBufIndex</i> points to (and not necessarily starting from the beginning of the <i>InpBuf</i>). | | | | | | | | |
| <i>OutToInp</i> | A cross-reference to each <i>InpBuf</i> code element from the transformed data. The cross-reference relates to the data in <i>InpBuf</i> starting with the first element that <i>InpBufIndex</i> | | | | | | | | |

points to (and not necessarily starting from the beginning of the *InpBuf*).

Property

A weighted value that represents peculiar input string transformation properties with different connotations as explained below. If this argument is not a nullpointer, it represents an array of values with the same number of elements as the source substring text before the transformation. Each byte will contain relevant “property” information of the corresponding element in *InpBuf* starting from the element pointed by *InpBufIndex*. The four rightmost bits of each “property” byte will contain information for bidirectional environments (when `ActiveDirectional` is `True`) and they will mean “`NestingLevels`.” The possible value from 0 to 15 represents the nesting level of the corresponding element in the *InpBuf* starting from the element pointed by *InpBufIndex*. If `ActiveDirectional` is false the content of `NestingLevel` bits will be ignored. The leftmost bit of each “property” byte will contain a “new cell indicator” for composed character environments, and will have a value of either 1 (for an element in *InpBuf* that is transformed to the beginning of a new cell) or 0 (for the “zero-length” composing character elements, when these are grouped into the same presentation cell with a non-composing character). Here again, each element of “property” pertains to the elements in the *InpBuf* starting from the element pointed by *InpBufIndex*. (Remember that this is not necessarily the beginning of *InpBuf*). If none of the transformation properties is required, the argument *Property* can be `NULL`. The use of “property” can be enhanced in the future to pertain to other possible usage in other environments.

The *InpBufIndex* argument is an offset value to the location of the transformed text. When `m_wtransform_layout()` is called, *InpBufIndex* contains the offset to the element in *InpBuf* that will be transformed first. (Note that this is not necessarily the first element in *InpBuf*). At the return from the transformation, *InpBufIndex* contains the offset to the first element in the *InpBuf* that has not been transformed. If the entire substring has been transformed successfully, *InpBufIndex* will be incremented by the amount defined by *InpSize*.

Each of these output arguments may be null to specify that no output is desired for the specific argument, but at least one of them should be set to a non-null value to perform any significant work.

In addition to the possible outputs above, *layout_object* maintains a directional state across calls to the transform functions. The directional state is reset to its initial state whenever any of the layout values `TypeOfText`, `Orientation`, or `ImplicitAlg` is modified by means of a call to `m_setvalues_layout()`.

The *layout_object* argument specifies a `LayoutObject` returned by the `m_create_layout()` function.

The *OutBuf* argument contains the transformed data. This argument can be specified as a null pointer to indicate that no transformed data is required.

The encoding of the *OutBuf* argument depends on the `ShapeCharset` layout value defined in *layout_object*. If the `ActiveShapeEditing` layout value is not set (False), the encoding of *OutBuf* is guaranteed to be the same as the codeset of the locale associated with the `LayoutObject` defined by *layout_object*.

On input, the *OutSize* argument specifies the size of the output buffer in number of wide characters. The output buffer should be large enough to contain the transformed result; otherwise, only a partial transformation is performed. If the `ActiveShapeEditing` layout value is set (True) the *OutBuf* should be allocated to contain at least the *InpSize* multiplied by `ShapeCharsetSize`.

On return, the *OutSize* argument is modified to the actual number of code elements in *OutBuf*.

When the *OutSize* argument is specified as zero, the function calculates the size of an output buffer large enough to contain the transformed text, and the result is returned in this field. The content of the buffers specified by *InpBuf* and *OutBuf*, and the value of *InpBufIndex*, remain unchanged. If *OutSize* = `NULL`, the `EINVAL` error condition should be returned.

If the *InpToOut* argument is not a null pointer, it points to an array of values with the same number of wide characters in *InpBuf* starting with the one pointed by *InpBufIndex* and up to the end of the substring in the buffer. On output, the *n*th value in *InpToOut* corresponds to the *n*th byte in *InpBuf*. This value is the index (in units of wide characters) in *OutBuf* that identifies the transformed `ShapeCharset` element of the *n*th byte in *InpBuf*.

InpToOut may be specified as `NULL` if no index array from *InpBuf* to *OutBuf* is desired.

If the *OutToInp* argument is not a null pointer, it points to an array of values with the same number of wide characters as contained in *OutBuf*. On output, the *n*th value in *OutToInp* corresponds to the *n*th byte in *OutBuf*. This value is the index in *InpBuf*, starting with wide character byte pointed to by *InpBufIndex*, that identifies the logical code element of the *n*th wide character in *OutBuf*.

OutToInp may be specified as `NULL` if no index array from *OutBuf* to *InpBuf* is desired.

| | |
|----------------------|--|
| | To perform shaping of a text string without reordering of code elements, the <i>layout_object</i> should be set with input and output layout value <i>TypeOfText</i> set to <code>TEXT_VISUAL</code> and both in and out of <i>Orientation</i> set to the same value. |
| RETURN VALUES | If successful, the <code>m_wtransform_layout()</code> function returns 0. If unsuccessful, the returned value is -1 and the <code>errno</code> is set to indicate the source of error. When the size of <i>OutBuf</i> is not large enough to contain the entire transformed text, the input text state at the end of the uncompleted transformation is saved internally and the error condition <code>E2BIG</code> is returned in <code>errno</code> . |
| ERRORS | The <code>m_wtransform_layout()</code> function may fail if: |
| <code>E2BIG</code> | The output buffer is full and the source text is not entirely processed. |
| <code>EBADF</code> | The layout values are set to a meaningless combination or the layout object is not valid. |
| <code>EILSEQ</code> | Transformation stopped due to an input code element that cannot be shaped or is invalid. The <i>InpBufIndex</i> argument is set to indicate the code element causing the error. The suspect code element is either a valid code element but cannot be shaped into the <i>ShapeCharset</i> layout value, or is an invalid code element not defined by the codeset of the locale of <i>layout_object</i> . The <code>mbtowc()</code> and <code>wctomb()</code> functions, when used in the same locale as the <i>LayoutObject</i> , can be used to determine if the code element is valid. |
| <code>EINVAL</code> | Transformation stopped due to an incomplete composite sequence at the end of the input buffer, or <i>OutSize</i> contains <code>NULL</code> . |
| <code>ERANGE</code> | More than 15 embedding levels are in source text or <i>InpBuf</i> contain unbalanced directional layout information (push/pop) or an incomplete composite sequence has been detected in the input buffer at the beginning of the string pointed to by <i>InpBufIndex</i> . An incomplete composite sequence at the end of the input buffer is not always detectable. Sometimes the fact that the sequence is incomplete will only be detected when additional character elements belonging to the composite sequence are found at the beginning of the next input buffer. |
| USAGE | A <i>LayoutObject</i> will have a meaningful combination of default layout values. Whoever chooses to change the default layout values is responsible for making sure that the combination of layout values is meaningful. Otherwise, the result of |

`m_wtransform_layout()` might be unpredictable or implementation-specific with `errno` set to `EBADF`.

EXAMPLES**EXAMPLE 1** Shaping and reordering input string into output buffer

The following example illustrates what the different arguments of `m_wtransform_layout()` look like when a string in *InpBuf* is shaped and reordered into *OutBuf*. Upper-case letters in the example represent left-to-right letters while lower-case letters represent right-to-left letters. *xyz* represents the shapes of *cde*.

```

Position:          0123456789
InpBuf:           AB cde 12z

Position:          0123456789
OutBuf:           AB 12 zyxZ

Position:          0123456789
OutToInp:         0127865439

Position:          0123456789
Property.NestLevel: 0001111220
Property.CelBdry:  1111111111

```

The values (encoded in binary) returned in the *Property* argument define the directionality of each code element in the source text as defined by the type of algorithm used within the *layout_object*. While the algorithm may be implementation dependent, the resulting values and levels are defined such as to allow a single method to be used in determining the directionality of the source text. The base rules are:

- Odd levels are always RTL.
- Even levels are always LTR.
- The *Orientation* layout value setting determines the initial level (0 or 1) used.

Within a *Property* array each increment in the level indicates the corresponding code elements should be presented in the opposite direction. Callers of this function should realize that the *Property* values for certain code elements is dependent on the context of the given character and the layout values: *Orientation* and *ImplicitAlg*. Callers should not assume that a given code element always has the same *Property* value in all cases.

EXAMPLE 2 Algorithm to handle nesting

The following is an example of a standard presentation algorithm that handles nesting correctly. The goal of this algorithm is ultimately to return to a zero nest level. Note that more efficient algorithms do exist; the following is provided for clarity rather than for efficiency.

1. Search for the highest next level in the string.
2. Reverse all surrounding code elements of the same level. Reduce the nest level of these code elements by 1.
3. Repeat 1 and 2 until all code elements are of level 0.

The following shows the progression of the example from above:

```

Position:          0123456789    0123456789    0123456789
InpBuf:           AB cde 12Z    AB cde 21Z    AB 12 edcZ
Property.NestLevel: 0001111220    0001111110    0000000000
Property.CellBdry: 1111111111    1111111111    1111111111
    
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`attributes(5)`

NAME | newDmiOctetString – create DmiOctetString in dynamic memory

SYNOPSIS | `cc [flag ...] file ... -ldmi -lnsl -lrwtool [library ...]`
`#include <dmi/util.hh>`
`DmiOctetString_t *newDmiOctetString(DmiOctetString_t *str);`

DESCRIPTION | The `newDmiOctetString()` function creates a `DmiOctetString` in dynamic memory and returns a pointer to the newly created `DmiOctetString`. The function returns `NULL` if no memory is available.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | MT-Safe |

SEE ALSO | `libdmi(3LIB)`, `attributes(5)`

NAME | newDmiString – create DmiString in dynamic memory

SYNOPSIS | `cc [flag ...] file ... -ldmi -lnsl -lrwtool [library ...]`
`#include <dmi/util.hh>`
`DmiString_t *newDmiString(char *str);`

DESCRIPTION | The `newDmiString()` function creates a `DmiString` in dynamic memory and returns a pointer to the newly created `DmiString`. The function returns `NULL` if no memory is available.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | MT-Safe |

SEE ALSO | `freeDmiString(3DMI)`, `libdmi(3LIB)`, `attributes(5)`

| NAME | nextafter – next representable double-precision floating-point number | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double nextafter (double <i>x</i> , double <i>y</i>); | | | | |
| DESCRIPTION | The <code>nextafter()</code> function computes the next representable double-precision floating-point value following <i>x</i> in the direction of <i>y</i> . Thus, if <i>y</i> is less than <i>x</i> , <code>nextafter()</code> returns the largest representable floating-point number less than <i>x</i> . | | | | |
| RETURN VALUES | The <code>nextafter()</code> function returns the next representable double-precision floating-point value following <i>x</i> in the direction of <i>y</i> . If <i>x</i> or <i>y</i> is NaN, then <code>nextafter()</code> returns NaN. If <i>x</i> is finite and the correct function value would overflow, <code>nextafter()</code> returns <code>±HUGE_VAL</code> (according to the sign of <i>x</i>) and sets <code>errno</code> to <code>ERANGE</code> . | | | | |
| ERRORS | The <code>nextafter()</code> function will fail if: <code>ERANGE</code> The correct value would overflow. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>attributes(5)</code> , | | | | |

| NAME | nlist – get entries from name list | | | | |
|----------------------|---|----------------|-----------------|----------|------|
| SYNOPSIS | <pre>cc [flag...] file ... -lelf [library ...] #include <nlist.h> int nlist(const char *filename, struct nlist *nl);</pre> | | | | |
| DESCRIPTION | <p>nlist() examines the name list in the executable file whose name is pointed to by <i>filename</i>, and selectively extracts a list of values and puts them in the array of nlist() structures pointed to by <i>nl</i>. The name list <i>nl</i> consists of an array of structures containing names of variables, types, and values. The list is terminated with a null name, that is, a null string is in the name position of the structure. Each variable name is looked up in the name list of the file. If the name is found, the type, value, storage class, and section number of the name are inserted in the other fields. The type field may be set to 0 if the file was not compiled with the <i>-g</i> option to cc(1B).</p> <p>nlist() will always return the information for an external symbol of a given name if the name exists in the file. If an external symbol does not exist, and there is more than one symbol with the specified name in the file (such as static symbols defined in separate files), the values returned will be for the last occurrence of that name in the file. If the name is not found, all fields in the structure except <i>n_name</i> are set to 0.</p> <p>This function is useful for examining the system name list kept in the file <i>/dev/ksyms</i>. In this way programs can obtain system addresses that are up to date.</p> | | | | |
| RETURN VALUES | <p>All value entries are set to 0 if the file cannot be read or if it does not contain a valid name list.</p> <p>nlist() returns 0 on success, -1 on error.</p> | | | | |
| ATTRIBUTES | <p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | Safe | | | | |
| SEE ALSO | <p>cc(1B), elf(3ELF), kvm_nlist(3KVM), kvm_open(3KVM), a.out(4), attributes(5), ksyms(7D), mem(7D)</p> | | | | |

| | |
|--------------------|--|
| NAME | NOTE, _NOTE – annotate source code with info for tools |
| SYNOPSIS | <pre>#include <note.h> NOTE(NoteInfo); or #include<sys/note.h> _NOTE(NoteInfo);</pre> |
| DESCRIPTION | <p>These macros are used to embed information for tools in program source. A use of one of these macros is called an "annotation". A tool may define a set of such annotations which can then be used to provide the tool with information that would otherwise be unavailable from the source code.</p> <p>Annotations should, in general, provide documentation useful to the human reader. If information is of no use to a human trying to understand the code but is necessary for proper operation of a tool, use another mechanism for conveying that information to the tool (one which does not involve adding to the source code), so as not to detract from the readability of the source. The following is an example of an annotation which provides information of use to a tool and to the human reader (in this case, which data are protected by a particular lock, an annotation defined by the static lock analysis tool <code>lock_lint</code>).</p> <pre>NOTE(MUTEX_PROTECTS_DATA(foo_lock, foo_list Foo))</pre> <p>Such annotations do not represent executable code; they are neither statements nor declarations. They should not be followed by a semicolon. If a compiler or tool that analyzes C source does not understand this annotation scheme, then the tool will ignore the annotations. (For such tools, <code>NOTE(X)</code> expands to nothing.)</p> <p>Annotations may only be placed at particular places in the source. These places are where the following C constructs would be allowed:</p> <ul style="list-style-type: none"> ■ a top-level declaration (that is, a declaration not within a function or other construct) ■ a declaration or statement within a block (including the block which defines a function) ■ a member of a <code>struct</code> or <code>union</code>. <p>Annotations are not allowed in any other place. For example, the following are illegal:</p> <pre>x = y + NOTE(...) z ; typedef NOTE(...) unsigned int uint ;</pre> |

NOTE vs _NOTE

While `NOTE` and `_NOTE` may be used in the places described above, a particular type of annotation may only be allowed in a subset of those places. For example, a particular annotation may not be allowed inside a `struct` or `union` definition. Ordinarily, `NOTE` should be used rather than `_NOTE`, since use of `_NOTE` technically makes a program non-portable. However, it may be inconvenient to use `NOTE` for this purpose in existing code if `NOTE` is already heavily used for another purpose. In this case one should use a different macro and write a header file similar to `/usr/include/note.h` which maps that macro to `_NOTE` in the same manner. For example, the following makes `FOO` such a macro:

```
#ifndef _FOO_H
#define _FOO_H
#define FOO _NOTE
#include <sys/note.h>
#endif
```

Public header files which span projects should use `_NOTE` rather than `NOTE`, since `NOTE` may already be used by a program which needs to include such a header file.

NoteInfo Argument

The actual *NoteInfo* used in an annotation should be specified by a tool that deals with program source (see the documentation for the tool to determine which annotations, if any, it understands).

NoteInfo must have one of the following forms:

```
NoteName
NoteName( Args )
```

where *NoteName* is simply an identifier which indicates the type of annotation, and *Args* is something defined by the tool that specifies the particular *NoteName*. The general restrictions on *Args* are that it be compatible with an ANSI C tokenizer and that unquoted parentheses be balanced (so that the end of the annotation can be determined without intimate knowledge of any particular annotation).

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Safe |

SEE ALSO

`note(4)`, `attributes(5)`

| | |
|----------------------|---|
| NAME | p2open, p2close – open, close pipes to and from a command |
| SYNOPSIS | <pre>cc [flag ...] file ... -lgen [library ...] #include <libgen.h> int p2open(const char *cmd, FILE *fp [2]); int p2close(FILE *fp [2]);</pre> |
| DESCRIPTION | <p>p2open() forks and execs a shell running the command line pointed to by <i>cmd</i>. On return, <i>fp</i>[0] points to a FILE pointer to write the command's standard input and <i>fp</i>[1] points to a FILE pointer to read from the command's standard output. In this way the program has control over the input and output of the command.</p> <p>The function returns 0 if successful; otherwise, it returns -1 .</p> <p>p2close() is used to close the file pointers that p2open() opened. It waits for the process to terminate and returns the process status. It returns 0 if successful; otherwise, it returns -1 .</p> |
| RETURN VALUES | A common problem is having too few file descriptors. p2close() returns -1 if the two file pointers are not from the same p2open() . |
| EXAMPLES | <p>EXAMPLE 1 Example of file descriptors.</p> <pre>#include <stdio.h> #include <libgen.h> main(argc,argv) int argc; char **argv; { FILE *fp[2]; pid_t pid; char buf[16]; pid=p2open("/usr/bin/cat", fp); if (pid == -1) { fprintf(stderr, "p2open failed\ "); exit(1); } write(fileno(fp[0]),"This is a test\ ", 16); if(read(fileno(fp[1]), buf, 16) <=0) fprintf(stderr, "p2open failed\ "); else write(1, buf, 16); (void)p2close(fp); }</pre> |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Unsafe |

SEE ALSO

`fclose(3C)`, `popen(3C)`, `setbuf(3C)`, `attributes(5)`

NOTES

Buffered writes on `fp[0]` can make it appear that the command is not listening. Judiciously placed `fflush()` calls or unbuffering `fp[0]` can be a big help; see `fclose(3C)`.

Many commands use buffered output when connected to a pipe. That, too, can make it appear as if things are not working.

Usage is not the same as for `popen()`, although it is closely related.

| | |
|---------------------------|---|
| NAME | pam – PAM (Pluggable Authentication Module) |
| SYNOPSIS | <pre>#include <security/pam_appl.h> cc [flag...] file ... -lpam [library ...]</pre> |
| DESCRIPTION | <p>The PAM framework, <code>libpam</code>, consists of an interface library and multiple authentication service modules. The PAM interface library is the layer implementing the Application Programming Interface (API). The authentication service modules are a set of dynamically loadable objects invoked by the PAM API to provide a particular type of user authentication. PAM gives system administrators the flexibility of choosing any authentication service available on the system to perform authentication. This framework also allows new authentication service modules to be plugged in and made available without modifying the applications.</p> |
| Interface Overview | <p>The PAM library interface consists of six categories of functions, the names for which all start with the prefix <code>pam_</code>.</p> <p>The first category contains functions for establishing and terminating an authentication activity, which are <code>pam_start(3PAM)</code> and <code>pam_end(3PAM)</code>. The functions <code>pam_set_data(3PAM)</code> and <code>pam_get_data(3PAM)</code> maintain module specific data. The functions <code>pam_set_item(3PAM)</code> and <code>pam_get_item(3PAM)</code> maintain state information. <code>pam_strerror(3PAM)</code> is the function that returns error status information.</p> <p>The second category contains the functions that authenticate an individual user and set the credentials of the user, <code>pam_authenticate(3PAM)</code> and <code>pam_setcred(3PAM)</code>.</p> <p>The third category of PAM interfaces is account management. The function <code>pam_acct_mgmt(3PAM)</code> checks for password aging and access-hour restrictions.</p> <p>Category four contains the functions that perform session management after access to the system has been granted. See <code>pam_open_session(3PAM)</code> and <code>pam_close_session(3PAM)</code>.</p> <p>The fifth category consists of the function that changes authentication tokens, <code>pam_chauthtok(3PAM)</code>. An authentication token is the object used to verify the identity of the user. In UNIX, an authentication token is a user's password.</p> <p>The sixth category of functions can be used to set values for PAM environment variables. See <code>pam_putenv(3PAM)</code>, <code>pam_getenv(3PAM)</code>, and <code>pam_getenvlist(3PAM)</code>.</p> <p>The <code>pam_*()</code> interfaces are implemented through the library <code>libpam</code>. For each of the categories listed above, excluding categories one and six, dynamically loadable shared modules exist that provides the appropriate service layer functionality upon demand. The functional entry points in the service layer</p> |

start with the `pam_sm_` prefix. The only difference between the `pam_sm_*`() interfaces and their corresponding `pam_` interfaces is that all the `pam_sm_*`() interfaces require extra parameters to pass service-specific options to the shared modules. Refer to `pam_sm(3PAM)` for an overview of the PAM service module APIs.

Stateful Interface

A sequence of calls sharing a common set of state information is referred to as an authentication transaction. An authentication transaction begins with a call to `pam_start()`. `pam_start()` allocates space, performs various initialization activities, and assigns a PAM authentication handle to be used for subsequent calls to the library.

After initiating an authentication transaction, applications can invoke `pam_authenticate()` to authenticate a particular user, and `pam_acct_mgmt()` to perform system entry management. For example, the application may want to determine if the user's password has expired.

If the user has been successfully authenticated, the application calls `pam_setcred()` to set any user credentials associated with the authentication service. Within one authentication transaction (between `pam_start()` and `pam_end()`), all calls to the PAM interface should be made with the same authentication handle returned by `pam_start()`. This is necessary because certain service modules may store module-specific data in a handle that is intended for use by other modules. For example, during the call to `pam_authenticate()`, service modules may store data in the handle that is intended for use by `pam_setcred()`.

To perform session management, applications call `pam_open_session()`. Specifically, the system may want to store the total time for the session. The function `pam_close_session()` closes the current session.

When necessary, applications can call `pam_get_item()` and `pam_set_item()` to access and to update specific authentication information. Such information may include the current username.

To terminate an authentication transaction, the application simply calls `pam_end()`, which frees previously allocated space used to store authentication information.

Application-Authentication Service Interactive Interface

The authentication service in PAM does not communicate directly with the user; instead it relies on the application to perform all such interactions. The application passes a pointer to the function, `conv()`, along with any associated application data pointers, through a `pam_conv` structure to the authentication service when it initiates an authentication transaction, via a call to `pam_start()`. The service will then use the function, `conv()`, to prompt the user for data, output error messages, and display text information. Refer to `pam_start(3PAM)` for more information.

Stacking Multiple Schemes

The PAM architecture enables authentication by multiple authentication services through *stacking*. System entry applications, such as `login(1)`, stack multiple service modules to authenticate users with multiple authentication services. The order in which authentication service modules are stacked is specified in the configuration file, `pam.conf(4)`. A system administrator determines this ordering, and also determines whether the same password can be used for all authentication services.

Administrative Interface

The authentication library, `/usr/lib/libpam.so.1`, implements the framework interface. Various authentication services are implemented by their own loadable modules whose paths are specified through the `pam.conf(4)` file.

RETURN VALUES

The PAM functions may return one of the following generic values, or one of the values defined in the specific man pages:

| | |
|------------------------------|---|
| <code>PAM_SUCCESS</code> | The function returned successfully. |
| <code>PAM_OPEN_ERR</code> | <code>dlopen()</code> failed when dynamically loading a service module. |
| <code>PAM_SYMBOL_ERR</code> | Symbol not found. |
| <code>PAM_SERVICE_ERR</code> | Error in service module. |
| <code>PAM_SYSTEM_ERR</code> | System error. |
| <code>PAM_BUF_ERR</code> | Memory buffer error. |
| <code>PAM_CONV_ERR</code> | Conversation failure. |
| <code>PAM_PERM_DENIED</code> | Permission denied. |

ATTRIBUTES

See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-------------------------|
| MT Level | MT-Safe with exceptions |

SEE ALSO

`login(1)`, `pam_authenticate(3PAM)`, `pam_chauthtok(3PAM)`, `pam_open_session(3PAM)`, `pam_set_item(3PAM)`, `pam_setcred(3PAM)`, `pam_sm(3PAM)`, `pam_start(3PAM)`, `pam_strerror(3PAM)`, `pam.conf(4)`, `attributes(5)`

NOTES

The interfaces in `libpam()` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME pam_acct_mgmt – perform PAM account validation procedures

SYNOPSIS `cc [flag ...] file ... -lpam [library ...]
#include <security/pam_appl.h>`

DESCRIPTION `int pam_acct_mgmt(pam_handle_t *pamh, int flags);`
 The `pam_acct_mgmt()` function is called to determine if the current user’s account is valid. It checks for password and account expiration, and verifies access hour restrictions. This function is typically called after the user has been authenticated with `pam_authenticate(3PAM)`.
 The `pamh` argument is an authentication handle obtained by a prior call to `pam_start()`. The following flags may be set in the `flags` field:

| | |
|--|---|
| <code>PAM_SILENT</code> | The account management service should not generate any messages. |
| <code>PAM_DISALLOW_NULL_AUTHTOK</code> | The account management service should return <code>PAM_NEW_AUTHTOK_REQD</code> if the user has a null authentication token. |

RETURN VALUES Upon successful completion, `PAM_SUCCESS` is returned. In addition to the error return values described in `pam(3PAM)`, the following values may be returned:

| | |
|-----------------------------------|---|
| <code>PAM_USER_UNKNOWN</code> | User not known to underlying account management module. |
| <code>PAM_AUTH_ERR</code> | Authentication failure. |
| <code>PAM_NEW_AUTHTOK_REQD</code> | New authentication token required. This is normally returned if the machine security policies require that the password should be changed because the password is NULL or has aged. |
| <code>PAM_ACCT_EXPIRED</code> | User account has expired. |

ATTRIBUTES See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO `pam(3PAM)`, `pam_authenticate(3PAM)`, `pam_start(3PAM)`, `libpam(3LIB)`, `attributes(5)`

NOTES

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME pam_authenticate – perform authentication within the PAM framework

SYNOPSIS
cc [flag ...] file ... -lpam [library ...]
#include <security/pam_appl.h>

int pam_authenticate(pam_handle_t *pamh, int flags);

DESCRIPTION
The pam_authenticate() function is called to authenticate the current user. The user is usually required to enter a password or similar authentication token depending upon the authentication service configured within the system. The user in question should have been specified by a prior call to pam_start() or pam_set_item().

The following flags may be set in the flags field:

PAM_SILENT Authentication service should not generate any messages.

PAM_DISALLOW_NULL_AUTHOK The authentication service should return PAM_AUTH_ERROR if the user has a null authentication token.

RETURN VALUES
Upon successful completion, PAM_SUCCESS is returned. In addition to the error return values described in pam(3PAM), the following values may be returned:

PAM_AUTH_ERR Authentication failure.

PAM_CRED_INSUFFICIENT Cannot access authentication data due to insufficient credentials.

PAM_AUTHINFO_UNAVAIL Underlying authentication service cannot retrieve authentication information.

PAM_USER_UNKNOWN User not known to the underlying authentication module.

PAM_MAXTRIES An authentication service has maintained a retry count which has been reached. No further retries should be attempted.

ATTRIBUTES See attributes(5) for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

pam(3PAM), pam_open_session(3PAM), pam_set_item(3PAM), pam_setcred(3PAM), pam_start(3PAM), libpam(3LIB), attributes(5)

NOTES

In the case of authentication failures due to an incorrect username or password, it is the responsibility of the application to retry `pam_authenticate()` and to maintain the retry count. An authentication service module may implement an internal retry count and return an error `PAM_MAXTRIES` if the module does not want the application to retry.

If the PAM framework cannot load the authentication module, then it will return `PAM_ABORT`. This indicates a serious failure, and the application should not attempt to retry the authentication.

For security reasons, the location of authentication failures is hidden from the user. Thus, if several authentication services are stacked and a single service fails, `pam_authenticate()` requires that the user re-authenticate each of the services.

A null authentication token in the authentication database will result in successful authentication unless `PAM_DISALLOW_NULL_AUTHTOK` was specified. In such cases, there will be no prompt to the user to enter an authentication token.

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

| | | | | | | | | | | | | | | | |
|-----------------------------|--|-----------------|--|-----------------------------|--|---------------------------|---|------------------------|---------------------------------|----------------------------|--------------------------------------|------------------|-----------------------------------|---------------|---|
| NAME | pam_chauthtok – perform password related functions within the PAM framework | | | | | | | | | | | | | | |
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lpam [<i>library ...</i>] #include <security/pam_appl.h> | | | | | | | | | | | | | | |
| DESCRIPTION | <p>int pam_chauthtok(pam_handle_t *pamh, const int flags);</p> <p>The pam_chauthtok() function is called to change the authentication token associated with a particular user referenced by the authentication handle <i>pamh</i>.</p> <p>The following flag may be passed in to pam_chauthtok():</p> <table border="0"> <tr> <td>PAM_SILENT</td> <td>The password service should not generate any messages.</td> </tr> <tr> <td>PAM_CHANGE_EXPIRED_AUTH Tok</td> <td>The password service should only update those passwords that have aged. If this flag is not passed, all password services should update their passwords.</td> </tr> </table> <p>Upon successful completion of the call, the authentication token of the user will be changed in accordance with the password service configured in the system through pam.conf(4).</p> | PAM_SILENT | The password service should not generate any messages. | PAM_CHANGE_EXPIRED_AUTH Tok | The password service should only update those passwords that have aged. If this flag is not passed, all password services should update their passwords. | | | | | | | | | | |
| PAM_SILENT | The password service should not generate any messages. | | | | | | | | | | | | | | |
| PAM_CHANGE_EXPIRED_AUTH Tok | The password service should only update those passwords that have aged. If this flag is not passed, all password services should update their passwords. | | | | | | | | | | | | | | |
| RETURN VALUES | <p>Upon successful completion, PAM_SUCCESS is returned. In addition to the error return values described in pam(3PAM), the following values may be returned:</p> <table border="0"> <tr> <td>PAM_PERM_DENIED</td> <td>No permission.</td> </tr> <tr> <td>PAM_AUTH Tok_ERR</td> <td>Authentication token manipulation error.</td> </tr> <tr> <td>PAM_AUTH Tok_RECOVERY_ERR</td> <td>Authentication information cannot be recovered.</td> </tr> <tr> <td>PAM_AUTH Tok_LOCK_BUSY</td> <td>Authentication token lock busy.</td> </tr> <tr> <td>PAM_AUTH Tok_DISABLE_AGING</td> <td>Authentication token aging disabled.</td> </tr> <tr> <td>PAM_USER_UNKNOWN</td> <td>User unknown to password service.</td> </tr> <tr> <td>PAM_TRY_AGAIN</td> <td>Preliminary check by password service failed.</td> </tr> </table> | PAM_PERM_DENIED | No permission. | PAM_AUTH Tok_ERR | Authentication token manipulation error. | PAM_AUTH Tok_RECOVERY_ERR | Authentication information cannot be recovered. | PAM_AUTH Tok_LOCK_BUSY | Authentication token lock busy. | PAM_AUTH Tok_DISABLE_AGING | Authentication token aging disabled. | PAM_USER_UNKNOWN | User unknown to password service. | PAM_TRY_AGAIN | Preliminary check by password service failed. |
| PAM_PERM_DENIED | No permission. | | | | | | | | | | | | | | |
| PAM_AUTH Tok_ERR | Authentication token manipulation error. | | | | | | | | | | | | | | |
| PAM_AUTH Tok_RECOVERY_ERR | Authentication information cannot be recovered. | | | | | | | | | | | | | | |
| PAM_AUTH Tok_LOCK_BUSY | Authentication token lock busy. | | | | | | | | | | | | | | |
| PAM_AUTH Tok_DISABLE_AGING | Authentication token aging disabled. | | | | | | | | | | | | | | |
| PAM_USER_UNKNOWN | User unknown to password service. | | | | | | | | | | | | | | |
| PAM_TRY_AGAIN | Preliminary check by password service failed. | | | | | | | | | | | | | | |
| ATTRIBUTES | See attributes(5) for description of the following attributes: | | | | | | | | | | | | | | |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

login(1), passwd(1), pam(3PAM), pam_authenticate(3PAM),
pam_start(3PAM), attributes

NOTES

The flag `PAM_CHANGE_EXPIRED_AUTHOK` is typically used by a login application which has determined that the user's password has aged or expired. Before allowing the user to login, the login application may invoke `pam_chauthtok()` with this flag to allow the user to update the password. Typically, applications such as `passwd(1)` should not use this flag.

The `pam_chauthtok()` functions performs a preliminary check before attempting to update passwords. This check is performed for each password module in the stack as listed in `pam.conf(4)`. The check may include pinging remote name services to determine if they are available. If `pam_chauthtok()` returns `PAM_TRY_AGAIN`, then the check has failed, and passwords are not updated.

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME pam_getenv - returns the value for a PAM environment name

SYNOPSIS cc [flag ...] file ... -lpam [library ...]
#include <security/pam_appl.h>

char *pam_getenv(pam_handle_t *pamh, const char *name);

DESCRIPTION The pam_getenv() function searches the PAM handle *pamh* for a value associated with *name*. If a value is present, pam_getenv() makes a copy of the value and returns a pointer to the copy back to the calling application. If no such entry exists, pam_getenv() returns NULL. It is the responsibility of the calling application to free the memory returned by pam_getenv().

RETURN VALUES If successful, pam_getenv() returns a copy of the *value* associated with *name* in the PAM handle; otherwise, it returns a NULL pointer.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO pam(3PAM), pam_getenvlist(3PAM), pam_putenv(3PAM), libpam(3LIB), attributes(5)

NOTES The interfaces in libpam are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME | pam_getenvlist – returns a list of all the PAM environment variables

SYNOPSIS | `cc [flag ...] file ... -lpam [library ...]`
`#include <security/pam_appl.h>`

DESCRIPTION | `char **pam_getenvlist(pam_handle_t *pamh);`
The `pam_getenvlist()` function returns a list of all the PAM environment variables stored in the PAM handle *pamh*. The list is returned as a null-terminated array of pointers to strings. Each string contains a single PAM environment variable of the form *name=value*. The list returned is a duplicate copy of all the environment variables stored in *pamh*. It is the responsibility of the calling application to free the memory returned by `pam_getenvlist()`.

RETURN VALUES | If successful, `pam_getenvlist()` returns in a null-terminated array a copy of all the PAM environment variables stored in *pamh*. Otherwise, `pam_getenvlist()` returns a null pointer.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO | `pam(3PAM)`, `pam_getenv(3PAM)`, `pam_putenv(3PAM)`, `libpam(3LIB)`, `attributes(5)`

NOTES | The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME pam_get_user - PAM routine to retrieve user name

SYNOPSIS cc [flag ...] file ... -lpam [library ...]
#include <security/pam_appl.h>

int pam_get_user(pam_handle_t *pamh, char **user, const char *prompt);

DESCRIPTION The pam_get_user() function is used by PAM service modules to retrieve the current user name from the PAM handle. If the user name has not been set with pam_start() or pam_set_item(), the PAM conversation function will be used to prompt the user for the user name with the string "prompt". If prompt is NULL, then pam_get_item() is called and the value of PAM_USER_PROMPT is used for prompting. If the value of PAM_USER_PROMPT is NULL, the following default prompt is used:

Please enter user name:

After the user name is gathered by the conversation function, pam_set_item() is called to set the value of PAM_USER. By convention, applications that need to prompt for a user name should call pam_set_item() and set the value of PAM_USER_PROMPT before calling pam_authenticate(). The service module's pam_sm_authenticate() function will then call pam_get_user() to prompt for the user name.

Note that certain PAM service modules, such as a smart card module, may override the value of PAM_USER_PROMPT and pass in their own prompt. Applications that call pam_authenticate() multiple times should set the value of PAM_USER to NULL with pam_set_item() before calling pam_authenticate(), if they want the user to be prompted for a new user name each time. The value of user retrieved by pam_get_user() should not be modified or freed. The item will be released by pam_end().

RETURN VALUES Upon success, pam_get_user() returns PAM_SUCCESS; otherwise it returns an error code. Refer to pam(3PAM) for information on error related return values.

ATTRIBUTES See attributes(5) for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO pam(3PAM), pam_authenticate(3PAM), pam_end(3PAM), pam_get_item(3PAM), pam_set_item(3PAM), pam_sm(3PAM), pam_sm_authenticate(3PAM), pam_start(3PAM), attributes(5)

NOTES

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME pam_open_session, pam_close_session – perform PAM session creation and termination operations

SYNOPSIS

```
cc [ flag ... ] file ... -lpam [ library ... ]
#include <security/pam_appl.h>
int pam_open_session(pam_handle_t *pamh, int flags);
int pam_close_session(pam_handle_t *pamh, int flags);
```

DESCRIPTION
 The pam_open_session() function is called after a user has been successfully authenticated. See pam_authenticate(3PAM) and pam_acct_mgmt(3PAM). It is used to notify the session modules that a new session has been initiated. All programs that use the pam(3PAM) library should invoke pam_open_session() when beginning a new session. Upon termination of this activity, pam_close_session() should be invoked to inform pam(3PAM) that the session has terminated.

The pamh argument is an authentication handle obtained by a prior call to pam_start(). The following flag may be set in the flags field for pam_open_session() and pam_close_session():
 PAM_SILENT The session service should not generate any messages.

RETURN VALUES
 Upon successful completion, PAM_SUCCESS is returned. In addition to the return values defined in pam(3PAM), the following value may be returned on error:
 PAM_SESSION_ERR Cannot make or remove an entry for the specified session.

ATTRIBUTES See attributes(5) for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO getutxent(3C), pam(3PAM), pam_acct_mgmt(3PAM), pam_authenticate(3PAM), pam_start(3PAM), attributes(5)

NOTES
 In many instances, the pam_open_session() and pam_close_session() calls may be made by different processes. For example, in UNIX the login process opens a session, while the init process closes the session. In this case, UTMP/WTMP entries may be used to link the call to pam_close_session() with an earlier call to pam_open_session(). This is possible because UTMP/WTMP entries are uniquely identified by a combination of attributes, including the user login name and device name, which are accessible through the PAM handle, pamh. The call to pam_open_session() should precede

UTMP/WTMP entry management, and the call to `pam_close_session()` should follow UTMP/WTMP exit management.

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

| | | | | | | | | | | | | | | | | | |
|----------------------|---|-------------|-------------------------------------|--------------|---|----------------|-------------------|-----------------|--------------------------|----------------|---------------|-------------|----------------------|--------------|-----------------------|-----------------|--------------------|
| NAME | pam_putenv – change or add a value to the PAM environment | | | | | | | | | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lpam [library ...] #include <security/pam_appl.h></pre> | | | | | | | | | | | | | | | | |
| DESCRIPTION | <p>int pam_putenv(pam_handle_t *pamh, const char *name_value);</p> <p>The <code>pam_putenv()</code> function sets the value of the PAM environment variable <i>name</i> equal to <i>value</i> either by altering an existing PAM variable or by creating a new one.</p> <p>The <i>name_value</i> argument points to a string of the form <i>name=value</i>. A call to <code>pam_putenv()</code> does not immediately change the environment. All <i>name_value</i> pairs are stored in the PAM handle <i>pamh</i>. An application such as <code>login(1)</code> may make a call to <code>pam_getenv(3PAM)</code> or <code>pam_getenvlist(3PAM)</code> to retrieve the PAM environment variables saved in the PAM handle and set them in the environment if appropriate. <code>login</code> will not set PAM environment values which overwrite the values for SHELL, HOME, LOGNAME, MAIL, CDPATH, IFS, and PATH. Nor will <code>login</code> set PAM environment values which overwrite any value that begins with LD_.</p> <p>If <i>name_value</i> equals <i>NAME=</i>, then the value associated with <i>NAME</i> in the PAM handle will be set to an empty value. If <i>name_value</i> equals <i>NAME</i>, then the environment variable <i>NAME</i> will be removed from the PAM handle.</p> | | | | | | | | | | | | | | | | |
| RETURN VALUES | <p>The <code>pam_putenv()</code> function may return one of the following values:</p> <table border="0"> <tr> <td>PAM_SUCCESS</td> <td>The function returned successfully.</td> </tr> <tr> <td>PAM_OPEN_ERR</td> <td><code>dlopen()</code> failed when dynamically loading a service module.</td> </tr> <tr> <td>PAM_SYMBOL_ERR</td> <td>Symbol not found.</td> </tr> <tr> <td>PAM_SERVICE_ERR</td> <td>Error in service module.</td> </tr> <tr> <td>PAM_SYSTEM_ERR</td> <td>System error.</td> </tr> <tr> <td>PAM_BUF_ERR</td> <td>Memory buffer error.</td> </tr> <tr> <td>PAM_CONV_ERR</td> <td>Conversation failure.</td> </tr> <tr> <td>PAM_PERM_DENIED</td> <td>Permission denied.</td> </tr> </table> | PAM_SUCCESS | The function returned successfully. | PAM_OPEN_ERR | <code>dlopen()</code> failed when dynamically loading a service module. | PAM_SYMBOL_ERR | Symbol not found. | PAM_SERVICE_ERR | Error in service module. | PAM_SYSTEM_ERR | System error. | PAM_BUF_ERR | Memory buffer error. | PAM_CONV_ERR | Conversation failure. | PAM_PERM_DENIED | Permission denied. |
| PAM_SUCCESS | The function returned successfully. | | | | | | | | | | | | | | | | |
| PAM_OPEN_ERR | <code>dlopen()</code> failed when dynamically loading a service module. | | | | | | | | | | | | | | | | |
| PAM_SYMBOL_ERR | Symbol not found. | | | | | | | | | | | | | | | | |
| PAM_SERVICE_ERR | Error in service module. | | | | | | | | | | | | | | | | |
| PAM_SYSTEM_ERR | System error. | | | | | | | | | | | | | | | | |
| PAM_BUF_ERR | Memory buffer error. | | | | | | | | | | | | | | | | |
| PAM_CONV_ERR | Conversation failure. | | | | | | | | | | | | | | | | |
| PAM_PERM_DENIED | Permission denied. | | | | | | | | | | | | | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | | | | | | | | | | | | | |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

dlopen(3DL), pam(3PAM), pam_getenv(3PAM), pam_getenvlist(3PAM), libpam(3LIB), attributes(5)

NOTES

The interfaces in libpam are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

| | | | | | | | | | | | |
|-----------------------|---|--------------------|--|------------------|--|-----------------------|--|------------------|--------------------------------------|------------|--|
| NAME | pam_setcred – modify/delete user credentials for an authentication service | | | | | | | | | | |
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lpam [<i>library ...</i>] #include <security/pam_appl.h> | | | | | | | | | | |
| DESCRIPTION | <p>int pam_setcred(pam_handle_t *<i>pamh</i>, int <i>flags</i>);</p> <p>The <code>pam_setcred()</code> function is used to establish, modify, or delete user credentials. It is typically called after the user has been authenticated and after a session has been opened. See <code>pam_authenticate(3PAM)</code>, <code>pam_acct_mgmt(3PAM)</code>, and <code>pam_open_session(3PAM)</code>.</p> <p>The user is specified by a prior call to <code>pam_start()</code> or <code>pam_set_item()</code>, and is referenced by the authentication handle, <i>pamh</i>. The following flags may be set in the <i>flags</i> field. Note that the first four flags are mutually exclusive:</p> <table border="0"> <tr> <td>PAM_ESTABLISH_CRED</td> <td>Set user credentials for an authentication service.</td> </tr> <tr> <td>PAM_DELETE_CRED</td> <td>Delete user credentials associated with an authentication service.</td> </tr> <tr> <td>PAM_REINITIALIZE_CRED</td> <td>Reinitialize user credentials.</td> </tr> <tr> <td>PAM_REFRESH_CRED</td> <td>Extend lifetime of user credentials.</td> </tr> <tr> <td>PAM_SILENT</td> <td>Authentication service should not generate any messages.</td> </tr> </table> <p>If no flag is set, PAM_ESTABLISH_CRED is used as the default.</p> | PAM_ESTABLISH_CRED | Set user credentials for an authentication service. | PAM_DELETE_CRED | Delete user credentials associated with an authentication service. | PAM_REINITIALIZE_CRED | Reinitialize user credentials. | PAM_REFRESH_CRED | Extend lifetime of user credentials. | PAM_SILENT | Authentication service should not generate any messages. |
| PAM_ESTABLISH_CRED | Set user credentials for an authentication service. | | | | | | | | | | |
| PAM_DELETE_CRED | Delete user credentials associated with an authentication service. | | | | | | | | | | |
| PAM_REINITIALIZE_CRED | Reinitialize user credentials. | | | | | | | | | | |
| PAM_REFRESH_CRED | Extend lifetime of user credentials. | | | | | | | | | | |
| PAM_SILENT | Authentication service should not generate any messages. | | | | | | | | | | |
| RETURN VALUES | <p>Upon success, <code>pam_setcred()</code> returns PAM_SUCCESS. In addition to the error return values described in <code>pam(3PAM)</code> the following values may be returned upon error:</p> <table border="0"> <tr> <td>PAM_CRED_UNAVAIL</td> <td>Underlying authentication service can not retrieve user credentials unavailable.</td> </tr> <tr> <td>PAM_CRED_EXPIRED</td> <td>User credentials expired.</td> </tr> <tr> <td>PAM_USER_UNKNOWN</td> <td>User unknown to underlying authentication service.</td> </tr> <tr> <td>PAM_CRED_ERR</td> <td>Failure setting user credentials.</td> </tr> </table> | PAM_CRED_UNAVAIL | Underlying authentication service can not retrieve user credentials unavailable. | PAM_CRED_EXPIRED | User credentials expired. | PAM_USER_UNKNOWN | User unknown to underlying authentication service. | PAM_CRED_ERR | Failure setting user credentials. | | |
| PAM_CRED_UNAVAIL | Underlying authentication service can not retrieve user credentials unavailable. | | | | | | | | | | |
| PAM_CRED_EXPIRED | User credentials expired. | | | | | | | | | | |
| PAM_USER_UNKNOWN | User unknown to underlying authentication service. | | | | | | | | | | |
| PAM_CRED_ERR | Failure setting user credentials. | | | | | | | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for description of the following attributes: | | | | | | | | | | |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

pam(3PAM), pam_acct_mgmt(3PAM), pam_authenticate(3PAM),
pam_open_session(3PAM), pam_set_item(3PAM), pam_start(3PAM),
libpam(3LIB), attributes(5)

NOTES

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME pam_set_data, pam_get_data – PAM routines to maintain module specific state

SYNOPSIS

```
cc [ flag ... ] file ... -lpam [ library ... ]
#include <security/pam_appl.h>
int pam_set_data(pam_handle_t *pamh, const char *module_data_name, void *data, void (*cleanup) (pam_handle_t *pamh, void *data, int pam_end_status));

int pam_get_data(const pam_handle_t *pamh, const char *module_data_name, const void **data);
```

DESCRIPTION

The `pam_set_data()` and `pam_get_data()` functions allow PAM service modules to access and update module specific information as needed. These functions should not be used by applications.

The `pam_set_data()` function stores module specific data within the PAM handle `pamh`. The `module_data_name` argument uniquely identifies the data, and the `data` argument represents the actual data. The `module_data_name` argument should be unique across all services.

The `cleanup` function frees up any memory used by the `data` after it is no longer needed, and is invoked by `pam_end()`. The `cleanup` function takes as its arguments a pointer to the PAM handle, `pamh`, a pointer to the actual data, `data`, and a status code, `pam_end_status`. The status code determines exactly what state information needs to be purged.

If `pam_set_data()` is called and module data already exists from a prior call to `pam_set_data()` under the same `module_data_name`, then the existing `data` is replaced by the new `data`, and the existing `cleanup` function is replaced by the new `cleanup` function.

The `pam_get_data()` function retrieves module-specific data stored in the PAM handle, `pamh`, identified by the unique name, `module_data_name`. The `data` argument is assigned the address of the requested data. The `data` retrieved by `pam_get_data()` should not be modified or freed. The `data` will be released by `pam_end()`.

RETURN VALUES

In addition to the return values listed in `pam(3PAM)`, the following value may also be returned:

PAM_NO_MODULE_DATA No module specific data is present.

ATTRIBUTES See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO `pam(3PAM)`, `pam_end(3PAM)`, `libpam(3LIB)`, `attributes(5)`

NOTES

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

| | | | | | | | | | | | | | | | | | | | |
|--------------------|---|-------------|-------------------|----------|----------------|-------------|--------------------------------|----------------|------------------------------------|---------|---------------|-----------|-----------------------|-----------|-----------------------|----------|--------------------------------------|-----------------|--|
| NAME | pam_set_item, pam_get_item – authentication information routines for PAM | | | | | | | | | | | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lpam [library ...] #include <security/pam_appl.h> int pam_set_item(pam_handle_t *pamh, int item_type, const void *item); int pam_get_item(const pam_handle_t *pamh, int item_type, void **item);</pre> | | | | | | | | | | | | | | | | | | |
| DESCRIPTION | <p>The <code>pam_get_item()</code> and <code>pam_set_item()</code> functions allow applications and PAM service modules to access and to update PAM information as needed. The information is specified by <i>item_type</i>, and can be one of the following:</p> <table border="0"> <tr> <td>PAM_SERVICE</td> <td>The service name.</td> </tr> <tr> <td>PAM_USER</td> <td>The user name.</td> </tr> <tr> <td>PAM_AUTHTOK</td> <td>The user authentication token.</td> </tr> <tr> <td>PAM_OLDAUTHTOK</td> <td>The old user authentication token.</td> </tr> <tr> <td>PAM_TTY</td> <td>The tty name.</td> </tr> <tr> <td>PAM_RHOST</td> <td>The remote host name.</td> </tr> <tr> <td>PAM_RUSER</td> <td>The remote user name.</td> </tr> <tr> <td>PAM_CONV</td> <td>The <code>pam_conv</code> structure.</td> </tr> <tr> <td>PAM_USER_PROMPT</td> <td>The default prompt used by <code>pam_get_user()</code>.</td> </tr> </table> <p>For security reasons, the <i>item_type</i> PAM_AUTHTOK and PAM_OLDAUTHTOK are available only to the module providers. The authentication module, account module, and session management module should treat PAM_AUTHTOK as the current authentication token and ignore PAM_OLDAUTHTOK. The password management module should treat PAM_OLDAUTHTOK as the current authentication token and PAM_AUTHTOK as the new authentication token.</p> <p>The <code>pam_set_item()</code> function is passed the authentication handle, <i>pamh</i>, returned by <code>pam_start()</code>, a pointer to the object, <i>item</i>, and its type, <i>item_type</i>. If successful, <code>pam_set_item()</code> copies the item to an internal storage area allocated by the authentication module and returns PAM_SUCCESS. An item that had been previously set will be overwritten by the new value.</p> <p>The <code>pam_get_item()</code> function is passed the authentication handle, <i>pamh</i>, returned by <code>pam_start()</code>, an <i>item_type</i>, and the address of the pointer, <i>item</i>, which is assigned the address of the requested object. The object data is valid until modified by a subsequent call to <code>pam_set_item()</code> for the same <i>item_type</i>, or unless it is modified by any of the underlying service modules. If the item has not been previously set, <code>pam_get_item()</code> returns a null pointer. An <i>item</i> retrieved by <code>pam_get_item()</code> should not be modified or freed. The item will be released by <code>pam_end()</code>.</p> | PAM_SERVICE | The service name. | PAM_USER | The user name. | PAM_AUTHTOK | The user authentication token. | PAM_OLDAUTHTOK | The old user authentication token. | PAM_TTY | The tty name. | PAM_RHOST | The remote host name. | PAM_RUSER | The remote user name. | PAM_CONV | The <code>pam_conv</code> structure. | PAM_USER_PROMPT | The default prompt used by <code>pam_get_user()</code> . |
| PAM_SERVICE | The service name. | | | | | | | | | | | | | | | | | | |
| PAM_USER | The user name. | | | | | | | | | | | | | | | | | | |
| PAM_AUTHTOK | The user authentication token. | | | | | | | | | | | | | | | | | | |
| PAM_OLDAUTHTOK | The old user authentication token. | | | | | | | | | | | | | | | | | | |
| PAM_TTY | The tty name. | | | | | | | | | | | | | | | | | | |
| PAM_RHOST | The remote host name. | | | | | | | | | | | | | | | | | | |
| PAM_RUSER | The remote user name. | | | | | | | | | | | | | | | | | | |
| PAM_CONV | The <code>pam_conv</code> structure. | | | | | | | | | | | | | | | | | | |
| PAM_USER_PROMPT | The default prompt used by <code>pam_get_user()</code> . | | | | | | | | | | | | | | | | | | |

RETURN VALUES

Upon success, `pam_get_item()` returns `PAM_SUCCESS`; otherwise it returns an error code. Refer to `pam(3PAM)` for information on error related return values.

ATTRIBUTES

See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

`pam(3PAM)`, `pam_acct_mgmt(3PAM)`, `pam_authenticate(3PAM)`, `pam_chauthtok(3PAM)`, `pam_get_user(3PAM)`, `pam_open_session(3PAM)`, `pam_setcred(3PAM)`, `pam_start(3PAM)`, `attributes(5)`

NOTES

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

| | |
|---------------------------|---|
| NAME | pam_sm – PAM Service Module APIs |
| SYNOPSIS | <pre>#include <security/pam_appl.h> #include <security/pam_modules.h> cc [flag ...] file ... -lpam [library ...]</pre> |
| DESCRIPTION | <p>PAM gives system administrators the flexibility of choosing any authentication service available on the system to perform authentication. The framework also allows new authentication service modules to be plugged in and made available without modifying the applications.</p> <p>The PAM framework, <code>libpam</code>, consists of an interface library and multiple authentication service modules. The PAM interface library is the layer implementing the Application Programming Interface (API). The authentication service modules are a set of dynamically loadable objects invoked by the PAM API to provide a particular type of user authentication.</p> |
| Interface Overview | <p>This manual page gives an overview of the PAM APIs for the service modules. The PAM service module interface consists of functions which can be grouped into four categories. The names for all the authentication library functions start with <code>pam_sm</code>. The only difference between the <code>pam_*()</code> interfaces and their corresponding <code>pam_sm_*()</code> interfaces is that all the <code>pam_sm_*()</code> interfaces require extra parameters to pass service-specific options to the shared modules. They are otherwise identical.</p> <p>The first category contains functions to authenticate an individual user, <code>pam_sm_authenticate(3PAM)</code>, and to set the credentials of the user, <code>pam_sm_setcred(3PAM)</code>. These back-end functions implement the functionality of <code>pam_authenticate(3PAM)</code> and <code>pam_setcred(3PAM)</code> respectively.</p> <p>The second category contains the function to do account management: <code>pam_sm_acct_mgmt(3PAM)</code>. This includes checking for password aging and access-hour restrictions. This back-end function implements the functionality of <code>pam_acct_mgmt(3PAM)</code>.</p> <p>The third category contains the functions <code>pam_sm_open_session(3PAM)</code> and <code>pam_sm_close_session(3PAM)</code> to perform session management after access to the system has been granted. These back-end functions implement the functionality of <code>pam_open_session(3PAM)</code> and <code>pam_close_session(3PAM)</code>, respectively.</p> <p>The fourth category consists a function to change authentication tokens <code>pam_sm_chauthtok(3PAM)</code>. This back-end function implements the functionality of <code>pam_chauthtok(3PAM)</code>.</p> |

Stateful Interface

A sequence of calls sharing a common set of state information is referred to as an authentication transaction. An authentication transaction begins with a call to `pam_start()`. `pam_start()` allocates space, performs various initialization activities, and assigns an authentication handle to be used for subsequent calls to the library. Note that the service modules do not get called or initialized when `pam_start()` is called. The modules are loaded and the symbols resolved upon first use of that function.

The PAM handle keeps certain information about the transaction that can be accessed through the `pam_get_item()` API. Though the modules can also use `pam_set_item()` to change any of the item information, it is recommended that nothing be changed except `PAM_AUTHTOK` and `PAM_OLDAUTHTOK`.

If the modules want to store any module specific state information then they can use the `pam_set_data(3PAM)` function to store that information with the PAM handle. The data should be stored with a name which is unique across all modules and module types. For example, `SUNW_PAM_UNIX_AUTH_userid` can be used as a name by the UNIX module to store information about the state of user's authentication. Some modules use this technique to share data across two different module types.

Also, during the call to `pam_authenticate()`, the UNIX module may store the authentication status (success or reason for failure) in the handle, using a unique name such as `SUNW_SECURE_RPC_DATA`. This information is intended for use by `pam_setcred()`.

During the call to `pam_acct_mgmt()`, the account modules may store data in the handle to indicate which passwords have aged. This information is intended for use by `pam_chauthtok()`.

The module can also store a cleanup function associated with the data. The PAM framework calls this cleanup function, when the application calls `pam_end()` to close the transaction.

Interaction with the User

The PAM service modules do not communicate directly with the user; instead they rely on the application to perform all such interactions. The application passes a pointer to the function, `conv()`, along with any associated application data pointers, through the `pam_conv` structure when it initiates an authentication transaction (via a call to `pam_start()`). The service module will then use the function, `conv()`, to prompt the user for data, output error messages, and display text information. Refer to `pam_start(3PAM)` for more information. The modules are responsible for the localization of all messages to the user.

CONVENTIONS

By convention, applications that need to prompt for a user name should call `pam_set_item()` and set the value of `PAM_USER_PROMPT` before calling `pam_authenticate()`. The service module's `pam_sm_authenticate()` function will then call `pam_get_user()` to prompt for the user name. Note

that certain PAM service modules (such as a smart card module) may override the value of PAM_USER_PROMPT and pass in their own prompt.

Though the PAM framework enforces no rules about the module's names, location, options and such, there are certain conventions that all module providers are expected to follow.

By convention, the modules should be located in the `/usr/lib/security` directory. Additional modules may be located in `/opt/<pkg>/lib`.

By convention, the modules are named `pam_<service_name>_<module_type>.so.1`. If the given module implements more than one module type (for example, `pam_unix.so.1` module), then the `module_type` suffix should be dropped.

For every such module, there should be a corresponding manual page in section 5 which should describe the *module_type* it supports, the functionality of the module, along with the options it supports. The dependencies should be clearly identified to the system administrator. For example, it should be made clear whether this module is a stand-alone module or depends upon the presence of some other module. One should also specify whether this module should come before or after some other module in the stack.

By convention, the modules should support the following options:

| | |
|---------------------|--|
| <code>debug</code> | Syslog debugging information at LOG_DEBUG level. Be careful as to not log any sensitive information such as passwords. |
| <code>nowarn</code> | Turn off warning messages such as "password is about to expire." |

In addition, it is recommended that the `auth` and the `password` module support the following options:

| | |
|-----------------------------|--|
| <code>use_first_pass</code> | Instead of prompting the user for the password, use the user's initial password (entered when the user was authenticated to the first authentication module in the stack) for authentication. If the passwords do not match, or if no password has been entered, return failure and do not prompt the user for a password. Support for this scheme allows the user to type only one password for multiple schemes. |
| <code>try_first_pass</code> | Instead of prompting the user for the password, use the user's initial password (entered when the user was authenticated to the first authentication module in the stack) for authentication. If the |

passwords do not match, or if no password has been entered, prompt the user for a password after identifying which type of password (ie. UNIX, etc.) is being requested. Support for this scheme allows the user to try to use only one password for multiple schemes, and type multiple passwords only if necessary.

If an unsupported option is passed to the modules, it should syslog the error at LOG_ERR level.

The permission bits on the service module should be set such that it is not writable by either "group" or "other." The PAM framework will not load the module if the above permission rules are not followed.

ERROR LOGGING

If there are any errors, the modules should log them using `syslog(3C)` at the LOG_ERR level.

RETURN VALUES

The PAM service module functions may return any of the PAM error numbers specified in the specific man pages. It can also return a PAM_IGNORE error number to mean that the PAM framework should ignore this module regardless of whether it is required, optional or sufficient. This error number is normally returned when the module does not want to deal with the given user at all.

ATTRIBUTES

See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

`pam(3PAM)`, `pam_authenticate(3PAM)`, `pam_chauthtok(3PAM)`, `pam_get_user(3PAM)`, `pam_open_session(3PAM)`, `pam_setcred(3PAM)`, `pam_set_item(3PAM)`, `pam_sm_authenticate(3PAM)`, `pam_sm_chauthtok(3PAM)`, `pam_sm_open_session(3PAM)`, `pam_sm_setcred(3PAM)`, `pam_start(3PAM)`, `pam_strerror(3PAM)`, `syslog(3C)`, `pam.conf(4)`, `attributes(5)`

NOTES

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

| | | | | | | | | | | | |
|--------------------------|---|------------|--|--------------------------|---|------------------|---|---------------------|------------------------------------|------------------|---------------------------|
| NAME | pam_sm_acct_mgmt – service provider implementation for pam_acct_mgmt | | | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lpam [library ...] #include <security/pam_appl.h> #include <security/pam_modules.h></pre> | | | | | | | | | | |
| DESCRIPTION | <p>int pam_sm_acct_mgmt(pam_handle_t *pamh, int flags, int argc, const char **argv);</p> <p>In response to a call to <code>pam_acct_mgmt(3PAM)</code>, the PAM framework calls <code>pam_sm_acct_mgmt()</code> from the modules listed in the <code>pam.conf(4)</code> file. The account management provider supplies the back-end functionality for this interface function. Applications should not call this API directly.</p> <p>The <code>pam_sm_acct_mgmt()</code> function determines whether or not the current user's account and password are valid. This includes checking for password and account expiration, and valid login times. The user in question is specified by a prior call to <code>pam_start()</code>, and is referenced by the authentication handle, <code>pamh</code>, which is passed as the first argument to <code>pam_sm_acct_mgmt()</code>. The following flags may be set in the <code>flags</code> field:</p> <table border="0"> <tr> <td style="padding-right: 20px;">PAM_SILENT</td> <td>The account management service should not generate any messages.</td> </tr> <tr> <td>PAM_DISALLOW_NULL_AUTHOK</td> <td>The account management service should return PAM_NEW_AUTHOK_REQD if the user has a null authentication token.</td> </tr> </table> <p>The <code>argc</code> argument represents the number of module options passed in from the configuration file <code>pam.conf(4)</code>. <code>argv</code> specifies the module options, which are interpreted and processed by the account management service. Please refer to the specific module man pages for the various available <i>options</i>. If an unknown option is passed to the module, an error should be logged through <code>syslog(3C)</code> and the option ignored.</p> <p>If an account management module determines that the user password has aged or expired, it should save this information as state in the authentication handle, <code>pamh</code>, using <code>pam_set_data()</code>. <code>pam_chauthok()</code> uses this information to determine which passwords have expired.</p> <p>If there are no restrictions to logging in, <code>PAM_SUCCESS</code> is returned. The following error values may also be returned upon error:</p> <table border="0"> <tr> <td style="padding-right: 20px;">PAM_USER_UNKNOWN</td> <td>User not known to underlying authentication module.</td> </tr> <tr> <td>PAM_NEW_AUTHOK_REQD</td> <td>New authentication token required.</td> </tr> <tr> <td>PAM_ACCT_EXPIRED</td> <td>User account has expired.</td> </tr> </table> | PAM_SILENT | The account management service should not generate any messages. | PAM_DISALLOW_NULL_AUTHOK | The account management service should return PAM_NEW_AUTHOK_REQD if the user has a null authentication token. | PAM_USER_UNKNOWN | User not known to underlying authentication module. | PAM_NEW_AUTHOK_REQD | New authentication token required. | PAM_ACCT_EXPIRED | User account has expired. |
| PAM_SILENT | The account management service should not generate any messages. | | | | | | | | | | |
| PAM_DISALLOW_NULL_AUTHOK | The account management service should return PAM_NEW_AUTHOK_REQD if the user has a null authentication token. | | | | | | | | | | |
| PAM_USER_UNKNOWN | User not known to underlying authentication module. | | | | | | | | | | |
| PAM_NEW_AUTHOK_REQD | New authentication token required. | | | | | | | | | | |
| PAM_ACCT_EXPIRED | User account has expired. | | | | | | | | | | |
| RETURN VALUES | | | | | | | | | | | |

PAM_PERM_DENIED User denied access to account at this time.

PAM_IGNORE Ignore underlying account module regardless of whether the control flag is *required*, *optional* or *sufficient*.

ATTRIBUTES

See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

`pam(3PAM)`, `pam_acct_mgmt(3PAM)`, `pam_set_data(3PAM)`, `pam_start(3PAM)`, `syslog(3C)`, `libpam(3LIB)`, `pam.conf(4)`, `attributes(5)`

NOTES

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

| | | | | | | | |
|-----------------------------------|---|-------------------|--|-----------------------------------|--|-----------------------|---|
| NAME | pam_sm_authenticate – service provider implementation for pam_authenticate | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lpam [library ...] #include <security/pam_appl.h> #include <security/pam_modules.h></pre> | | | | | | |
| DESCRIPTION | <p>int pam_sm_authenticate(pam_handle_t *pamh, int flags, int argc, const char **argv);</p> <p>In response to a call to pam_authenticate(3PAM), the PAM framework calls pam_sm_authenticate() from the modules listed in the pam.conf(4) file. The authentication provider supplies the back-end functionality for this interface function.</p> <p>The pam_sm_authenticate() function is called to verify the identity of the current user. The user is usually required to enter a password or similar authentication token depending upon the authentication scheme configured within the system. The user in question is specified by a prior call to pam_start(), and is referenced by the authentication handle <i>pamh</i>.</p> <p>If the user is unknown to the authentication service, the service module should mask this error and continue to prompt the user for a password. It should then return the error, PAM_USER_UNKNOWN.</p> <p>The following flag may be passed in to pam_sm_authenticate():</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">PAM_SILENT</td> <td>The authentication service should not generate any messages.</td> </tr> <tr> <td>PAM_DISALLOW_NULL_AUTH Tok</td> <td>The authentication service should return</td> </tr> <tr> <td>PAM_AUTH_ERROR</td> <td>The user has a null authentication token.</td> </tr> </table> <p>The <i>argc</i> argument represents the number of module options passed in from the configuration file pam.conf(4). <i>argv</i> specifies the module options, which are interpreted and processed by the authentication service. Please refer to the specific module man pages for the various available <i>options</i>. If any unknown option is passed in, the module should log the error and ignore the option.</p> <p>Before returning, pam_sm_authenticate() should call pam_get_item() and retrieve PAM_AUTHTOK. If it has not been set before and the value is NULL, pam_sm_authenticate() should set it to the password entered by the user using pam_set_item().</p> <p>An authentication module may save the authentication status (success or reason for failure) as state in the authentication handle using pam_set_data(3PAM). This information is intended for use by pam_setcred().</p> | PAM_SILENT | The authentication service should not generate any messages. | PAM_DISALLOW_NULL_AUTH Tok | The authentication service should return | PAM_AUTH_ERROR | The user has a null authentication token. |
| PAM_SILENT | The authentication service should not generate any messages. | | | | | | |
| PAM_DISALLOW_NULL_AUTH Tok | The authentication service should return | | | | | | |
| PAM_AUTH_ERROR | The user has a null authentication token. | | | | | | |

RETURN VALUES

Upon successful completion, `PAM_SUCCESS` must be returned. In addition, the following values may be returned:

| | |
|------------------------------------|---|
| <code>PAM_MAXTRIES</code> | Maximum number of authentication attempts exceeded. |
| <code>PAM_AUTH_ERR</code> | Authentication failure. |
| <code>PAM_CRED_INSUFFICIENT</code> | Cannot access authentication data due to insufficient credentials. |
| <code>PAM_AUTHINFO_UNAVAIL</code> | Underlying authentication service can not retrieve authentication information. |
| <code>PAM_USER_UNKNOWN</code> | User not known to underlying authentication module. |
| <code>PAM_IGNORE</code> | Ignore underlying authentication module regardless of whether the control flag is <i>required</i> , <i>optional</i> , or <i>sufficient</i> ¹ . |

ATTRIBUTES

See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

`pam(3PAM)`, `pam_authenticate(3PAM)`, `pam_get_item(3PAM)`, `pam_set_data(3PAM)`, `pam_set_item(3PAM)`, `pam_setcred(3PAM)`, `pam_start(3PAM)`, `libpam(3LIB)`, `pam.conf(4)`, `attributes(5)`

NOTES

Modules should not retry the authentication in the event of a failure. Applications handle authentication retries and maintain the retry count. To limit the number of retries, the module can return a `PAM_MAXTRIES` error.

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

| | | | | | | | | | |
|--|---|-------------------------|--|--|---|-------------------------------|--|----------------------------------|---|
| NAME | pam_sm_chauthtok – service provider implementation for pam_chauthtok | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lpam [library ...] #include <security/pam_appl.h> #include <security/pam_modules.h> int pam_sm_chauthtok(pam_handle_t *pamh, int flags, int argc, const char **argv);</pre> | | | | | | | | |
| DESCRIPTION | <p>In response to a call to <code>pam_chauthtok()</code> the PAM framework calls <code>pam_sm_chauthtok(3PAM)</code> from the modules listed in the <code>pam.conf(4)</code> file. The password management provider supplies the back-end functionality for this interface function.</p> <p>The <code>pam_sm_chauthtok()</code> function changes the authentication token associated with a particular user referenced by the authentication handle <code>pamh</code>.</p> <p>The following flag may be passed to <code>pam_chauthtok()</code>:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>PAM_SILENT</code></td> <td>The password service should not generate any messages.</td> </tr> <tr> <td style="padding-right: 20px;"><code>PAM_CHANGE_EXPIRED_AUTH Tok</code></td> <td>The password service should only update those passwords that have aged. If this flag is not passed, the password service should update all passwords.</td> </tr> <tr> <td style="padding-right: 20px;"><code>PAM_PRELIM_CHECK</code></td> <td>The password service should only perform preliminary checks. No passwords should be updated.</td> </tr> <tr> <td style="padding-right: 20px;"><code>PAM_UPDATE_AUTH Tok</code></td> <td>The password service should update passwords.</td> </tr> </table> <p>Note that <code>PAM_PRELIM_CHECK</code> and <code>PAM_UPDATE_AUTH Tok</code> cannot be set at the same time.</p> <p>Upon successful completion of the call, the authentication token of the user will be ready for change or will be changed, depending upon the flag, in accordance with the authentication scheme configured within the system.</p> <p>The <code>argc</code> argument represents the number of module options passed in from the configuration file <code>pam.conf(4)</code>. The <code>argv</code> argument specifies the module options, which are interpreted and processed by the password management service. Please refer to the specific module man pages for the various available <i>options</i>.</p> <p>It is the responsibility of <code>pam_sm_chauthtok()</code> to determine if the new password meets certain strength requirements. <code>pam_sm_chauthtok()</code> may continue to re-prompt the user (for a limited number of times) for a new password until the password entered meets the strength requirements.</p> | <code>PAM_SILENT</code> | The password service should not generate any messages. | <code>PAM_CHANGE_EXPIRED_AUTH Tok</code> | The password service should only update those passwords that have aged. If this flag is not passed, the password service should update all passwords. | <code>PAM_PRELIM_CHECK</code> | The password service should only perform preliminary checks. No passwords should be updated. | <code>PAM_UPDATE_AUTH Tok</code> | The password service should update passwords. |
| <code>PAM_SILENT</code> | The password service should not generate any messages. | | | | | | | | |
| <code>PAM_CHANGE_EXPIRED_AUTH Tok</code> | The password service should only update those passwords that have aged. If this flag is not passed, the password service should update all passwords. | | | | | | | | |
| <code>PAM_PRELIM_CHECK</code> | The password service should only perform preliminary checks. No passwords should be updated. | | | | | | | | |
| <code>PAM_UPDATE_AUTH Tok</code> | The password service should update passwords. | | | | | | | | |

Before returning, `pam_sm_chauthtok()` should call `pam_get_item()` and retrieve both `PAM_AUTHTOK` and `PAM_OLDAUTHOK`. If both are `NULL`, `pam_sm_chauthtok()` should set them to the new and old passwords as entered by the user.

RETURN VALUES

Upon successful completion, `PAM_SUCCESS` must be returned. The following values may also be returned:

| | |
|--|---|
| <code>PAM_PERM_DENIED</code> | No permission. |
| <code>PAM_AUTHTOK_ERR</code> | Authentication token manipulation error. |
| <code>PAM_AUTHTOK_RECOVERY_ERR</code> | Old authentication token cannot be recovered. |
| <code>PAM_AUTHTOK_LOCK_BUSY</code> | Authentication token lock busy. |
| <code>PAM_AUTHTOK_DISABLE_AGING</code> | Authentication token aging disabled. |
| <code>PAM_USER_UNKNOWN</code> | User unknown to password service. |
| <code>PAM_TRY_AGAIN</code> | Preliminary check by password service failed. |

ATTRIBUTES

See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

`ping(1M)`, `pam(3PAM)`, `pam_chauthtok(3PAM)`, `pam_get_data(3PAM)`, `pam_get_item(3PAM)`, `pam_set_data(3PAM)`, `libpam(3LIB)`, `pam.conf(4)`, `attributes(5)`

NOTES

The PAM framework invokes the password services twice. The first time the modules are invoked with the flag, `PAM_PRELIM_CHECK`. During this stage, the password modules should only perform preliminary checks. For example, they may `ping` remote name services to see if they are ready for updates. If a password module detects a transient error such as a remote name service temporarily down, it should return `PAM_TRY_AGAIN` to the PAM framework, which will immediately return the error back to the application. If all password modules pass the preliminary check, the PAM framework invokes the password services again with the flag, `PAM_UPDATE_AUTHTOK`. During this stage, each password module should proceed to update the appropriate password. Any error will again be reported back to application.

If a service module receives the flag `PAM_CHANGE_EXPIRED_AUTH Tok`, it should check whether the password has aged or expired. If the password has aged or expired, then the service module should proceed to update the password. If the status indicates that the password has not yet aged or expired, then the password module should return `PAM_IGNORE`.

If a user's password has aged or expired, a PAM account module could save this information as state in the authentication handle, *pamh*, using `pam_set_data()`. The related password management module could retrieve this information using `pam_get_data()` to determine whether or not it should prompt the user to update the password for this particular module.

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME pam_sm_open_session, pam_sm_close_session – service provider implementation for pam_open_session and pam_close_session

SYNOPSIS

```
cc [ flag ... ] file ... -lpam [ library ... ]
#include <security/pam_appl.h>
#include <security/pam_modules.h>
int pam_sm_open_session(pam_handle_t *pamh, int flags, int argc, const char **argv);
int pam_sm_close_session(pam_handle_t *pamh, int flags, int argc, const char **argv);
```

DESCRIPTION

In response to a call to pam_open_session(3PAM) and pam_close_session(3PAM), the PAM framework calls pam_sm_open_session() and pam_sm_close_session(), respectively from the modules listed in the pam.conf(4) file. The session management provider supplies the back-end functionality for this interface function.

The pam_sm_open_session() function is called to initiate session management. The pam_sm_close_session() function is invoked when a session has terminated. The argument pamh is an authentication handle. The following flag may be set in the flags field:

PAM_SILENT Session service should not generate any messages.

The argc argument represents the number of module options passed in from the configuration file pam.conf(4). argv specifies the module options, which are interpreted and processed by the session management service. If an unknown option is passed in, an error should be logged through syslog(3C) and the option ignored.

RETURN VALUES Upon successful completion, PAM_SUCCESS should be returned. The following values may also be returned upon error:

PAM_SESSION_ERR Cannot make or remove an entry for the specified session.

PAM_IGNORE Ignore underlying session module regardless of whether the control flag is required, optional or sufficient.

ATTRIBUTES See attributes(5) for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO pam(3PAM), pam_open_session(3PAM), syslog(3C), libpam(3LIB), pam.conf(4), attributes(5)

NOTES

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

| | | | | | | | | | | | |
|-----------------------|--|--------------------|--|-----------------|---|-----------------------|--------------------------------|------------------|--------------------------------------|------------|---|
| NAME | pam_sm_setcred – service provider implementation for pam_setcred | | | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lpam [library ...] #include <security/pam_appl.h> #include <security/pam_modules.h> int pam_sm_setcred(pam_handle_t *pamh, int flags, int argc, const char **argv);</pre> | | | | | | | | | | |
| DESCRIPTION | <p>In response to a call to pam_setcred(3PAM), the PAM framework calls pam_sm_setcred() from the modules listed in the pam.conf(4) file. The authentication provider supplies the back-end functionality for this interface function.</p> <p>The pam_sm_setcred() function is called to set the credentials of the current user associated with the authentication handle, pamh. The following flags may be set in the flags field. Note that the first four flags are mutually exclusive:</p> <table border="0"> <tr> <td>PAM_ESTABLISH_CRED</td> <td>Set user credentials for the authentication service.</td> </tr> <tr> <td>PAM_DELETE_CRED</td> <td>Delete user credentials associated with the authentication service.</td> </tr> <tr> <td>PAM_REINITIALIZE_CRED</td> <td>Reinitialize user credentials.</td> </tr> <tr> <td>PAM_REFRESH_CRED</td> <td>Extend lifetime of user credentials.</td> </tr> <tr> <td>PAM_SILENT</td> <td>Authentication service should not generate messages</td> </tr> </table> <p>If no flag is set, PAM_ESTABLISH_CRED is used as the default.</p> <p>The argc argument represents the number of module options passed in from the configuration file pam.conf(4). argv specifies the module options, which are interpreted and processed by the authentication service. If an unknown option is passed to the module, an error should be logged and the option ignored.</p> <p>If the PAM_SILENT flag is not set, then pam_sm_setcred() should print any failure status from the corresponding pam_sm_authenticate() function using the conversation function.</p> <p>The authentication status (success or reason for failure) is saved as module-specific state in the authentication handle by the authentication module. The status should be retrieved using pam_get_data(), and used to determine if user credentials should be set.</p> | PAM_ESTABLISH_CRED | Set user credentials for the authentication service. | PAM_DELETE_CRED | Delete user credentials associated with the authentication service. | PAM_REINITIALIZE_CRED | Reinitialize user credentials. | PAM_REFRESH_CRED | Extend lifetime of user credentials. | PAM_SILENT | Authentication service should not generate messages |
| PAM_ESTABLISH_CRED | Set user credentials for the authentication service. | | | | | | | | | | |
| PAM_DELETE_CRED | Delete user credentials associated with the authentication service. | | | | | | | | | | |
| PAM_REINITIALIZE_CRED | Reinitialize user credentials. | | | | | | | | | | |
| PAM_REFRESH_CRED | Extend lifetime of user credentials. | | | | | | | | | | |
| PAM_SILENT | Authentication service should not generate messages | | | | | | | | | | |
| RETURN VALUES | <p>Upon successful completion, PAM_SUCCESS should be returned. The following values may also be returned upon error:</p> <table border="0"> <tr> <td>PAM_CRED_UNAVAIL</td> <td>Underlying authentication service can not retrieve user credentials.</td> </tr> </table> | PAM_CRED_UNAVAIL | Underlying authentication service can not retrieve user credentials. | | | | | | | | |
| PAM_CRED_UNAVAIL | Underlying authentication service can not retrieve user credentials. | | | | | | | | | | |

| | |
|------------------|--|
| PAM_CRED_EXPIRED | User credentials have expired. |
| PAM_USER_UNKNOWN | User unknown to the authentication service. |
| PAM_CRED_ERR | Failure in setting user credentials. |
| PAM_IGNORE | Ignore underlying authentication module regardless of whether the control flag is <i>required</i> , <i>optional</i> , or <i>sufficient</i> . |

ATTRIBUTES

See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

`pam(3PAM)`, `pam_authenticate(3PAM)`, `pam_get_data(3PAM)`, `pam_setcred(3PAM)`, `pam_sm_authenticate(3PAM)`, `libpam(3LIB)`, `pam.conf(4)`, `attributes(5)`

NOTES

The `pam_sm_setcred()` function is passed the same module options that are used by `pam_sm_authenticate()`.

The interfaces in `libpam` are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME | pam_start, pam_end – authentication transaction routines for PAM

SYNOPSIS |

```
cc [ flag ... ] file ... -lpam [ library ... ]
#include <security/pam_appl.h>
int pam_start(const char *service, const char *user, const struct pam_conv *pam_conv,
pam_handle_t **pamh);

int pam_end(pam_handle_t *pamh, int status);
```

DESCRIPTION | The `pam_start()` function is called to initiate an authentication transaction. `pam_start()` takes as arguments the name of the current service, `service`, the name of the user to be authenticated, `user`, the address of the conversation structure, `pam_conv`, and the address of a variable to be assigned the authentication handle `pamh`. Upon successful completion, `pamh` refers to a PAM handle for use with subsequent calls to the authentication library.

The `pam_conv` structure contains the address of the conversation function provided by the application. The underlying PAM service module invokes this function to output information to and retrieve input from the user. The `pam_conv` structure has the following entries:

```
struct pam_conv {
    int (*conv)(); /* Conversation function */
    void *appdata_ptr; /* Application data */
};

int conv(int num_msg, const struct pam_message **msg,
        struct pam_response **resp, void *appdata_ptr);
```

The `conv()` function is called by a service module to hold a PAM conversation with the application or user. For window applications, the application can create a new pop-up window to be used by the interaction.

The `num_msg` parameter is the number of messages associated with the call. The parameter `msg` is a pointer to an array of length `num_msg` of the `pam_message` structure.

The `pam_message` structure is used to pass prompt, error message, or any text information from the authentication service to the application or user. It is the responsibility of the PAM service modules to localize the messages. The memory used by `pam_message` has to be allocated and freed by the PAM modules. The `pam_message` structure has the following entries:

```
struct pam_message{
    int msg_style;
    char *msg;
};
```

The message style, `msg_style`, can be set to one of the following values:

| | |
|---------------------|---|
| PAM_PROMPT_ECHO_OFF | Prompt user, disabling echoing of response. |
| PAM_PROMPT_ECHO_ON | Prompt user, enabling echoing of response. |
| PAM_ERROR_MSG | Print error message. |
| PAM_TEXT_INFO | Print general text information. |
| PAM_MSG_NOCONF | Print general text information without user acknowledgment. |
| PAM_CONV_INTERRUPT | Return from the conversation function. |

The maximum size of the message and the response string is PAM_MAX_MSG_SIZE as defined in <security/pam.appl.h> .

The structure *pam_response* is used by the authentication service to get the user's response back from the application or user. The storage used by *pam_response* has to be allocated by the application and freed by the PAM modules. The *pam_response* structure has the following entries:

```
struct pam_response{
    char *resp;
    int  resp_retcode; /* currently not used, */
                          /* should be set to 0 */
};
```

It is the responsibility of the conversation function to strip off NEWLINE characters for PAM_PROMPT_ECHO_OFF and PAM_PROMPT_ECHO_ON message styles, and to add NEWLINE characters (if appropriate) for PAM_ERROR_MSG and PAM_TEXT_INFO message styles.

The *appdata_ptr* argument is an application data pointer which is passed by the application to the PAM service modules. Since the PAM modules pass it back through the conversation function, the applications can use this pointer to point to any application-specific data.

The *pam_end()* function is called to terminate the authentication transaction identified by *pamh* and to free any storage area allocated by the authentication module. The argument, *status*, is passed to the *cleanup()* function stored within the *pam* handle, and is used to determine what module-specific state must be purged. A cleanup function is attached to the handle by the underlying PAM modules through a call to *pam_set_item(3PAM)* to free module specific data.

RETURN VALUES

Refer to *pam(3PAM)* for information on error related return values.

ATTRIBUTES

See *attributes(5)* for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO

libpam(3LIB) , pam(3PAM) , pam_acct_mgmt(3PAM) ,
pam_authenticate(3PAM) , pam_chauthtok(3PAM) ,
pam_open_session(3PAM) , pam_setcred(3PAM) , pam_set_item(3PAM) ,
pam_strerror(3PAM) , attributes(5)

NOTES

The interfaces in libpam are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME pam_strerror – get PAM error message string

SYNOPSIS cc [flag ...] file ... -lpam [library ...]
#include <security/pam_appl.h>

const char *pam_strerror(pam_handle_t*pamh, int errnum);

DESCRIPTION The pam_strerror() function maps the PAM error number in *errnum* to a PAM error message string, and returns a pointer to that string. The application should not free or modify the string returned.

The *pamh* argument is the PAM handle obtained by a prior call to pam_start(). If pam_start() returns an error, a null PAM handle should be passed.

ERRORS The pam_strerror() function returns NULL if *errnum* is out-of-range.

ATTRIBUTES See attributes(5) for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-------------------------|
| Interface Stability | Stable |
| MT-Level | MT-Safe with exceptions |

SEE ALSO pam(3PAM), pam_start(3PAM), attributes(5)

NOTES The interfaces in libpam are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

NAME pathfind – search for named file in named directories

SYNOPSIS cc [*flag* ...] *file* ... -lgen [*library* ...]
#include <libgen.h>
char ***pathfind**(const char **path*, const char **name*, const char **mode*);

DESCRIPTION The `pathfind()` function searches the directories named in *path* for the file *name*. The directories named in *path* are separated by colons (:). The *mode* argument is a string of option letters chosen from the set [*rwxfbcdpugks*] :

| Letter | Meaning |
|--------|-------------------|
| r | readable |
| w | writable |
| x | executable |
| f | normal file |
| b | block special |
| c | character special |
| d | directory |
| p | FIFO (pipe) |
| u | set user ID bit |
| g | set group ID bit |
| k | sticky bit |
| s | size non-zero |

Options read, write, and execute are checked relative to the real (not the effective) user ID and group ID of the current process.

If *name* begins with a slash, it is treated as an absolute path name, and *path* is ignored.

An empty *path* member is treated as the current directory. A slash (/) character is not prepended at the occurrence of the first match; rather, the unadorned *name* is returned.

EXAMPLES

EXAMPLE 1 Example of finding the `ls` command using the `PATH` environment variable.

To find the `ls` command using the `PATH` environment variable:

```
pathfind (getenv ("PATH"), "ls", "rx")
```

RETURN VALUES

The `pathfind()` function returns a `(char *)` value containing static, thread-specific data that will be overwritten upon the next call from the same thread.

If the file *name* with all characteristics specified by *mode* is found in any of the directories specified by *path*, then `pathfind()` returns a pointer to a string containing the member of *path*, followed by a slash character (`/`), followed by *name*.

If no match is found, `pathfind()` returns a null pointer, `((char *) 0)`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`sh(1)`, `test(1)`, `access(2)`, `mknod(2)`, `stat(2)`, `getenv(3C)`, `attributes(5)`

NOTES

The string pointed to by the returned pointer is stored in an area that is reused on subsequent calls to `pathfind()`. The string should not be deallocated by the caller.

When compiling multithreaded applications, the `_REENTRANT` flag must be defined on the compile line. This flag should only be used in multithreaded applications.

| | |
|---|--|
| NAME | pctx_capture, pctx_create, pctx_run, pctx_release – process context library |
| SYNOPSIS | <pre>cc [flag...] file... -lpctx [library...] #include <libpctx.h> typedef void (pctx_errfn_t)(const char *fn, const char *fmt, va_list ap); pctx_t *pctx_create(const char *filename, char *const *argv, void *arg, int verbose, pctx_errfn_t *errfn); pctx_t *pctx_capture(pid_t pid, void *arg, int verbose, pctx_errfn_t *errfn); int pctx_run(pctx_t *pctx, uint_t sample, uint_t nsamples, int (*tick)(pctx *, pid_t, id_t, void *)); void pctx_release(pctx_t *pctx);</pre> |
| DESCRIPTION | <p>This family of functions allows a controlling process (the process that invokes them) to create or capture controlled processes. The functions allow the occurrence of various events of interest in the controlled process to cause the controlled process to be stopped, and to cause callback routines to be invoked in the controlling process.</p> <p>There are two ways a process can be acquired by the process context functions. First, a named application can be invoked with the usual <i>argv</i> [] array using <code>pctx_create()</code>, which forks the caller and <code>exec</code>s the application in the child. Alternatively, an existing process can be captured by its process ID using <code>pctx_capture()</code>.</p> <p>Both functions accept a pointer to an opaque handle, <i>arg</i>; this is saved and treated as a caller-private handle that is passed to the other functions in the library. Both functions accept a pointer to a <code>fork(3C)</code>-like error routine <i>errfn</i>; a default version is provided if <code>NULL</code> is specified.</p> <p>A freshly-created process is created stopped; similarly, a process that has been successfully captured is stopped by the act of capturing it, thereby allowing the caller to specify the handlers that should be called when various events occur in the controlled process. The set of handlers is listed on the <code>pctx_set_events(3CPC)</code> manual page.</p> <p>Once the callback handlers have been set with <code>pctx_set_events()</code>, the application can be set running using <code>pctx_run()</code>. This function starts the event handling loop; it returns only when either the process has exited, the number of time samples has expired, or an error has occurred (for example, if the controlling process is not privileged, and the controlled process has <code>exec</code>-ed a <code>setuid</code> program).</p> <p>Every <i>sample</i> milliseconds the process is stopped and the <code>tick()</code> routine is called so that, for example, the performance counters can be sampled by the caller. No periodic sampling is performed if <i>sample</i> is 0.</p> |
| <code>pctx_create()</code> and <code>pctx_capture()</code> | |
| <code>pctx_run()</code> | |

pctx_release()

Once `pctx_run()` has returned, the process can be released and the underlying storage freed using `pctx_release()`. Releasing the process will either allow the controlled process to continue (in the case of an existing captured process and its children) or kill the process (if it and its children were created using `pctx_create()`).

RETURN VALUES

Upon successful completion, `pctx_capture()` and `pctx_create()` return a valid handle. Otherwise, the functions print a diagnostic message and return `NULL`.

Upon successful completion, `pctx_run()` returns 0 with the controlled process either stopped or exited (if the controlled process has invoked `exit(2)`.) If an error has occurred (for example, if the controlled process has `exec`-ed a set-ID executable, if certain callbacks have returned error indications, or if the process was unable to respond to `proc(4)` requests) an error message is printed and the function returns -1.

USAGE

Within an event handler in the controlling process, the controlled process can be made to perform various system calls on its behalf. No system calls are directly supported in this version of the API, though system calls are executed by the `cpc_pctx` family of interfaces in `libcpc` such as `cpc_pctx_bind_event(3CPC)`. A specially created agent LWP is used to execute these system calls in the controlled process. See `proc(4)` for more details.

While executing the event handler functions, the library arranges for the signals `SIGTERM`, `SIGQUIT`, `SIGABRT`, and `SIGINT` to be blocked to reduce the likelihood of a keyboard signal killing the controlling process prematurely, thereby leaving the controlled process permanently stopped while the agent LWP is still alive inside the controlled process.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|---|
| MT-Level | Unsafe |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) |
| Interface Stability | Evolving |

SEE ALSO

`fork(2)`, `cpc(3CPC)`, `pctx_set_events(3CPC)`, `proc(4)`, `attributes(5)`.

| | |
|--------------------|--|
| NAME | pctx_set_events – associate callbacks with process events |
| SYNOPSIS | <pre>cc [flag...] file... -lpctx [library...] #include <libpctx.h> typedef enum { PCTX_NULL_EVENT = 0, PCTX_SYSC_EXEC_EVENT, PCTX_SYSC_FORK_EVENT, PCTX_SYSC_EXIT_EVENT, PCTX_SYSC_LWP_CREATE_EVENT, PCTX_INIT_LWP_EVENT, PCTX_FINI_LWP_EVENT, PCTX_SYSC_LWP_EXIT_EVENT } pctx_event_t; typedef int pctx_sysc_execfn_t(pctx_t *pctx, pid_t pid, id_t lwpid, char *cmd, void *arg); typedef void pctx_sysc_forkfn_t(pctx_t *pctx, pid_t pid, id_t lwpid, pid_t child, void *arg); typedef void pctx_sysc_exitfn_t(pctx_t *pctx, pid_t pid, id_t lwpid, void *arg); typedef int pctx_sysc_lwp_createfn_t(pctx_t *pctx, pid_t pid, id_t lwpid, void *arg); typedef int pctx_init_lwpfn_t(pctx_t *pctx, pid_t pid, id_t lwpid, void *arg); typedef int pctx_fini_lwpfn_t(pctx_t *pctx, pid_t pid, id_t lwpid, void *arg); typedef int pctx_sysc_lwp_exitfn_t(pctx_t *pctx, pid_t pid, id_t lwpid, void *arg); int pctx_set_events(pctx_t *pctx, ...);</pre> |
| DESCRIPTION | <p>The <code>pctx_set_events()</code> function allows the caller (the controlling process) to express interest in various events in the controlled process. See <code>pctx_capture(3CPC)</code> for information about how the controlling process is able to create, capture and manipulate the controlled process.</p> <p>The <code>pctx_set_events()</code> function takes a <code>pctx_t</code> handle, followed by a variable length list of pairs of <code>pctx_event_t</code> tags and their corresponding handlers, terminated by a <code>PCTX_NULL_EVENT</code> tag.</p> <p>Most of the events correspond closely to various classes of system calls, though two additional pseudo-events (<i>init_lwp</i> and <i>fini_lwp</i>) are provided to allow callers to perform various housekeeping tasks. The <i>init_lwp</i> handler is called as soon as the library identifies a new LWP, while <i>fini_lwp</i> is called just before the LWP disappears. Thus the classic "hello world" program would see an <i>init_lwp</i> event, a <i>fini_lwp</i> event and (process) <i>exit</i> event, in that order. The table below displays the interactions between the states of the controlled process and the handlers executed by users of the library.</p> |

| System Calls and pctx Handlers | | |
|--------------------------------|------------|--|
| System call | Handler | Comments |
| exec(2), execve(2) | fini_lwp | Invoked serially on all lwps in the process. |
| | exec | Only invoked if the exec() system call succeeded. |
| | init_lwp | If the exec succeeds, only invoked on lwp 1. If the exec fails, invoked serially on all lwps in the process. |
| fork(2), vfork(2), fork1(2) | fork | Only invoked if the fork() system call succeeded. |
| exit(2) | fini_lwp | Invoked on all lwps in the process. |
| | exit | Invoked on the exiting lwp. |
| _lwp_create(2) | init_lwp | Only if the corresponding _lwp_create() system call succeeded. |
| | lwp_create | |
| _lwp_exit(2) | fini_lwp | |
| | lwp_exit | |

Each of the handlers is passed the caller's opaque handle, a pctx_t handle, the pid, and lwpid of the process and lwp generating the event. The *lwp_exit*, and (process) *exit* events are delivered *before* the underlying system calls begin, while the *exec*, *fork*, and *lwp_create* events are only delivered after the relevant system calls complete successfully. The *exec* handler is passed a string that describes the command being executed. Catching the *fork* event causes the calling process to *fork(2)*, then capture the child of the controlled process using *pctx_capture()* before handing control to the *fork* handler. The process is released on return from the handler.

RETURN VALUES

Upon successful completion, *pctx_set_events()* returns 0. Otherwise, the function returns -1.

EXAMPLES

CODE EXAMPLE 1 HandleExec example.

This example captures an existing process whose process identifier is *pid*, and arranges to call the *HandleExec* routine when the process performs an *exec(2)*.

```
static void
HandleExec(pctx_t *pctx, pid_t pid, id_t lwpid, char *cmd, void *arg)
{
    (void) printf("pid %d execed '%s'\n", (int)pid, cmd);
}
int
main()
```

```

{
  ...
  pctx = pctx_capture(pid, NULL, 1, NULL);
  (void) pctx_set_events(pctx,
    PCTX_SYSC_EXEC_EVENT, HandleExec,
    ...
    PCTX_NULL_EVENT);
  (void) pctx_run(pctx, 0, 0, NULL);
  pctx_release(pctx);
}

```

ATTRIBUTES

See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|---|
| MT-Level | Unsafe |
| Availability | SUNWcpcu (32-bit) SUNWcpcux (64-bit) |
| Interface Stability | Evolving |

SEE ALSO

[exec\(2\)](#), [exit\(2\)](#), [fork\(2\)](#), [vfork\(2\)](#), [forkl\(2\)](#), [_lwp_create\(2\)](#), [_lwp_exit\(2\)](#), [cpc\(3CPC\)](#), [proc\(4\)](#), [attributes\(5\)](#).

NAME pow – power function

SYNOPSIS cc [flag ...] file ... -lm [library ...]
#include <math.h>
double pow(double x, double y);

DESCRIPTION The pow() function computes the value of x raised to the power y, x^y. If x is negative, y must be an integer value.

RETURN VALUES Upon successful completion, pow() returns the value of x raised to the power y. If x is 0 and y is 0, 1.0 is returned. If y is NaN, or y is non-zero and x is NaN, NaN is returned. If y is 0.0 and x is NaN, NaN is returned. If x is 0.0 and y is negative, -HUGE_VAL is returned and errno may be set to EDOM or ERANGE. If the correct value would cause overflow, ±HUGE_VAL is returned, and errno is set to ERANGE. If the correct value would cause underflow to 0, 0 is returned and errno may be set to ERANGE. For exceptional cases, matherr(3M) tabulates the values to be returned as dictated by Standards other than XPG4.

ERRORS The pow() function will fail if:
EDOM The value of x is negative and y is non-integral.
ERANGE The value to be returned would have caused overflow.
The pow() function may fail if:
EDOM The value of x is 0.0 and y is negative.
ERANGE The correct value would cause underflow.

USAGE An application wishing to check for error situations should set errno to 0 before calling pow(). If errno is non-zero on return, or the return value is NaN, an error has occurred.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO exp(3M), isnan(3M), matherr(3M), attributes(5), standards(5)

NAME | printDmiAttributeValues – print data in input DmiAttributeValues list

SYNOPSIS | `cc [flag ...] file ... -ldmi -lnsl -lrwtool [library ...]`
`#include <dmi/util.hh>`
`void printDmiAttributeValues(DmiAttributeValues_t *values);`

DESCRIPTION | The `printDmiAttributeValues()` function prints the data in the input `DmiAttributeValues` list. The function prints "unknown data" for those *values* that contain invalid data.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | MT-Safe |

SEE ALSO | `libdmi(3LIB)`, `attributes(5)`

NAME | printDmiDataUnion – print data in input data union

SYNOPSIS | `cc [flag ...] file ... -ldmi -lnsl -lrwtool [library ...]`
`#include <dmi/util.hh>`
`void printDmiDataUnion(DmiDataUnion_t *data);`

DESCRIPTION | The `printDmiDataUnion()` function prints the data in the input data union. The output depends on the type of DMI data in the union.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | MT-Safe |

SEE ALSO | `libdmi(3LIB)`, `attributes(5)`

NAME | printDmiString – print a DmiString

SYNOPSIS | `cc [flag ...] file ... -ldmi -lnsl -lrwtool [library ...]`
`#include <dmi/util.hh>`
`void printDmiString(DmiString_t *dstr);`

DESCRIPTION | The `printDmiString()` function prints a DmiString.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-level | MT-Safe |

SEE ALSO | `newDmiString(3DMI)`, `libdmi(3LIB)`, `attributes(5)`

| NAME | read_vtoc, write_vtoc – read and write a disk's VTOC | | | | |
|----------------------|---|----------------|-----------------|----------|--------|
| SYNOPSIS | <pre>cc [flag ...] file ... -ladm [library ...] #include <sys/vtoc.h> int read_vtoc(int fd, struct vtoc *vtoc); int write_vtoc(int fd, struct vtoc *vtoc);</pre> | | | | |
| DESCRIPTION | <p>The <code>read_vtoc()</code> function returns the VTOC (volume table of contents) structure that is stored on the disk associated with the open file descriptor <code>fd</code>.</p> <p>The <code>write_vtoc()</code> function stores the VTOC structure on the disk associated with the open file descriptor <code>fd</code>.</p> <p>The <code>fd</code> argument refers to any slice on a raw disk.</p> | | | | |
| RETURN VALUES | <p>Upon successful completion, <code>read_vtoc()</code> returns a positive number indicating the slice index associated with the open file descriptor. Otherwise, it returns a negative number indicating one of the following errors:</p> <p>VT_EIO An I/O error occurred.</p> <p>VT_ERROR An unknown error occurred.</p> <p>Upon successful completion, <code>write_vtoc()</code> returns 0. Otherwise, it returns a negative number indicating one of the following errors:</p> <p>VT_EIO An I/O error occurred.</p> <p>VT_ERROR An unknown error occurred.</p> <p>VT_EINVAL The VTOC contains an incorrect field.</p> | | | | |
| ATTRIBUTES | <p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Unsafe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | Unsafe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | Unsafe | | | | |
| SEE ALSO | <code>fmthard(1M)</code> , <code>format(1M)</code> , <code>prtvtoc(1M)</code> , <code>ioctl(2)</code> , <code>attributes(5)</code> , <code>dkio(7I)</code> | | | | |
| BUGS | The <code>write_vtoc()</code> function cannot write a VTOC on an unlabeled disk. Use <code>format(1M)</code> for this purpose. | | | | |

| NAME | reg_ci_callback – provide a component instrumentation with a transient program number | | | | |
|----------------------|--|----------------|-----------------|----------|-------|
| SYNOPSIS | <pre>cc [flag ...] file ... -ldmici [library ...] #include <dmi/ci_callback_svc.hh> u_long reg_ci_callback();</pre> | | | | |
| DESCRIPTION | The reg_ci_callback() function provides a component instrumentation with a transient program number. The instrumentation uses this number to register its RPC service provider. The prognum member of the DmiRegisterInfo structure is populated with the return value of this function | | | | |
| RETURN VALUES | Upon successful completion, the reg_ci_callback() function returns a transient program number of type u_long. | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-level</td> <td>Unafe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-level | Unafe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-level | Unafe | | | | |
| SEE ALSO | attributes(5) | | | | |

| | |
|--------------------|--|
| NAME | regexpr, compile, step, advance – regular expression compile and match routines |
| SYNOPSIS | <pre>cc [flag...] [file...] -lgen [library...]</pre> <pre>#include <regexpr.h></pre> <pre>char *compile(char *instring, char *expbuf, const char *endbuf);</pre> <pre>int step(const char *string, const char *expbuf);</pre> <pre>int advance(const char *string, const char *expbuf);</pre> <pre>extern char *loc1, loc2, locs;</pre> <pre>extern int nbra, regerrno, reglength;</pre> <pre>extern char *braslist[], *braelist[];</pre> |
| DESCRIPTION | <p>These routines are used to compile regular expressions and match the compiled expressions against lines. The regular expressions compiled are in the form used by ed(1).</p> <p>The parameter <i>instring</i> is a null-terminated string representing the regular expression.</p> <p>The parameter <i>expbuf</i> points to the place where the compiled regular expression is to be placed. If <i>expbuf</i> is NULL, <code>compile()</code> uses <code>malloc(3C)</code> to allocate the space for the compiled regular expression. If an error occurs, this space is freed. It is the user's responsibility to free unneeded space after the compiled regular expression is no longer needed.</p> <p>The parameter <i>endbuf</i> is one more than the highest address where the compiled regular expression may be placed. This argument is ignored if <i>expbuf</i> is NULL. If the compiled expression cannot fit in (<i>endbuf</i> - <i>expbuf</i>) bytes, <code>compile()</code> returns NULL and <code>regerrno</code> (see below) is set to 50.</p> <p>The parameter <i>string</i> is a pointer to a string of characters to be checked for a match. This string should be null-terminated.</p> <p>The parameter <i>expbuf</i> is the compiled regular expression obtained by a call of the function <code>compile()</code>.</p> <p>The function <code>step()</code> returns non-zero if the given string matches the regular expression, and zero if the expressions do not match. If there is a match, two external character pointers are set as a side effect to the call to <code>step()</code>. The variables set in <code>step()</code> are <code>loc1</code> and <code>loc2</code>. <code>loc1</code> is a pointer to the first character that matched the regular expression. The variable <code>loc2</code> points to the character after the last character that matches the regular expression. Thus if the regular expression matches the entire line, <code>loc1</code> points to the first character of <i>string</i> and <code>loc2</code> points to the null at the end of <i>string</i>.</p> |

The purpose of `step()` is to step through the *string* argument until a match is found or until the end of *string* is reached. If the regular expression begins with `^`, `step()` tries to match the regular expression at the beginning of the string only.

The `advance()` function is similar to `step()`; but, it only sets the variable `loc2` and always restricts matches to the beginning of the string.

If one is looking for successive matches in the same string of characters, `locs` should be set equal to `loc2`, and `step()` should be called with *string* equal to `loc2`. `locs` is used by commands like `ed` and `sed` so that global substitutions like `s/y*/g` do not loop forever, and is `NULL` by default.

The external variable `nbra` is used to determine the number of subexpressions in the compiled regular expression. `braslist` and `braelist` are arrays of character pointers that point to the start and end of the `nbra` subexpressions in the matched string. For example, after calling `step()` or `advance()` with *string* `sabcdefg` and regular expression `\\(abcdef\\)`, `braslist[0]` will point at `a` and `braelist[0]` will point at `g`. These arrays are used by commands like `ed` and `sed` for substitute replacement patterns that contain the `\\ n` notation for subexpressions.

Note that it is not necessary to use the external variables `regerrno`, `nbra`, `loc1`, `loc2`, `locs`, `braelist`, and `braslist` if one is only checking whether or not a string matches a regular expression.

EXAMPLES

EXAMPLE 1 The following is similar to the regular expression code from `grep`:

```
#include<regexpr.h>
.
.
.
if(compile(*argv, (char *)0, (char *)0) == (char *)0)
    regerr(regerrno);
.
.
.
if (step(linebuf, expbuf))
    succeed();
```

RETURN VALUES

If `compile()` succeeds, it returns a non-`NULL` pointer whose value depends on *expbuf*. If *expbuf* is non-`NULL`, `compile()` returns a pointer to the byte after the last byte in the compiled regular expression. The length of the compiled regular expression is stored in `reglength`. Otherwise, `compile()` returns a pointer to the space allocated by `malloc(3C)`.

The functions `step()` and `advance()` return non-zero if the given string matches the regular expression, and zero if the expressions do not match.

ERRORS

If an error is detected when compiling the regular expression, a `NULL` pointer is returned from `compile()` and `regerrno` is set to one of the non-zero error numbers indicated below:

| ERROR | MEANING |
|--------------|---|
| 11 | Range endpoint too large. |
| 16 | Bad Number. |
| 25 | "\\digit" out of range. |
| 36 | Illegal or missing delimiter. |
| 41 | No remembered string search. |
| 42 | \\(~\\) imbalance. |
| 43 | Too many \\(. |
| 44 | More than 2 numbers given in \\[~\\]. |
| 45 | } expected after \\. |
| 46 | First number exceeds second in \\{~\\}. |
| 49 | [] imbalance. |
| 50 | Regular expression overflow. |

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`ed(1)`, `grep(1)`, `sed(1)`, `malloc(3C)`, `attributes(5)`, `regexp(5)`

NOTES

When compiling multi-threaded applications, the `_REENTRANT` flag must be defined on the compile line. This flag should only be used in multi-threaded applications.

| NAME | remainder – remainder function | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>#include <math.h> double remainder(double x, double y);</pre> | | | | |
| DESCRIPTION | The <code>remainder()</code> function returns the floating point remainder $r = x - ny$ when y is non-zero. The value n is the integral value nearest the exact value x/y . When $ n - x/y = \frac{1}{2}$, the value n is chosen to be even. The behavior of <code>remainder()</code> is independent of the rounding mode. | | | | |
| RETURN VALUES | The <code>remainder()</code> function returns the floating point remainder $r = x - ny$ when y is non-zero. When y is 0, <code>remainder()</code> returns NaN. and sets <code>errno</code> to <code>EDOM</code> . If the value of x is $\pm\text{Inf}$, <code>remainder()</code> returns NaN and sets <code>errno</code> to <code>EDOM</code> . If x or y is NaN, then the function returns NaN. | | | | |
| ERRORS | The <code>remainder()</code> function will fail if: <code>EDOM</code> The y argument is 0 or the x argument is positive or negative infinity. | | | | |
| USAGE | The <code>remainder()</code> function computes the remainder $x \text{ REM } y$ required by ANSI/IEEE 754 (IEC 559). | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>fmod(3M)</code> , <code>attributes(5)</code> | | | | |

| NAME | rint – round-to-nearest integral value | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double rint(double <i>x</i>); | | | | |
| DESCRIPTION | The rint() function returns the integral value (represented as a double) nearest <i>x</i> in the direction of the current IEEE754 rounding mode. If the current rounding mode rounds toward negative infinity, then rint() is identical to floor(3M). If the current rounding mode rounds toward positive infinity, then rint() is identical to ceil(3M). | | | | |
| RETURN VALUES | Upon successful completion, the rint() function returns the integer (represented as a double precision number) nearest <i>x</i> in the direction of the current IEEE754 rounding mode. When <i>x</i> is ±Inf, rint() returns <i>x</i> . If the value of <i>x</i> is NaN, NaN is returned. | | | | |
| ERRORS | No errors will occur. | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | ceil(3M), floor(3M), isnan(3M), attributes(5) | | | | |

NAME | scalb – load exponent of a radix-independent floating-point number

SYNOPSIS | #include <math.h>
 double **scalb**(double *x*, double *n*);

DESCRIPTION | The `scalb()` function computes $x * r^n$, where *r* is the radix of the machine's floating point arithmetic. When *r* is 2, `scalb()` is equivalent to `ldexp(3C)`.

RETURN VALUES | Upon successful completion, the `scalb()` function returns $x * r^n$.
 If the correct value would overflow, `scalb()` returns `±HUGE_VAL` (according to the sign of *x*) and sets `errno` to `ERANGE`.
 If the correct value would underflow to 0.0, `scalb()` returns 0 and sets `errno` to `ERANGE`.
 The `scalb()` function returns *x* when *x* is `±Inf`.
 If *x* or *n* is NaN, then `scalb()` returns NaN.
 For exceptional cases, `matherr(3M)` tabulates the values to be returned as dictated by Standards other than XPG4.

ERRORS | The `scalb()` function will fail if:
`ERANGE` The correct value would overflow or underflow.

USAGE | An application wishing to check for error situations should set `errno` to 0 before calling `scalb()`. If `errno` is non-zero on return, or the return value is NaN, an error has occurred.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `ldexp(3C)`, `matherr(3M)`, `attributes(5)`

| NAME | scalbn – load exponent of a radix-independent floating-point number | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double scalbn (double <i>x</i> , int <i>n</i>); | | | | |
| DESCRIPTION | The <code>scalbn()</code> function computes $x * r^n$, where r is the radix of the machine's floating point arithmetic. | | | | |
| RETURN VALUES | Upon successful completion, the <code>scalbn()</code> function returns $x * r^n$. If the correct value would overflow, <code>scalbn()</code> returns \pm HUGE_VAL (according to the sign of x). The <code>scalbn()</code> function returns x when x is \pm Inf. If x is NaN, then <code>scalbn()</code> returns NaN. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>attributes(5)</code> | | | | |

NAME | significand – significand function

SYNOPSIS | `cc [flag ...] file ... -lm [library ...]`
`#include <math.h>`
`double significand(double x);`

DESCRIPTION | The `significand()` function, along with the `logb(3M)` and `scalb(3M)` functions, allows users to verify compliance to ANSI/IEEE Std 754-1985 by running certain test vectors distributed by the University of California.

If x equals $sig * 2^n$ with $1 < sig < 2$, then `significand(x)` returns sig for exercising the fraction-part(F) test vector. `significand(x)` is not defined when x is either 0, $\pm Inf$ or NaN.

RETURN VALUES | For exceptional cases, `matherr(3M)` tabulates the values to be returned as dictated by various Standards.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `logb(3M)`, `matherr(3M)`, `scalb(3M)`, `attributes(5)`

NAME sin – sine function

SYNOPSIS cc [*flag ...*] *file ...* -lm [*library ...*]
#include <math.h>
double **sin**(double *x*);

DESCRIPTION The `sin()` function computes the sine of its argument *x*, measured in radians.

RETURN VALUES Upon successful completion, `sin()` returns the sine of *x*.
If *x* is NaN or \pm Inf, NaN is returned.

ERRORS No errors will occur.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO `asin(3M)`, `isnan(3M)`, `attributes(5)`

| NAME | sinh – hyperbolic sine function | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double sinh (double <i>x</i>); | | | | |
| DESCRIPTION | The <code>sinh()</code> function computes the hyperbolic sine of <i>x</i> . | | | | |
| RETURN VALUES | Upon successful completion, <code>sinh()</code> returns the hyperbolic sine of <i>x</i> . If the result would cause an overflow, <code>±HUGE_VAL</code> is returned and <code>errno</code> is set to <code>ERANGE</code> . If <i>x</i> is NaN, NaN is returned. For exceptional cases, <code>matherr(3M)</code> tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The <code>sinh()</code> function will fail if: <code>ERANGE</code> The result would cause overflow. | | | | |
| USAGE | An application wishing to check for error situations should set <code>errno</code> to 0 before calling <code>sinh()</code> . If <code>errno</code> is non-zero on return, or the return value is NaN, an error has occurred. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>asinh(3M)</code> , <code>cosh(3M)</code> , <code>isnan(3M)</code> , <code>matherr(3M)</code> , <code>tanh(3M)</code> , <code>attributes(5)</code> , <code>standards(5)</code> | | | | |

| NAME | sqrt – square root function | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] <i>file ...</i> -lm [<i>library ...</i>] #include <math.h> double sqrt (double <i>x</i>); | | | | |
| DESCRIPTION | The <code>sqrt()</code> function computes the square root of <i>x</i> . | | | | |
| RETURN VALUES | Upon successful completion, <code>sqrt()</code> returns the square root of <i>x</i> . If <i>x</i> is NaN, NaN is returned. If <i>x</i> is negative, NaN is returned and <code>errno</code> is set to EDOM. | | | | |
| ERRORS | The <code>sqrt()</code> function will fail if: EDOM The value of <i>x</i> is negative. | | | | |
| USAGE | An application wishing to check for error situations should set <code>errno</code> to 0 before calling <code>sqrt()</code> . If <code>errno</code> is non-zero on return, or the return value is NaN, an error has occurred. | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>isnan(3M)</code> , <code>attributes(5)</code> | | | | |

NAME SSAAgentIsAlive, SSAGetTrapPort, SSAREgSubtable, SSAREgSubagent, SSAREgSubtree, SSASendTrap, SSASubagentOpen – Sun Solstice Enterprise Agent registration and communication helper functions

SYNOPSIS

```
cc [ flag ... ] file ... -lssagent -lssasnmp [ library .. ]
#include <impl.h>
extern int SSAAgentIsAlive(IPAddress *agent_addr, int *port, char *community, struct
timeval *timeout);

extern int SSAGetTrapPort();

extern int *SSAREgSubagent(Agent* agent);

int SSAREgSubtable(SSA_Table *table);

int SSAREgSubtree(SSA_Subtree *subtree);

extern void SSASendTrap(char *name);

extern int SSASubagentOpen(int *num_of_retry, char *agent_name);
```

DESCRIPTION

The **SSAAgentIsAlive**() function returns TRUE if the master agent is alive, otherwise returns FALSE . The *agent_addr* parameter is the address of the agent. Specify the security token in the *community* parameter. You can specify the maximum amount of time to wait for a response with the *timeout* parameter.

The **SSAGetTrapPort**() function returns the port number used by the Master Agent to communicate with the subagent.

The **SSAREgSubagent**() function enables a subagent to register and unregister with a Master Agent. The *agent* parameter is a pointer to an Agent structure containing the following members:

```
int      timeout;          /* optional */
int      agent_id;        /* required */
int      agent_status;    /* required */
char     *personal_file;  /* optional */
char     *config_file;    /* optional */
char     *executable;    /* optional */
char     *version_string; /* optional */
char     *protocol;       /* optional */
int      process_id;      /* optional */
char     *name;           /* optional */
int      system_up_time;  /* optional */
int      watch_dog_time;  /* optional */
Address  address;         /* required */
struct   _Agent;         /* reserved */
struct   _Subtree;       /* reserved */
```

The `agent_id` member is an integer value returned by the `SSASubagentOpen()` function. After calling `SSASubagentOpen()`, you pass the `agent_id` in the `SSARegSubagent()` call to register the subagent with the Master Agent.

The following values are supported for `agent_status`:

```
SSA_OPER_STATUS_ACTIVE
SSA_OPER_STATUS_NOT_IN_SERVICE
SSA_OPER_STATUS_DESTROY
```

You pass `SSA_OPER_STATUS_DESTROY` as the value in a `SSARegSubagent()` function call when you want to unregister the agent from the Master Agent.

`Address` has the same structure as `sockaddr_in`, that is a common UNIX structure containing the following members:

```
short    sin_family;
ushort_t sin_port;
struct   in_addr sin_addr;
char     sin_zero[8];
```

The `SSARegSubtable()` function registers a MIB table with the Master Agent. If this function is successful, an index number is returned, otherwise 0 is returned. The `table` parameter is a pointer to a `SSA_Table` structure containing the following members:

```
int regTblIndex;      /* index value */
int regTblAgentID;    /* current agent ID */
Oid regTblOID;        /* Object ID of the table */
int regTblStartColumn; /* start column index */
int regTblEndColumn;  /* end column index */
int regTblStartRow;   /* start row index */
int regTblEndRow;     /* end row index */
int regTblStatus;     /* status */
```

The `regTblStatus` can have one of the following values:

```
SSA_OPER_STATUS_ACTIVE
SSA_OPER_STATUS_NOT_IN_SERVICE
```

The `SSARegSubtree()` function registers a MIB subtree with the master agent. If successful this function returns an index number, otherwise 0 is returned. The `subtree` parameter is a pointer to a `SSA_Subtree` structure containing the following members:

```
int regTreeIndex;     /* index value */
int regTreeAgentID;   /* current agent ID */
Oid name;             /* Object ID to register */
int regtreeStatus;    /* status */
```

The `regtreeStatus` can have one of the following values:

```
SSA_OPER_STATUS_ACTIVE
SSA_OPER_STATUS_NOT_IN_SERVICE
```

The `SSASendTrap()` function instructs the Master Agent to send a trap notification, based on the keyword passed with *name*. When your subagent MIB is compiled by `mibcodegen`, it creates a lookup table of the trap notifications defined in the MIB. By passing the name of the trap notification type as *name*, the subagent instructs the Master Agent to construct the type of trap defined in the MIB.

The `SSASubagentOpen()` function initializes communication between the subagent and the Master Agent. You must call this function before calling `SSARegSubagent()` to register the subagent with the Master Agent. The `SSASubagentOpen()` function returns a unique agent ID that is passed in the `SSARegSubagent()` call to register the subagent. If 0 is returned as the agent ID, the attempt to initialize communication with the Master Agent was unsuccessful. Since UDP is used to initialize communication with the Master Agent, you may want to set the value of *num_of_retry* to make multiple attempts.

The value for *agent_name* must be unique within the domain for which the Master Agent is responsible.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Unsafe |

SEE ALSO

`attributes(5)`

| | |
|--------------------|---|
| NAME | SSAOidCmp, SSAOidCpy, SSAOidDup, SSAOidFree, SSAOidInit, SSAOidNew, SSAOidString, SSAOidStrToOid, SSAOidZero – Sun Solstice Enterprise Agent OID helper functions |
| SYNOPSIS | <pre>cc [flag ...] file ... -lssasnmp [library ..] #include <impl.h> int SSAOidCmp(Oid *oid1, Oid *oid2); int SSAOidCpy(Oid *oid1, Oid *oid2, char *error_label); Oid *SSAOidDup(Oid *oid, char *error_label); void SSAOidFree(Oid *oid); int SSAOidInit(Oid *oid, Subid *subids, int len, char *error_label); Oid *SSAOidNew(); char *SSAOidString(Oid *oid); Oid *SSAOidStrToOid(char* name, char *error_label); void SSAOidZero(Oid *oid);</pre> |
| DESCRIPTION | <p>The SSAOidCmp() function performs a comparison of the given OIDs. This function returns:</p> <ul style="list-style-type: none"> 0 if <i>oid1</i> is equal to <i>oid2</i> 1 if <i>oid1</i> is greater than <i>oid2</i> -1 if <i>oid1</i> is less than <i>oid2</i> <p>The SSAOidCpy() function makes a deep copy of <i>oid2</i> to <i>oid1</i> . This function assumes <i>oid1</i> has been processed by the SSAOidZero() function. Memory is allocated inside <i>oid1</i> and the contents of <i>oid2</i> , not just the pointer, is copied to <i>oid1</i> . If an error is encountered, an error message is stored in the <i>error_label</i> buffer.</p> <p>The SSAOidDup() function returns a clone of <i>oid</i> , by using the deep copy. Error information is stored in the <i>error_label</i> buffer.</p> <p>The SSAOidFree() function frees the OID instance, with its content.</p> <p>The SSAOidNew() function returns a new OID.</p> <p>The SSAOidInit() function copies the Subid array from <i>subids</i> to the OID instance with the specified length <i>len</i> . This function assumes that the OID instance has been processed by the SSAOidZero() function or no memory is allocated inside the OID instance. If an error is encountered, an error message is stored in the <i>error_label</i> buffer.</p> |

The `SSAOidString()` function returns a char pointer for the printable form of the given *oid*.

The `SSAOidStrToOid()` function returns a new OID instance from *name*. If an error is encountered, an error message is stored in the *error_label* buffer.

The `SSAOidZero()` function frees the memory used by the OID object for buffers, but not the OID instance itself.

RETURN VALUES

The `SSAOidNew()` and `SSAOidStrToOid()` functions return 0 if an error is detected.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Unsafe |

SEE ALSO

`attributes(5)`

NAME SSAStringCpy, SSAStringInit, SSAStringToChar, SSAStringZero – Sun Solstice Enterprise Agent string helper functions

SYNOPSIS

```
cc [ flag ... ] file ... -lssasnmp [ library .. ]
#include <impl.h>
void *SSAStringZero(String *string);

int SSAStringInit(String *string, uchar_t *chars, int len, char *error_label);

int SSAStringCpy(String *string1, String *string2, char *error_label);

char *SSAStringToChar(String string);
```

DESCRIPTION

The SSAStringCpy() function makes a deep copy of *string2* to *string1*. This function assumes that *string1* has been processed by the SSAStringZero() function. Memory is allocated inside the *string1* and the contents of *string2*, not just the pointer, is copied to the *string1*. If an error is encountered, an error message is stored in the *error_label* buffer.

The SSAStringInit() function copies the char array from *chars* to the string instance with the specified length *len*. This function assumes that the string instance has been processed by the SSAStringZero() function or no memory is allocated inside the string instance. If an error is encountered, an error message is stored in the *error_label* buffer.

The SSAStringToChar() function returns a temporary char array buffer for printing purposes.

The SSAStringZero() function frees the memory inside of the String instance, but not the string object itself.

RETURN VALUES The SSAStringInit() and SSAStringCpy() functions return 0 if successful and -1 if error.

ATTRIBUTES See attributes (5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | Unsafe |

SEE ALSO attributes(5)

| | |
|--------------------|--|
| NAME | strccpy, streadd, strcadd, strecpy – copy strings, compressing or expanding escape codes |
| SYNOPSIS | <pre>cc [flag ...] file ... -lgen [library ...] #include <libgen.h> char *strccpy(char *output, const char *input); char *strcadd(char *output, const char *input); char *strecpy(char *output, const char *input, const char *exceptions); char *streadd(char *output, const char *input, const char *exceptions);</pre> |
| DESCRIPTION | <p>strccpy() copies the <i>input</i> string, up to a null byte, to the <i>output</i> string, compressing the C-language escape sequences (for example, \, \\001) to the equivalent character. A null byte is appended to the output. The <i>output</i> argument must point to a space big enough to accommodate the result. If it is as big as the space pointed to by <i>input</i> it is guaranteed to be big enough. strccpy() returns the <i>output</i> argument.</p> <p>strcadd() is identical to strccpy() , except that it returns the pointer to the null byte that terminates the output.</p> <p>strecpy() copies the <i>input</i> string, up to a null byte, to the <i>output</i> string, expanding non-graphic characters to their equivalent C-language escape sequences (for example, \, \\001). The <i>output</i> argument must point to a space big enough to accommodate the result; four times the space pointed to by <i>input</i> is guaranteed to be big enough (each character could become \\ and 3 digits). Characters in the <i>exceptions</i> string are not expanded. The <i>exceptions</i> argument may be zero, meaning all non-graphic characters are expanded. strecpy() returns the <i>output</i> argument.</p> <p>streadd() is identical to strecpy() , except that it returns the pointer to the null byte that terminates the output.</p> |
| EXAMPLES | <p>EXAMPLE 1 Example of expanding and compressing escape codes.</p> <pre>/* expand all but newline and tab */ strecpy(output, input, "\ \t"); /* concatenate and compress several strings */ cp = strcadd(output, input1); cp = strcadd(cp, input2); cp = strcadd(cp, input3);</pre> |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`string(3C)`, `strfind(3GEN)`, `attributes(5)`

NOTES

When compiling multi-thread applications, the `_REENTRANT` flag must be defined on the compile line. This flag should only be used in multi-thread applications.

| NAME | strfind, strstrn, strtrns, str – string manipulations | | | | |
|--------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lgen [library ...] #include <libgen.h> int strfind(const char *as1, const char *as2); char *strstrn(const char *string, const char *tc); char *strtrns(const char *string, const char *old, const char *new, char *result);</pre> | | | | |
| DESCRIPTION | <p>The <code>strfind()</code> function returns the offset of the first occurrence of the second string, <code>as2</code>, if it is a substring of string <code>as1</code>. If the second string is not a substring of the first string <code>strfind()</code> returns <code>-1</code>.</p> <p>The <code>strstrn()</code> function trims characters from a string. It searches from the end of <code>string</code> for the first character that is not contained in <code>tc</code>. If such a character is found, <code>strstrn()</code> returns a pointer to the next character; otherwise, it returns a pointer to <code>string</code>.</p> <p>The <code>strtrns()</code> function transforms <code>string</code> and copies it into <code>result</code>. Any character that appears in <code>old</code> is replaced with the character in the same position in <code>new</code>. The <code>new</code> result is returned.</p> | | | | |
| USAGE | When compiling multithreaded applications, the <code>_REENTRANT</code> flag must be defined on the compile line. This flag should only be used in multithreaded applications. | | | | |
| EXAMPLES | <p>EXAMPLE 1 An example of the <code>strfind()</code> function.</p> <pre>/* find offset to substring "hello" within as1 */ i = strfind(as1, "hello"); /* trim junk from end of string */ s2 = strstrn(s1, ".*#\$\$%"); *s2 = '\\0'; /* transform lower case to upper case */ a1[] = "abcdefghijklmnopqrstuvwxyz"; a2[] = "ABCDEFGHIJKLMNOPQRSTUVWXYZ"; s2 = strtrns(s1, a1, a2, s2);</pre> | | | | |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | <code>string(3C)</code> , <code>attributes(5)</code> | | | | |

NAME tan – tangent function

SYNOPSIS cc [*flag ...*] *file ...* -lm [*library ...*]
#include <math.h>
double **tan**(double *x*);

DESCRIPTION The `tan()` function computes the tangent of its argument *x*, measured in radians.

RETURN VALUES Upon successful completion, `tan()` returns the tangent of *x*.
If *x* is NaN or \pm Inf, NaN is returned.

ERRORS No errors will occur.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO `atan(3M)`, `isnan(3M)`, `attributes(5)`

NAME | tanh – hyperbolic tangent function

SYNOPSIS | cc [*flag ...*] *file ...* -lm [*library ...*]
 | #include <math.h>
 | double **tanh**(double *x*);

DESCRIPTION | The `tanh()` function computes the hyperbolic tangent of *x*.

RETURN VALUES | Upon successful completion, `tanh()` returns the hyperbolic tangent of *x*.
 | If *x* is NaN, NaN is returned.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `atanh(3M)`, `isnan(3M)`, `tan(3M)`, `attributes(5)`

| | |
|----------------------|--|
| NAME | tnfctl_buffer_alloc, tnfctl_buffer_dealloc – allocate or deallocate a buffer for trace data |
| SYNOPSIS | <pre>cc [flag ...] file ... -ltnfctl [library ...] #include <tnf/tnfctl.h> tnfctl_errcode_t tnfctl_buffer_alloc(tnfctl_handle_t *hdl, const char *trace_file_name, size_t trace_buffer_size); tnfctl_buffer_dealloc(tnfctl_handle_t *hdl);</pre> |
| DESCRIPTION | <p>tnfctl_buffer_alloc() allocates a buffer to which trace events are logged. When tracing a process using a tnfctl handle returned by tnfctl_pid_open(3TNF), tnfctl_exec_open(3TNF), tnfctl_indirect_open(3TNF), and tnfctl_internal_open(3TNF), <i>trace_file_name</i> is the name of the trace file to which trace events should be logged. It can be an absolute path specification or a relative path specification. If it is relative, the current working directory of the process that is calling tnfctl_buffer_alloc() is prefixed to <i>trace_file_name</i>. If the named trace file already exists, it is overwritten. For kernel tracing, that is, for a tnfctl handle returned by tnfctl_kernel_open(3TNF), trace events are logged to a trace buffer in memory; therefore, <i>trace_file_name</i> is ignored. Use tnfextract(1) to extract a kernel buffer into a file.</p> <p><i>trace_buffer_size</i> is the size in bytes of the trace buffer that should be allocated. An error is returned if an attempt is made to allocate a buffer when one already exists. tnfctl_buffer_alloc() affects the trace attributes; use tnfctl_trace_attrs_get(3TNF) to get the latest trace attributes after a buffer is allocated.</p> <p>tnfctl_buffer_dealloc() is used to deallocate a kernel trace buffer that is no longer needed. <i>hdl</i> must be a kernel handle, returned by tnfctl_kernel_open(3TNF). A process's trace file cannot be deallocated using tnfctl_buffer_dealloc(). Instead, once the trace file is no longer needed for analysis and after the process being traced exits, use rm(1) to remove the trace file. Do not remove the trace file while the process being traced is still alive. tnfctl_buffer_dealloc() affects the trace attributes; use tnfctl_trace_attrs_get(3TNF) to get the latest trace attributes after a buffer is deallocated.</p> <p>For a complete discussion of tnf tracing, see tracing(3TNF).</p> |
| RETURN VALUES | tnfctl_buffer_alloc() and tnfctl_buffer_dealloc() return TNFCTL_ERR_NONE upon success. |
| ERRORS | The following error codes apply to tnfctl_buffer_alloc(): TNFCTL_ERR_BUFEXISTS A buffer already exists. |

| | |
|---|--|
| TNFCTL_ERR_ACCES | Permission denied; could not create a trace file. |
| TNFCTL_ERR_SIZETOOSMALL | The <i>trace_buffer_size</i> requested is smaller than the minimum trace buffer size needed. Use <i>trace_min_size</i> of trace attributes in <code>tnfctl_trace_attrs_get(3TNF)</code> to determine the minimum size of the buffer. |
| TNFCTL_ERR_SIZETOOBIG | The requested trace file size is too big. |
| TNFCTL_ERR_BADARG | <i>trace_file_name</i> is NULL or the absolute path name is longer than MAXPATHLEN. |
| TNFCTL_ERR_ALLOCFAIL | A memory allocation failure occurred. |
| TNFCTL_ERR_INTERNAL | An internal error occurred. |
| The following error codes apply to <code>tnfctl_buffer_dealloc()</code> : | |
| TNFCTL_ERR_BADARG | <i>hndl</i> is not a kernel handle. |
| TNFCTL_ERR_NOBUF | No buffer exists to deallocate. |
| TNFCTL_ERR_BADDEALLOC | Cannot deallocate a trace buffer unless tracing is stopped. Use <code>tnfctl_trace_state_set(3TNF)</code> to stop tracing. |
| TNFCTL_ERR_INTERNAL | An internal error occurred. |

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO

`prex(1)`, `rm(1)`, `tnfextract(1)`, `TNF_PROBE(3TNF)`, `libtnfctl(3TNF)`, `tnfctl_exec_open(3TNF)`, `tnfctl_indirect_open(3TNF)`, `tnfctl_internal_open(3TNF)`, `tnfctl_kernel_open(3TNF)`, `tnfctl_pid_open(3TNF)`, `tnfctl_trace_attrs_get(3TNF)`, `tracing(3TNF)`, `attributes(5)`

| | | | | | | | | | |
|----------------------|--|---------------------|---|---------------------|-----------------------------|--------------------|--|---------------------|---|
| NAME | tnfctl_close – close a tnfctl handle | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -ltnfctl [library ...] #include <tnf/tnfctl.h></pre> | | | | | | | | |
| DESCRIPTION | <p>tnfctl_errcode_t tnfctl_close(tnfctl_handle_t *hndl, tnfctl_targ_op_t action);</p> <p>tnfctl_close() is used to close a tnfctl handle and to free up the memory associated with the handle. When the handle is closed, the tracing state and the states of the probes are not changed. tnfctl_close() can be used to close handles in any mode, that is, whether they were created by tnfctl_internal_open(3TNF), tnfctl_pid_open(3TNF), tnfctl_exec_open(3TNF), tnfctl_indirect_open(3TNF), or tnfctl_kernel_open(3TNF).</p> <p>The <i>action</i> argument is only used in direct mode, that is, if <i>hndl</i> was created by tnfctl_exec_open(3TNF) or tnfctl_pid_open(3TNF). In direct mode, <i>action</i> specifies whether the process will proceed, be killed, or remain suspended. <i>action</i> may have the following values:</p> <table border="0"> <tr> <td style="vertical-align: top;">TNFCTL_TARG_DEFAULT</td> <td>Kills the target process if <i>hndl</i> was created with tnfctl_exec_open(3TNF), but lets it continue if it was created with tnfctl_pid_open(3TNF).</td> </tr> <tr> <td style="vertical-align: top;">TNFCTL_TARG_KILL</td> <td>Kills the target process.</td> </tr> <tr> <td style="vertical-align: top;">TNFCTL_TARG_RESUME</td> <td>Allows the target process to continue.</td> </tr> <tr> <td style="vertical-align: top;">TNFCTL_TARG_SUSPEND</td> <td>Leaves the target process suspended. This is not a job control suspend. It is possible to attach to the process again with a debugger or with the tnfctl_pid_open(3TNF) interface. The target process can also be continued with prun(1).</td> </tr> </table> | TNFCTL_TARG_DEFAULT | Kills the target process if <i>hndl</i> was created with tnfctl_exec_open(3TNF), but lets it continue if it was created with tnfctl_pid_open(3TNF). | TNFCTL_TARG_KILL | Kills the target process. | TNFCTL_TARG_RESUME | Allows the target process to continue. | TNFCTL_TARG_SUSPEND | Leaves the target process suspended. This is not a job control suspend. It is possible to attach to the process again with a debugger or with the tnfctl_pid_open(3TNF) interface. The target process can also be continued with prun(1). |
| TNFCTL_TARG_DEFAULT | Kills the target process if <i>hndl</i> was created with tnfctl_exec_open(3TNF), but lets it continue if it was created with tnfctl_pid_open(3TNF). | | | | | | | | |
| TNFCTL_TARG_KILL | Kills the target process. | | | | | | | | |
| TNFCTL_TARG_RESUME | Allows the target process to continue. | | | | | | | | |
| TNFCTL_TARG_SUSPEND | Leaves the target process suspended. This is not a job control suspend. It is possible to attach to the process again with a debugger or with the tnfctl_pid_open(3TNF) interface. The target process can also be continued with prun(1). | | | | | | | | |
| RETURN VALUES | tnfctl_close() returns TNFCTL_ERR_NONE upon success. | | | | | | | | |
| ERRORS | <p>The following error codes apply to tnfctl_close():</p> <table border="0"> <tr> <td style="vertical-align: top;">TNFCTL_ERR_BADARG</td> <td>A bad argument was sent in <i>action</i>.</td> </tr> <tr> <td style="vertical-align: top;">TNFCTL_ERR_INTERNAL</td> <td>An internal error occurred.</td> </tr> </table> | TNFCTL_ERR_BADARG | A bad argument was sent in <i>action</i> . | TNFCTL_ERR_INTERNAL | An internal error occurred. | | | | |
| TNFCTL_ERR_BADARG | A bad argument was sent in <i>action</i> . | | | | | | | | |
| TNFCTL_ERR_INTERNAL | An internal error occurred. | | | | | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | | | | | |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO

prex(1), prun(1), TNF_PROBE(3TNF), libtnfctl(3TNF),
tnfctl_exec_open(3TNF), tnfctl_indirect_open(3TNF),
tnfctl_kernel_open(3TNF), tnfctl_pid_open(3TNF), tracing(3TNF),
attributes(5)

| | |
|--------------------|--|
| NAME | tnfctl_indirect_open, tnfctl_check_libs – control probes of another process where caller provides /proc functionality |
| SYNOPSIS | <pre>cc [flag ...] file ... -ltnfctl [library ...] #include <tnf/tnfctl.h> tnfctl_errcode_t tnfctl_indirect_open(void *prochandle, tnfctl_ind_config_t *config, tnfctl_handle_t **ret_val); tnfctl_errcode_t tnfctl_check_libs(tnfctl_handle_t *hndl);</pre> |
| DESCRIPTION | <p>The interfaces <code>tnfctl_indirect_open()</code> and <code>tnfctl_check_libs()</code> are used to control probes in another process where the <code>libtnfctl(3TNF)</code> client has already opened <code>proc(4)</code> on the target process. An example of this is when the client is a debugger. Since these clients already use <code>/proc</code> on the target, <code>libtnfctl(3TNF)</code> cannot use <code>/proc</code> directly. Therefore, these clients must provide callback functions that can be used to inspect and to update the target process. The target process must load <code>libtnfprobe.so.1</code> (defined in <code><tnf/tnfctl.h></code> as macro <code>TNFCTL_LIBTNFPROBE</code>).</p> <p>The first argument <code>prochandle</code> is a pointer to an opaque structure that is used in the callback functions that inspect and update the target process. This structure should encapsulate the state that the caller needs to use <code>/proc</code> on the target process (the <code>/proc</code> file descriptor). The second argument, <code>config</code>, is a pointer to</p> <pre>typedef struct tnfctl_ind_config { int (*p_read)(void *prochandle, paddr_t addr, char *buf, size_t size); int (*p_write)(void *prochandle, paddr_t addr, char *buf, size_t size); pid_t (*p_getpid)(void *prochandle); int (*p_obj_iter)(void *prochandle, tnfctl_ind_obj_f *func, void *client_data); } tnfctl_ind_config_t;</pre> <p>The first field <code>p_read</code> is the address of a function that can read <code>size</code> bytes at address <code>addr</code> in the target image into the buffer <code>buf</code>. The function should return 0 upon success.. The second field <code>p_write</code> is the address of a function that can write <code>size</code> bytes at address <code>addr</code> in the target image from the buffer <code>buf</code>. The function should return 0 upon success. The third field <code>p_getpid</code> is the address of a function that should return the process id of the target process (<code>prochandle</code>). The fourth field <code>p_obj_iter</code> is the address of a function that iterates over all load objects and the executable by calling the callback function <code>func</code> with <code>client_data</code>. If <code>func</code> returns 0, <code>p_obj_iter</code> should continue processing link objects. If <code>func</code> returns any other value, <code>p_obj_iter</code> should stop calling the callback function and return that value. <code>p_obj_iter</code> should return 0 if it iterates over all load objects.</p> |

If a failure is returned by any of the functions in *config*, the error is propagated back as `PREX_ERR_INTERNAL` by the `libtnfctl` interface that called it.

The definition of `tnfctl_ind_obj_f` is:

```
typedef int
tnfctl_ind_obj_f(void *prochandle,
  const struct tnfctl_ind_obj_info *obj
  void *client_data);
typedef struct tnfctl_ind_obj_info {
  int    objfd; /* -1 indicates fd not available */
  paddr_t text_base; /* virtual addr of text segment */
  paddr_t data_base; /* virtual addr of data segment */
  const char *objname; /* null-term. pathname to loadobj */
} tnfctl_ind_obj_info_t;
```

objfd should be the file descriptor of the load object or executable. If it is `-1`, then *objname* should be an absolute pathname to the load object or executable. If *objfd* is not closed by `libtnfctl`, it should be closed by the load object iterator function. *text_base* and *data_base* are the addresses where the text and data segments of the load object are mapped in the target process.

Whenever the target process opens or closes a dynamic object, the set of available probes may change. See `dlopen(3DL)` and `dlclose(3DL)`. In indirect mode, call `tnfctl_check_libs()` when such events occur to make `libtnfctl` aware of any changes. In other modes this is unnecessary but harmless. It is also harmless to call `tnfctl_check_libs()` when no such events have occurred.

RETURN VALUES

`tnfctl_indirect_open()` and `tnfctl_check_libs()` return `TNFCTL_ERR_NONE` upon success.

ERRORS

The following error codes apply to `tnfctl_indirect_open()`:

| | |
|---------------------------------------|--|
| <code>TNFCTL_ERR_ALLOCFAIL</code> | A memory allocation failure occurred. |
| <code>TNFCTL_ERR_BUSY</code> | Internal tracing is being used. |
| <code>TNFCTL_ERR_NOLIBTNFPROBE</code> | <code>libtnfprobe.so.1</code> is not loaded in the target process. |
| <code>TNFCTL_ERR_INTERNAL</code> | An internal error occurred. |

The following error codes apply to `tnfctl_check_libs()`:

| | |
|-----------------------------------|---------------------------------------|
| <code>TNFCTL_ERR_ALLOCFAIL</code> | A memory allocation failure occurred. |
| <code>TNFCTL_ERR_INTERNAL</code> | An internal error occurred. |

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO

prex(1), TNF_PROBE(3TNF), dlclose(3DL), dlopen(3DL), libtnfctl(3TNF), tnfctl_probe_enable(3TNF), tnfctl_probe_trace(3TNF), tracing(3TNF), proc(4), attributes(5)

Linker and Libraries Guide

NOTES

tnfctl_indirect_open() should only be called after the dynamic linker has mapped in all the libraries (rtld sync point) and called only after the process is stopped. Indirect process probe control assumes the target process is stopped whenever any libtnfctl interface is used on it. For example, when used for indirect process probe control, tnfctl_probe_enable(3TNF) and tnfctl_probe_trace(3TNF) should be called only for a process that is stopped.

NAME | tnfctl_internal_open – create handle for internal process probe control

SYNOPSIS |

```
cc [ flag ... ] file ... -ltnfctl [ library ... ]
#include <tnf/tnfctl.h>
```

DESCRIPTION |

```
tnfctl_errcode_t tnfctl_internal_open(tnfctl_handle_t **ret_val);
```

tnfctl_internal_open() returns in *ret_val* a pointer to an opaque handle that can be used to control probes in the same process as the caller (internal process probe control). The process must have libtnfprobe.so.1 loaded. Probes in libraries that are brought in by dlopen(3DL) will be visible after the library has been opened. Probes in libraries closed by a dlclose(3DL) will not be visible after the library has been disassociated. See the NOTES section for more details.

RETURN VALUES | tnfctl_internal_open() returns TNFCTL_ERR_NONE upon success.

ERRORS |

| | |
|--------------------------|--|
| TNFCTL_ERR_ALLOCFAIL | A memory allocation failure occurred. |
| TNFCTL_ERR_BUSY | Another client is already tracing this program (internally or externally). |
| TNFCTL_ERR_NOLIBTNFPROBE | libtnfprobe.so.1 is not linked in the target process. |
| TNFCTL_ERR_INTERNAL | An internal error occurred. |

ATTRIBUTES | See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO | ld(1), prex(1), TNF_PROBE(3TNF), dlopen(3DL), dlclose(3DL), libtnfctl(3TNF), tracing(3TNF), attributes(5)

Linker and Libraries Guide

NOTES | libtnfctl interposes on dlopen(3DL) and dlclose(3DL) in order to be notified of libraries being dynamically opened and closed. This interposition is necessary for internal process probe control to update its list of probes. In these interposition functions, a lock is acquired to synchronize on traversal of the library list maintained by the runtime linker. To avoid deadlocking on this lock, tnfctl_internal_open() should not be called from within the init section of a library that can be opened by dlopen(3DL).

Since interposition does not work as expected when a library is opened dynamically, tnfctl_internal_open() should not be used if the client

opened `libtnfctl` through `dlopen(3DL)`. In this case, the client program should be built with a static dependency on `libtnfctl`. Also, if the client program is explicitly linking in `-ldl`, it should link `-ltnfctl` before `-ldl`.

Probes in filtered libraries (see `ld(1)`) will not be seen because the filtee (backing library) is loaded lazily on the first symbol reference and not at process startup or `dlopen(3DL)` time. A workaround is to call `tnfctl_check_libs(3TNF)` once the caller is sure that the filtee has been loaded.

| NAME | tnfctl_kernel_open – create handle for kernel probe control | | | | | | | | | | |
|-------------------------|--|------------------|--|-----------------|---|----------------------|---------------------------|-------------------------|------------------------|---------------------|------------------------------|
| SYNOPSIS | <pre>cc [flag ...] file ... -ltnfctl [library ...] #include <tnf/tnfctl.h></pre> | | | | | | | | | | |
| DESCRIPTION | <pre>tnfctl_errcode_t tnfctl_kernel_open(tnfctl_handle_t **ret_val);</pre> <p>tnfctl_kernel_open() starts a kernel tracing session and returns in ret_val an opaque handle that can be used to control tracing and probes in the kernel. Only one kernel tracing session is possible at a time on a given machine. An error code of TNFCTL_ERR_BUSY is returned if there is another process using kernel tracing. Use the command</p> <pre>fuser -f /dev/tnfctl</pre> <p>to print the process id of the process currently using kernel tracing. Only a superuser may use tnfctl_kernel_open(). An error code of TNFCTL_ERR_ACCES is returned if the caller does not have the necessary privileges.</p> | | | | | | | | | | |
| RETURN VALUES | tnfctl_kernel_open returns TNFCTL_ERR_NONE upon success. | | | | | | | | | | |
| ERRORS | <table border="0"> <tr> <td>TNFCTL_ERR_ACCES</td> <td>Permission denied. Superuser privileges are needed for kernel tracing.</td> </tr> <tr> <td>TNFCTL_ERR_BUSY</td> <td>Another client is currently using kernel tracing.</td> </tr> <tr> <td>TNFCTL_ERR_ALLOCFAIL</td> <td>Memory allocation failed.</td> </tr> <tr> <td>TNFCTL_ERR_FILENOTFOUND</td> <td>/dev/tnfctl not found.</td> </tr> <tr> <td>TNFCTL_ERR_INTERNAL</td> <td>Some other failure occurred.</td> </tr> </table> | TNFCTL_ERR_ACCES | Permission denied. Superuser privileges are needed for kernel tracing. | TNFCTL_ERR_BUSY | Another client is currently using kernel tracing. | TNFCTL_ERR_ALLOCFAIL | Memory allocation failed. | TNFCTL_ERR_FILENOTFOUND | /dev/tnfctl not found. | TNFCTL_ERR_INTERNAL | Some other failure occurred. |
| TNFCTL_ERR_ACCES | Permission denied. Superuser privileges are needed for kernel tracing. | | | | | | | | | | |
| TNFCTL_ERR_BUSY | Another client is currently using kernel tracing. | | | | | | | | | | |
| TNFCTL_ERR_ALLOCFAIL | Memory allocation failed. | | | | | | | | | | |
| TNFCTL_ERR_FILENOTFOUND | /dev/tnfctl not found. | | | | | | | | | | |
| TNFCTL_ERR_INTERNAL | Some other failure occurred. | | | | | | | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | | | | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWtnfc</td> </tr> <tr> <td>MT Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | Availability | SUNWtnfc | MT Level | MT-Safe | | | | |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | | | | | | | |
| Availability | SUNWtnfc | | | | | | | | | | |
| MT Level | MT-Safe | | | | | | | | | | |
| SEE ALSO | prex(1), fuser(1M), TNF_PROBE(3TNF), libtnfctl(3TNF), tracing(3TNF), tnf_kernel_probes (4), attributes(5) | | | | | | | | | | |

| | |
|--------------------|---|
| NAME | tnfctl_pid_open, tnfctl_exec_open, tnfctl_continue – interfaces for direct probe and process control for another process |
| SYNOPSIS | <pre>cc [flag ...] file ... -ltnfctl [library ...] #include <tnf/tnfctl.h> tnfctl_errcode_t tnfctl_pid_open(pid_t pid, tnfctl_handle_t **ret_val); tnfctl_errcode_t tnfctl_exec_open(const char *pgm_name, char * const *argv, char * const *envp, const char *libtnfprobe_path, const char *ld_preload, tnfctl_handle_t **ret_val); tnfctl_errcode_t tnfctl_continue(tnfctl_handle_t *hndl, tnfctl_event_t *evt, tnfctl_handle_t **child_hndl);</pre> |
| DESCRIPTION | <p>tnfctl_pid_open(), tnfctl_exec_open(), and tnfctl_continue() are the interfaces used to create handles to control probes in another process (direct process probe control). Either tnfctl_pid_open() or tnfctl_exec_open() will return a handle in <i>ret_val</i> that can be used for probe control. On return of these calls, the process is stopped. tnfctl_continue() allows the process specified by <i>hndl</i> to continue execution.</p> <p>tnfctl_pid_open() attaches to a running process with process id of <i>pid</i>. The process is stopped on return of this call. tnfctl_pid_open() returns an error message if <i>pid</i> is the same as the calling process. See tnfctl_internal_open(3TNF) for information on internal process probe control. A pointer to an opaque handle is returned in <i>ret_val</i>, which can be used to control the process and the probes in the process. The target process must have libtnfprobe.so.1 (defined in <tnf/tnfctl.h> as macro TNFCTL_LIBTNFPROBE) linked in for probe control to work.</p> <p>tnfctl_exec_open() is used to exec(2) a program and obtain a probe control handle. For probe control to work, the process image to be exec'd must load libtnfprobe.so.1. The interface tnfctl_exec_open() makes it simple for the library to be loaded at process start up time. <i>pgm_name</i> is the command to exec. If <i>pgm_name</i> is not an absolute path, then the \$PATH environment variable is used to find the <i>pgm_name</i>. <i>argv</i> is a null-terminated argument pointer, that is, it is a null-terminated array of pointers to null-terminated strings. These strings constitute the argument list available to the new process image. <i>argv</i> must have at least one member, and it should point to a string that is the same as <i>pgm_name</i>. See execve(2). <i>libtnfprobe_path</i> is an optional argument, and if set, it should be the path to the directory that contains libtnfprobe.so.1. There is no need for a trailing "/" in this argument. This argument is useful if libtnfprobe.so.1 is not installed in /usr/lib. <i>ld_preload</i> is a space-separated list of libraries to preload into the target program. This string should follow the syntax guidelines of the LD_PRELOAD environment variable. See ld.so.1(1). The following illustrates how strings are concatenated to form the LD_PRELOAD environment variable in the new process image:</p> |

```
<current value of $LD_PRELOAD> + <space> +
libtnfprobe_path + "/libtnfprobe.so.1" +<space> +
ld_preload
```

This option is useful for preloading interposition libraries that have probes in them.

envp is an optional argument, and if set, it is used for the environment of the target program. It is a null-terminated array of pointers to null-terminated strings. These strings constitute the environment of the new process image. See `execve(2)`. If *envp* is set, it overrides *ld_preload*. In this case, it is the caller's responsibility to ensure that `libtnfprobe.so.1` is loaded into the target program. If *envp* is not set, the new process image inherits the environment of the calling process, except for `LD_PRELOAD`.

ret_val is the return argument which is the handle that can be used to control the process and the probes within the process. Upon return, the process is stopped before any user code, including `.init` sections, has been executed.

`tnfctl_continue()` is a blocking call and lets the target process referenced by *hdl* continue running. It can only be used on handles returned by `tnfctl_pid_open()` and `tnfctl_exec_open()` (direct process probe control). It returns when the target stops; the reason that the process stopped is returned in *evt*. This call is interruptible by signals. If it is interrupted, the process is stopped, and `TNFCTL_EVENT_EINTR` is returned in *evt*. The client of this library will have to decide which signal implies a stop to the target and catch that signal. Since a signal interrupts `tnfctl_continue()`, it will return, and the caller can decide whether or not to call `tnfctl_continue()` again.

`tnfctl_continue()` returns with an event of `TNFCTL_EVENT_DLOPEN`, `TNFCTL_EVENT_DLCLOSE`, `TNFCTL_EVENT_EXEC`, `TNFCTL_EVENT_FORK`, `TNFCTL_EVENT_EXIT`, or `TNFCTL_EVENT_TARGGONE`, respectively, when the target program does a `dlopen(3DL)`, `dlclose(3DL)`, any flavor of `exec(2)`, `fork(2)` (or `fork1(2)`), `exit(2)`, or terminates unexpectedly. If the target program did an `exec(2)`, then the client needs to call `tnfctl_close(3TNF)` on the current handle leaving the target resumed, suspended, or killed (second argument to `tnfctl_close(3TNF)`). No other `libtnfctl` interface call can be used on the existing handle. If the client wants to control the `exec`'ed image, it should leave the old handle suspended, and use `tnfctl_pid_open()` to reattach to the same process. This new handle can then be used to control the `exec`'ed image. See `EXAMPLES` below for sample code. If the target process did a `fork(2)` or `fork1(2)`, and if control of the child process is not needed, then *child_hdl* should be `NULL`. If control of the child process is needed, then *child_hdl* should be set. If it is set, a pointer to a handle that can be used to

control the child process is returned in *child_hdl*. The child process is stopped at the end of the `fork()` system call. See `EXAMPLES` for an example of this event.

RETURN VALUES

`tnfctl_pid_open()`, `tnfctl_exec_open()`, and `tnfctl_continue()` return `TNFCTL_ERR_NONE` upon success.

ERRORS

The following error codes apply to `tnfctl_pid_open()`:

| | |
|---------------------------------------|---|
| <code>TNFCTL_ERR_BADARG</code> | The <i>pid</i> specified is the same process. Use <code>tnfctl_internal_open(3TNF)</code> instead. |
| <code>TNFCTL_ERR_ACCES</code> | Permission denied. No privilege to connect to a <code>setuid</code> process. |
| <code>TNFCTL_ERR_ALLOCFAIL</code> | A memory allocation failure occurred. |
| <code>TNFCTL_ERR_BUSY</code> | Another client is already using <code>/proc</code> to control this process or internal tracing is being used. |
| <code>TNFCTL_ERR_NOTDYNAMIC</code> | The process is not a dynamic executable. |
| <code>TNFCTL_ERR_NOPROCESS</code> | No such target process exists. |
| <code>TNFCTL_ERR_NOLIBTNFPROBE</code> | <code>libtnfprobe.so.1</code> is not linked in the target process. |
| <code>TNFCTL_ERR_INTERNAL</code> | An internal error occurred. |

The following error codes apply to `tnfctl_exec_open()`:

| | |
|---------------------------------------|--|
| <code>TNFCTL_ERR_ACCES</code> | Permission denied. |
| <code>TNFCTL_ERR_ALLOCFAIL</code> | A memory allocation failure occurred. |
| <code>TNFCTL_ERR_NOTDYNAMIC</code> | The target is not a dynamic executable. |
| <code>TNFCTL_ERR_NOLIBTNFPROBE</code> | <code>libtnfprobe.so.1</code> is not linked in the target process. |
| <code>TNFCTL_ERR_FILENOTFOUND</code> | The program is not found. |
| <code>TNFCTL_ERR_INTERNAL</code> | An internal error occurred. |

The following error codes apply to `tnfctl_continue()`:

| | |
|----------------------------------|--|
| <code>TNFCTL_ERR_BADARG</code> | Bad input argument. <i>hdl</i> is not a direct process probe control handle. |
| <code>TNFCTL_ERR_INTERNAL</code> | An internal error occurred. |

TNFCTL_ERR_NOPROCESS No such target process exists.

EXAMPLES

EXAMPLE 1 Using `tnfctl_pid_open()`

These examples do not include any error-handling code. Only the initial example includes the declaration of the variables that are used in all of the examples.

The following example shows how to preload `libtnfprobe.so.1` from the normal location and inherit the parent's environment.

```
const char *pgm;
char * const *argv;
tnfctl_handle_t *hndl, *new_hndl, *child_hndl;
tnfctl_errcode_t err;
char * const *envp;
extern char **environ;
tnfctl_event_t evt;
int pid;

/* assuming argv has been allocated */
argv[0] = pgm;
/* set up rest of argument vector here */
err = tnfctl_exec_open(pgm, argv, NULL, NULL, NULL, &hndl);
```

This example shows how to preload two user-supplied libraries `libc_probe.so.1` and `libthread_probe.so.1`. They interpose on the corresponding `libc.so` and `libthread.so` interfaces and have probes for function entry and exit. `libtnfprobe.so.1` is preloaded from the normal location and the parent's environment is inherited.

```
/* assuming argv has been allocated */
argv[0] = pgm;
/* set up rest of argument vector here */
err = tnfctl_exec_open(pgm, argv, NULL, NULL,
    "libc_probe.so.1 libthread_probe.so.1", &hndl);
```

This example preloads an interposition library `libc_probe.so.1`, and specifies a different location from which to preload `libtnfprobe.so.1`.

```
/* assuming argv has been allocated */
argv[0] = pgm;
/* set up rest of argument vector here */
err = tnfctl_exec_open(pgm, argv, NULL, "/opt/SUNWXXX/lib",
    "libc_probe.so.1", &hndl);
```

To set up the environment explicitly for probe control to work, the target process must link `libtnfprobe.so.1`. If using `envp`, it is the caller's responsibility to do so.

```

/* assuming argv has been allocated */
argv[0] = pgm;
/* set up rest of argument vector here */
/* envpnr set up to caller's needs */
err = tnfctl_exec_open(pgm, argv, envpnr, NULL, NULL, &hndl);

```

Use this example to resume a process that does an `exec(2)` without controlling it.

```

err = tnfctl_continue(hndl, &evt, NULL);
switch (evt) {
case TNFCTL_EVENT_EXEC:
/* let target process continue without control */
err = tnfctl_close(hndl, TNFCTL_TARG_RESUME);
...
break;
}

```

Alternatively, use the next example to control a process that does an `exec(2)`.

```

/*
 * assume the pid variable has been set by calling
 * tnfctl_trace_attrs_get()
 */
err = tnfctl_continue(hndl, &evt, NULL);
switch (evt) {
case TNFCTL_EVENT_EXEC:
/* suspend the target process */
err = tnfctl_close(hndl, TNFCTL_TARG_SUSPEND);
/* re-open the exec'ed image */
err = tnfctl_pid_open(pid, &new_hndl);
/* new_hndl now controls the exec'ed image */
...
break;
}

```

To let `fork`'ed children continue without control, use `NULL` as the last argument to `tnfctl_continue()`.

```
err = tnfctl_continue(hndl, &evt, NULL);
```

The next example is how to control child processes that `fork(2)` or `fork1(2)` create.

```

err = tnfctl_continue(hndl, &evt, &child_hndl);
switch (evt) {
case TNFCTL_EVENT_FORK:
/* spawn a new thread or process to control child_hndl */
...
break;
}

```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO

`ld(1)`, `prex(1)`, `proc(1)`, `exec(2)`, `execve(2)`, `exit(2)`, `fork(2)`, `TNF_PROBE(3TNF)`, `dlclose(3DL)`, `dlopen(3DL)`, `libtnfctl(3TNF)`, `tnfctl_close(3TNF)`, `tnfctl_internal_open(3TNF)`, `tracing(3TNF)`, `attributes(5)`

Linker and Libraries Guide

NOTES

After a `tnfctl_continue()` returns, a client should use `tnfctl_trace_attrs_get(3TNF)` to check the `trace_buf_state` member of the trace attributes and make sure that there is no internal error in the target.

| | |
|----------------------|--|
| NAME | tnfctl_probe_apply, tnfctl_probe_apply_ids – iterate over probes |
| SYNOPSIS | <pre>cc [flag ...] file ... -ltnfctl [library ...] #include <tnf/tnfctl.h> tnfctl_errcode_t tnfctl_probe_apply(tnfctl_handle_t *hdl, tnfctl_probe_op_t probe_op, void *clientdata); tnfctl_errcode_t tnfctl_probe_apply_ids(tnfctl_handle_t *hdl, ulong_t probe_count, ulong_t *probe_ids, tnfctl_probe_op_t probe_op, void *clientdata);</pre> |
| DESCRIPTION | <p>tnfctl_probe_apply() is used to iterate over the probes controlled by <i>hdl</i>. For every probe, the <i>probe_op</i> function is called:</p> <pre>typedef tnfctl_errcode_t (*tnfctl_probe_op_t)(tnfctl_handle_t *hdl, tnfctl_probe_t *probe_hdl, void *clientdata);</pre> <p>Several predefined functions are available for use as <i>probe_op</i>. These functions are described in tnfctl_probe_state_get(3TNF).</p> <p>The <i>clientdata</i> supplied in tnfctl_probe_apply() is passed in as the last argument of <i>probe_op</i>. The <i>probe_hdl</i> in the probe operation function can be used to query or change the state of the probe. See tnfctl_probe_state_get(3TNF). The <i>probe_op</i> function should return TNFCTL_ERR_NONE upon success. It can also return an error code, which will cause tnfctl_probe_apply() to stop processing the rest of the probes and return with the same error code. Note that there are five (5) error codes reserved that the client can use for its own semantics. See ERRORS.</p> <p>The lifetime of <i>probe_hdl</i> is the same as the lifetime of <i>hdl</i>. It is good until <i>hdl</i> is closed by tnfctl_close(3TNF). Do not confuse a <i>probe_hdl</i> with <i>hdl</i>. The <i>probe_hdl</i> refers to a particular probe, while <i>hdl</i> refers to a process or the kernel. If <i>probe_hdl</i> is used in another libtnfctl(3TNF) interface, and it references a probe in a library that has been dynamically closed (see dlclose(3DL)), then the error code TNFCTL_ERR_INVALIDPROBE will be returned by that interface.</p> <p>tnfctl_probe_apply_ids() is very similar to tnfctl_probe_apply(). The difference is that <i>probe_op</i> is called only for probes that match a probe id specified in the array of integers referenced by <i>probe_ids</i>. The number of probe ids in the array should be specified in <i>probe_count</i>. Use tnfctl_probe_state_get() to get the <i>probe_id</i> that corresponds to the <i>probe_hdl</i>.</p> |
| RETURN VALUES | tnfctl_probe_apply() and tnfctl_probe_apply_ids() return TNFCTL_ERR_NONE upon success. |

ERRORS

The following errors apply to both `tnfctl_probe_apply()` and `tnfctl_probe_apply_ids()`:

| | |
|----------------------------------|-------------------------------|
| <code>TNFCTL_ERR_INTERNAL</code> | An internal error occurred. |
| <code>TNFCTL_ERR_USR1</code> | Error code reserved for user. |
| <code>TNFCTL_ERR_USR2</code> | Error code reserved for user. |
| <code>TNFCTL_ERR_USR3</code> | Error code reserved for user. |
| <code>TNFCTL_ERR_USR4</code> | Error code reserved for user. |
| <code>TNFCTL_ERR_USR5</code> | Error code reserved for user. |

`tnfctl_probe_apply()` and `tnfctl_probe_apply_ids()` also return any error returned by the callback function `probe_op`.

The following errors apply only to `tnfctl_probe_apply_ids()`:

| | |
|--------------------------------------|---|
| <code>TNFCTL_ERR_INVALIDPROBE</code> | The probe handle is no longer valid. For example, the probe is in a library that has been closed by <code>dlclose(3DL)</code> . |
|--------------------------------------|---|

EXAMPLES**EXAMPLE 1** Enabling Probes

To enable all probes:

```
tnfctl_probe_apply(hndl, tnfctl_probe_enable, NULL);
```

EXAMPLE 2 Disabling Probes

To disable the probes that match a certain pattern in the probe attribute string:

```
/* To disable all probes that contain the string "vm" */
tnfctl_probe_apply(hndl, select_disable, "vm");
static tnfctl_errcode_t
select_disable(tnfctl_handle_t *hndl, tnfctl_probe_t *probe_hndl,
void *client_data)
{
    char *pattern = client_data;
    tnfctl_probe_state_t probe_state;
    tnfctl_probe_state_get(hndl, probe_hndl, &probe_state);
    if (strstr(probe_state.attr_string, pattern)) {
        tnfctl_probe_disable(hndl, probe_hndl, NULL);
    }
}
```

Note that these examples do not have any error handling code.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT-Level | MT-Safe |

SEE ALSO

prex(1) , TNF_PROBE(3TNF) , dlclose(3DL) ,
dlopen(3DL) , libtnfctl(3TNF) , tnfctl_close(3TNF)
, tnfctl_probe_state_get(3TNF) , tracing(3TNF) ,
tnf_kernel_probes(4) , attributes(5)

Linker and Libraries Guide

NAME | tnfctl_probe_state_get, tnfctl_probe_enable, tnfctl_probe_disable, tnfctl_probe_trace, tnfctl_probe_untrace, tnfctl_probe_connect, tnfctl_probe_disconnect_all – interfaces to query and to change the state of a probe

SYNOPSIS

```
cc [ flag ... ] file ... -ltnfctl [ library ... ]
#include <tnf/tnfctl.h>
tnfctl_errcode_t tnfctl_probe_state_get(tnfctl_handle_t *hdl, tnfctl_probe_t
*probe_hdl, tnfctl_probe_state_t *state);

tnfctl_errcode_t tnfctl_probe_enable(tnfctl_handle_t *hdl, tnfctl_probe_t
*probe_hdl, void *ignored);

tnfctl_errcode_t tnfctl_probe_disable(tnfctl_handle_t *hdl, tnfctl_probe_t
*probe_hdl, void *ignored);

tnfctl_errcode_t tnfctl_probe_trace(tnfctl_handle_t *hdl, tnfctl_probe_t *probe_hdl,
void *ignored);

tnfctl_errcode_t tnfctl_probe_untrace(tnfctl_handle_t *hdl, tnfctl_probe_t
*probe_hdl, void *ignored);

tnfctl_errcode_t tnfctl_probe_disconnect_all(tnfctl_handle_t *hdl, tnfctl_probe_t
*probe_hdl, void *ignored);

tnfctl_errcode_t tnfctl_probe_connect(tnfctl_handle_t *hdl, tnfctl_probe_t
*probe_hdl, const char *lib_base_name, const char *func_name);
```

DESCRIPTION

tnfctl_probe_state_get() returns the state of the probe specified by *probe_hdl* in the process or kernel specified by *hdl*. The user will pass these in to an apply iterator. The caller must also allocate *state* and pass in a pointer to it. The semantics of the individual members of *state* are:

| | |
|-------------|--|
| id | The unique integer assigned to this probe. This number does not change over the lifetime of this probe. A <i>probe_hdl</i> can be obtained by using the calls <code>tnfctl_apply()</code> , <code>tnfctl_apply_ids()</code> , or <code>tnfctl_register_funcs()</code> . |
| attr_string | A string that consists of <i>attribute value</i> pairs separated by semicolons. For the syntax of this string, see the syntax of the <code>detail</code> argument of the <code>TNF_PROBE(3TNF)</code> macro. The attributes <i>name</i> , <i>slots</i> , <i>keys</i> , <i>file</i> , and <i>line</i> are defined for every probe. Additional user-defined attributes can be added by using the <i>detail</i> argument of the <code>TNF_PROBE(3TNF)</code> macro. An example of <i>attr_string</i> follows: |

| | |
|------------------------|--|
| | <code>"name pageout;slots vnode pages_pageout ; keys vm pageio io;file vm.c;line 25;"</code> |
| enabled | B_TRUE if the probe is enabled, or B_FALSE if the probe is disabled. Probes are disabled by default. Use <code>tnfctl_probe_enable()</code> or <code>tnfctl_probe_disable()</code> to change this state. |
| traced | B_TRUE if the probe is traced, or B_FALSE if the probe is not traced. Probes in user processes are traced by default. Kernel probes are untraced by default. Use <code>tnfctl_probe_trace()</code> or <code>tnfctl_probe_untrace()</code> to change this state. |
| new_probe | B_TRUE if this is a new probe brought in since the last change in libraries. See <code>dlopen(3DL)</code> or <code>dlclose(3DL)</code> . Otherwise, the value of <code>new_probe</code> will be B_FALSE. This field is not meaningful for kernel probe control. |
| obj_name | The name of the shared object or executable in which the probe is located. This string can be freed, so the client should make a copy of the string if it needs to be saved for use by other <code>libtnfctl</code> interfaces. In kernel mode, this string is always NULL. |
| func_names | A null-terminated array of pointers to strings that contain the names of functions connected to this probe. Whenever an enabled probe is encountered at runtime, these functions are executed. This array also will be freed by the library when the state of the probe changes. Use <code>tnfctl_probe_connect()</code> or <code>tnfctl_probe_disconnect_all()</code> to change this state. |
| func_addrs | A null-terminated array of pointers to addresses of functions in the target image connected to this probe. This array also will be freed by the library when the state of the probe changes. |
| client_registered_data | Data that was registered by the client for this probe by the creator function in <code>tnfctl_register_funcs(3TNF)</code> . |
| | <code>tnfctl_probe_enable()</code> , <code>tnfctl_probe_disable()</code> , <code>tnfctl_probe_trace()</code> , <code>tnfctl_probe_untrace()</code> , and |

`tnfctl_probe_disconnect_all()` ignore the last argument. This convenient feature permits these functions to be used in the *probe_op* field of `tnfctl_probe_apply(3TNF)` and `tnfctl_probe_apply_ids(3TNF)`. `tnfctl_probe_enable()` enables the probe specified by *probe_hdl*. This is the master switch on a probe. A probe does not perform any action until it is enabled.

`tnfctl_probe_disable()` disables the probe specified by *probe_hdl*.

`tnfctl_probe_trace()` turns on tracing for the probe specified by *probe_hdl*. Probes emit a trace record only if the probe is traced.

`tnfctl_probe_untrace()` turns off tracing for the probe specified by *probe_hdl*. This is useful if you want to connect probe functions to a probe without tracing it.

`tnfctl_probe_connect()` connects the function *func_name* which exists in the library *lib_base_name*, to the probe specified by *probe_hdl*. `tnfctl_probe_connect()` returns an error code if used on a kernel `tnfctl` handle. *lib_base_name* is the base name (not a path) of the library. If it is `NULL`, and multiple functions in the target process match *func_name*, one of the matching functions is chosen arbitrarily. A probe function is a function that is in the target's address space and is written to a certain specification. The specification is not currently published.

`tnf_probe_debug()` is one function exported by `libtnfprobe.so.1` and is the debug function that `prex(1)` uses. When the debug function is executed, it prints out the probe arguments and the value of the `sunw%debug` attribute of the probe to `stderr`.

`tnfctl_probe_disconnect_all()` disconnects all probe functions from the probe specified by *probe_hdl*.

Note that no `libtnfctl` call returns a probe handle (`tnfctl_probe_t`), yet each of the routines described here takes a *probe_hdl* as an argument. These routines may be used by passing them to one of the `tnfctl_probe_apply(3TNF)` iterators as the "op" argument. Alternatively, probe handles may be obtained and saved by a user's "op" function, and they can be passed later as the *probe_hdl* argument when using any of the functions described here.

RETURN VALUES

`tnfctl_probe_state_get()`, `tnfctl_probe_enable()`, `tnfctl_probe_disable()`, `tnfctl_probe_trace()`, `tnfctl_probe_untrace()`, `tnfctl_probe_disconnect_all()` and `tnfctl_probe_connect()` return `TNFCTL_ERR_NONE` upon success.

ERRORS

The following error codes apply to `tnfctl_probe_state_get()`:

TNFCTL_ERR_INVALIDPROBE *probe_hdl* is no longer valid. The library that the probe was in could have been dynamically closed by `dlclose(3DL)`.

The following error codes apply to `tnfctl_probe_enable()`, `tnfctl_probe_disable()`, `tnfctl_probe_trace()`, `tnfctl_probe_untrace()`, and `tnfctl_probe_disconnect_all()`

TNFCTL_ERR_INVALIDPROBE *probe_hdl* is no longer valid. The library that the probe was in could have been dynamically closed by `dlclose(3DL)`.

TNFCTL_ERR_BUFBROKEN Cannot do probe operations because tracing is broken in the target.

TNFCTL_ERR_NOBUF Cannot do probe operations until a buffer is allocated. See `tnfctl_buffer_alloc(3TNF)`. This error code does not apply to kernel probe control.

The following error codes apply to `tnfctl_probe_connect()`:

TNFCTL_ERR_INVALIDPROBE *probe_hdl* is no longer valid. The library that the probe was in could have been dynamically closed by `dlclose(3DL)`.

TNFCTL_ERR_BADARG The handle is a kernel handle, or *func_name* could not be found.

TNFCTL_ERR_BUFBROKEN Cannot do probe operations because tracing is broken in the target.

TNFCTL_ERR_NOBUF Cannot do probe operations until a buffer is allocated. See `tnfctl_buffer_alloc(3TNF)`.

ATTRIBUTES

See `attributes(5)` for description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO

`prex(1)`, `TNF_PROBE(3TNF)`, `libtnfctl(3TNF)`, `tnfctl_check_libs(3TNF)`, `tnfctl_continue(3TNF)`,

```
tnfctl_probe_apply(3TNF), tnfctl_probe_apply_ids(3TNF),  
tracing(3TNF), tnf_kernel_probes(4), attributes(5)
```

NAME tnfctl_register_funcs – register callbacks for probe creation and destruction

SYNOPSIS

```
cc [ flag ... ] file ... -ltnfctl [ library ... ]
#include <tnf/tnfctl.h>
```



```
tnfctl_errcode_t tnfctl_register_funcs(tnfctl_handle_t *hndl, void *
(*create_func)(tnfctl_handle_t *, tnfctl_probe_t *), void (*destroy_func)(void *));
```

DESCRIPTION
The function `tnfctl_register_funcs()` is used to store client-specific data on a per-probe basis. It registers a creator and a destructor function with `hndl`, either of which can be NULL. The creator function is called for every probe that currently exists in `hndl`. Every time a new probe is discovered, that is brought in by `dlopen(3DL)`, `create_func` is called.

The return value of the creator function is stored as part of the probe state and can be retrieved by `tnfctl_probe_state_get(3TNF)` in the member field `client_registered_data`.

`destroy_func` is called for every probe handle that is freed. This does not necessarily happen at the time `dlclose(3DL)` frees the shared object. The probe handles are freed only when `hndl` is closed by `tnfctl_close(3TNF)`. If `tnfctl_register_funcs()` is called a second time for the same `hndl`, then the previously registered destructor function is called first for all of the probes.

RETURN VALUES
`tnfctl_register_funcs()` returns `TNFCTL_ERR_NONE` upon success.

ERRORS
`TNFCTL_ERR_INTERNAL` An internal error occurred.

ATTRIBUTES
See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO

```
prex(1), TNF_PROBE(3TNF), dlclose(3DL), dlopen(3DL), libtnfctl(3TNF),
tnfctl_close(3TNF), tnfctl_probe_state_get(3TNF), tracing(3TNF),
tnf_kernel_probes(4), attributes(5)
```


Linker and Libraries Guide

NAME tnfctl_strerror - map a tnfctl error code to a string

SYNOPSIS

```
cc [ flag ... ] file ... -ltnfctl [ library ... ]
#include <tnf/tnfctl.h>
```

```
const char * tnfctl_strerror(tnfctl_errcode_t errcode);
```

DESCRIPTION tnfctl_strerror() maps the error number in *errcode* to an error message string, and it returns a pointer to that string. The returned string should not be overwritten or freed.

ERRORS tnfctl_strerror() returns the string "unknown libtnfctl.so error code" if the error number is not within the legal range.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO prex(1), TNF_PROBE(3TNF), libtnfctl(3TNF), tracing(3TNF), attributes(5)

| | | | | | | | | | |
|--------------------|--|----------|---|-----------------|--|----------------|--|----------------|---|
| NAME | tnfctl_trace_attrs_get – get the trace attributes from a tnfctl handle | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -ltnfctl [library ...] #include <tnf/tnfctl.h></pre> <p>tnfctl_errcode_t tnfctl_trace_attrs_get(tnfctl_handle_t *hdl, tnfctl_trace_attrs_t *attrs);</p> | | | | | | | | |
| DESCRIPTION | <p>tnfctl_trace_attrs_get() returns the trace attributes associated with <i>hdl</i> in <i>attrs</i>. The trace attributes can be changed by some of the other interfaces in libtnfctl(3TNF). It is the client's responsibility to use tnfctl_trace_attrs_get() to get the new trace attributes after use of interfaces that change them. Typically, a client will use tnfctl_trace_attrs_get() after a call to tnfctl_continue(3TNF) in order to make sure that tracing is still working. See the discussion of trace_buf_state that follows.</p> <p>Trace attributes are represented by the struct tnfctl_trace_attrs structure defined in <tnf/tnfctl.h>:</p> <pre>struct tnfctl_trace_attrs { pid_t targ_pid; /* not kernel mode */ const char *trace_file_name; /* not kernel mode */ size_t trace_buf_size; size_t trace_min_size; tnfctl_bufstate_t trace_buf_state; boolean_t trace_state; boolean_t filter_state; /* kernel mode only */ long pad; };</pre> <p>The semantics of the individual members of <i>attrs</i> are:</p> <table border="0"> <tr> <td style="vertical-align: top;">targ_pid</td> <td>The process id of the target process. This is not valid for kernel tracing.</td> </tr> <tr> <td style="vertical-align: top;">trace_file_name</td> <td>The name of the trace file to which the target writes. trace_file_name will be NULL if no trace file exists or if kernel tracing is implemented. This pointer should not be used after calling other libtnfctl interfaces. The client should copy this string if it should be saved for the use of other libtnfctl interfaces.</td> </tr> <tr> <td style="vertical-align: top;">trace_buf_size</td> <td>The size of the trace buffer or file in bytes.</td> </tr> <tr> <td style="vertical-align: top;">trace_min_size</td> <td>The minimum size in bytes of the trace buffer that can be allocated by using the tnfctl_buffer_alloc(3TNF) interface.</td> </tr> </table> | targ_pid | The process id of the target process. This is not valid for kernel tracing. | trace_file_name | The name of the trace file to which the target writes. trace_file_name will be NULL if no trace file exists or if kernel tracing is implemented. This pointer should not be used after calling other libtnfctl interfaces. The client should copy this string if it should be saved for the use of other libtnfctl interfaces. | trace_buf_size | The size of the trace buffer or file in bytes. | trace_min_size | The minimum size in bytes of the trace buffer that can be allocated by using the tnfctl_buffer_alloc(3TNF) interface. |
| targ_pid | The process id of the target process. This is not valid for kernel tracing. | | | | | | | | |
| trace_file_name | The name of the trace file to which the target writes. trace_file_name will be NULL if no trace file exists or if kernel tracing is implemented. This pointer should not be used after calling other libtnfctl interfaces. The client should copy this string if it should be saved for the use of other libtnfctl interfaces. | | | | | | | | |
| trace_buf_size | The size of the trace buffer or file in bytes. | | | | | | | | |
| trace_min_size | The minimum size in bytes of the trace buffer that can be allocated by using the tnfctl_buffer_alloc(3TNF) interface. | | | | | | | | |

| | |
|-----------------|--|
| trace_buf_state | The state of the trace buffer. TNFCTL_BUF_OK indicates that a trace buffer has been allocated. TNFCTL_BUF_NONE indicates that no buffer has been allocated. TNFCTL_BUF_BROKEN indicates that there is an internal error in the target for tracing. The target will continue to run correctly, but no trace records will be written. To fix tracing, restart the process. For kernel tracing, deallocate the existing buffer with <code>tnfctl_buffer_dealloc(3TNF)</code> and allocate a new one with <code>tnfctl_buffer_alloc(3TNF)</code> . |
| trace_state | The global tracing state of the target. Probes that are enabled will not write out data unless this state is on. This state is off by default for the kernel and can be changed by <code>tnfctl_trace_state_set(3TNF)</code> . For a process, this state is on by default and can only be changed by <code>tnf_process_disable(3TNF)</code> and <code>tnf_process_enable(3TNF)</code> . |
| filter_state | The state of process filtering. For kernel probe control, it is possible to select a set of processes for which probes are enabled. See <code>tnfctl_filter_list_get(3TNF)</code> , <code>tnfctl_filter_list_add(3TNF)</code> , and <code>tnfctl_filter_list_delete(3TNF)</code> . No trace output will be written when other processes traverse these probe points. By default process filtering is off, and all processes cause the generation of trace records when they hit an enabled probe. Use <code>tnfctl_filter_state_set(3TNF)</code> to change the filter state. |

RETURN VALUES

`tnfctl_trace_attrs_get()` returns `TNFCTL_ERR_NONE` upon success.

ERRORS

The following error codes apply to `tnfctl_trace_attrs_get()`
`TNFCTL_ERR_INTERNAL` An internal error occurred.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO

prex(1), TNF_PROBE(3TNF), libtnfctl(3TNF),
tnfctl_buffer_alloc(3TNF), tnfctl_continue(3TNF),
tnfctl_filter_list_get (3TNF), tnf_process_disable(3TNF),
tracing(3TNF), attributes(5)

| | |
|----------------------|--|
| NAME | tnfctl_trace_state_set, tnfctl_filter_state_set, tnfctl_filter_list_get, tnfctl_filter_list_add, tnfctl_filter_list_delete – control kernel tracing and process filtering |
| SYNOPSIS | <pre>cc [flag ...] file ... -ltnfctl [library ...] #include <tnf/tnfctl.h> tnfctl_errcode_t tnfctl_trace_state_set(tnfctl_handle_t *hdl, boolean_t trace_state); tnfctl_errcode_t tnfctl_filter_state_set(tnfctl_handle_t *hdl, boolean_t filter_state); tnfctl_errcode_t tnfctl_filter_list_get(tnfctl_handle_t *hdl, pid_t **pid_list, int *pid_count); tnfctl_errcode_t tnfctl_filter_list_add(tnfctl_handle_t *hdl, pid_t pid_to_add); tnfctl_errcode_t tnfctl_filter_list_delete(tnfctl_handle_t *hdl, pid_t pid_to_delete);</pre> |
| DESCRIPTION | <p>The interfaces to control kernel tracing and process filtering are used only with kernel handles, handles created by <code>tnfctl_kernel_open(3TNF)</code>. These interfaces are used to change the tracing and filter states for kernel tracing.</p> <p><code>tnfctl_trace_state_set()</code> sets the kernel global tracing state to "on" if <code>trace_state</code> is <code>B_TRUE</code>, or to "off" if <code>trace_state</code> is <code>B_FALSE</code>. For the kernel, <code>trace_state</code> is off by default. Probes that are enabled will not write out data unless this state is on. Use <code>tnfctl_trace_attrs_get(3TNF)</code> to retrieve the current tracing state.</p> <p><code>tnfctl_filter_state_set()</code> sets the kernel process filtering state to "on" if <code>filter_state</code> is <code>B_TRUE</code>, or to "off" if <code>filter_state</code> is <code>B_FALSE</code>. <code>filter_state</code> is off by default. If it is on, only probe points encountered by processes in the process filter set by <code>tnfctl_filter_list_add()</code> will generate trace points. Use <code>tnfctl_trace_attrs_get(3TNF)</code> to retrieve the current process filtering state.</p> <p><code>tnfctl_filter_list_get()</code> returns the process filter list as an array in <code>pid_list</code>. The count of elements in the process filter list is returned in <code>pid_count</code>. The caller should use <code>free(3C)</code> to free memory allocated for the array <code>pid_list</code>.</p> <p><code>tnfctl_filter_list_add()</code> adds <code>pid_to_add</code> to the process filter list. The process filter list is maintained even when the process filtering state is off, but it has no effect unless the process filtering state is on.</p> <p><code>tnfctl_filter_list_delete()</code> deletes <code>pid_to_delete</code> from the process filter list. It returns an error if the process does not exist or is not in the filter list.</p> |
| RETURN VALUES | The interfaces <code>tnfctl_trace_state_set()</code> , <code>tnfctl_filter_state_set()</code> , <code>tnfctl_filter_list_add()</code> , |

tnfctl_filter_list_delete(), and tnfctl_filter_list_get() return TNFCTL_ERR_NONE upon success.

ERRORS

The following error codes apply to tnfctl_trace_state_set:

- TNFCTL_ERR_BADARG The handle is not a kernel handle.
- TNFCTL_ERR_NOBUF Cannot turn on tracing without a buffer being allocated.
- TNFCTL_ERR_BUFBROKEN Tracing is broken in the target.
- TNFCTL_ERR_INTERNAL An internal error occurred.

The following error codes apply to tnfctl_filter_state_set:

- TNFCTL_ERR_BADARG The handle is not a kernel handle.
- TNFCTL_ERR_INTERNAL An internal error occurred.

The following error codes apply to tnfctl_filter_list_add:

- TNFCTL_ERR_BADARG The handle is not a kernel handle.
- TNFCTL_ERR_NOPROCESS No such process exists.
- TNFCTL_ERR_ALLOCFAIL A memory allocation failure occurred.
- TNFCTL_ERR_INTERNAL An internal error occurred.

The following error codes apply to tnfctl_filter_list_delete:

- TNFCTL_ERR_BADARG The handle is not a kernel handle.
- TNFCTL_ERR_NOPROCESS No such process exists.
- TNFCTL_ERR_INTERNAL An internal error occurred.

The following error codes apply to tnfctl_filter_list_get:

- TNFCTL_ERR_BADARG The handle is not a kernel handle.
- TNFCTL_ERR_ALLOCFAIL A memory allocation failure occurred.
- TNFCTL_ERR_INTERNAL An internal error occurred.

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfc |
| MT Level | MT-Safe |

SEE ALSO

prex(1), TNF_PROBE(3TNF), free(3C), libtnfctl(3TNF), tnfctl_kernel_open(3TNF), tnfctl_trace_attrs_get (3TNF), tracing(3TNF), tnf_kernel_probes(4), attributes(5)

| | |
|--------------------|---|
| NAME | TNF_DECLARE_RECORD, TNF_DEFINE_RECORD_1, TNF_DEFINE_RECORD_2, TNF_DEFINE_RECORD_3, TNF_DEFINE_RECORD_4, TNF_DEFINE_RECORD_5 – TNF type extension interface for probes |
| SYNOPSIS | <pre>cc [flag ...] file ...[-ltnfprobe] [library ...] #include <tnf/probe.h> TNF_DECLARE_RECORD(c_type, tnf_type); TNF_DEFINE_RECORD_1(c_type, tnf_type, tnf_member_type_1, c_member_name_1); TNF_DEFINE_RECORD_2(c_type, tnf_type, tnf_member_type_1, c_member_name_1, tnf_member_type_2, c_member_name_2); TNF_DEFINE_RECORD_3(c_type, tnf_type, tnf_member_type_1, c_member_name_1, tnf_member_type_2, c_member_name_2, tnf_member_type_3, c_member_name_3); TNF_DEFINE_RECORD_4(c_type, tnf_type, tnf_member_type_1, c_member_name_1, tnf_member_type_2, c_member_name_2, tnf_member_type_3, c_member_name_3, tnf_member_type_4, c_member_name_4); TNF_DEFINE_RECORD_5(c_type, tnf_type, tnf_member_type_1, c_member_name_1, tnf_member_type_2, c_member_name_2, tnf_member_type_3, c_member_name_3, tnf_member_type_4, c_member_name_4, tnf_member_type_5, c_member_name_5);</pre> |
| DESCRIPTION | <p>This macro interface is used to extend the TNF (Trace Normal Form) types that can be used in TNF_PROBE(3TNF) .</p> <p>There should be only one TNF_DECLARE_RECORD and one TNF_DEFINE_RECORD per new type being defined. The TNF_DECLARE_RECORD should precede the TNF_DEFINE_RECORD . It can be in a header file that multiple source files share if those source files need to use the <i>tnf_type</i> being defined. The TNF_DEFINE_RECORD should only appear in one of the source files.</p> <p>The TNF_DEFINE_RECORD macro interface defines a function as well as a couple of data structures. Hence, this interface has to be used in a source file (.c or .cc file) at file scope and not inside a function.</p> <p>Note that there is no semicolon after the TNF_DEFINE_RECORD interface. Having one will generate a compiler warning.</p> <p>Compiling with the preprocessor option <code>-DNPROBE</code> (see <code>cc(1B)</code>), or with the preprocessor control statement <code>#define NPROBE</code> ahead of the <code>#include <tnf/probe.h></code> statement, will stop the TNF type extension code from being compiled into the program.</p> <p>c_type <i>c_type</i> must be a C struct type. It is the template from which the new <i>tnf_type</i> is being created. Not all elements of the C struct need be provided in the TNF type being defined.</p> |

tnf_type | *tnf_type* is the name being given to the newly created type. Use of this interface uses the name space prefixed by *tnf_type*. So, if a new type called "xxx_type" is defined by a library, then the library should not use "xxx_type" as a prefix in any other symbols it defines. The policy on managing the type name space is the same as managing any other name space in a library i.e., prefix any new TNF types by the unique prefix that the rest of the symbols in the library use. This would prevent name space collisions when linking multiple libraries that define new TNF types. For example, if a library libpalloc.so uses the prefix "pal" for all symbols it defines, then it should also use the prefix "pal" for all new TNF types being defined.

tnf_member_type_n | *tnf_member_type_n* is the TNF type of the *n* th provided member of the C structure.

tnf_member_name_n | *tnf_member_name_n* is the name of the *n* th provided member of the C structure.

EXAMPLES

EXAMPLE 1 Defining and using a TNF type.

This example shows how a new TNF type is defined and used in a probe. This code is assumed to be part of a fictitious library called "libpalloc.so" which uses the prefix "pal" for all it's symbols.

```
#include <tnf/probe.h>
typedef struct pal_header {
    long    size;
    char *  descriptor;
    struct pal_header *next;
} pal_header_t;
TNF_DECLARE_RECORD(pal_header_t, pal_tnf_header);
TNF_DEFINE_RECORD_2(pal_header_t, pal_tnf_header,
                   tnf_long,    size,
                   tnf_string, descriptor)

/*
 * Note: name space prefixed by pal_tnf_header should not be used by this
 * client anymore.
 */
void
pal_free(pal_header_t *header_p)
{
    int state;
    TNF_PROBE_2(pal_free_start, "palloc pal_free",
               "sunw%debug entering pal_free",
               tnf_long,    state_var, state,
               pal_tnf_header, header_var, header_p);
    . . .
}
```

ATTRIBUTES

See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfd |
| MT-Level | MT-Safe |

SEE ALSO

`prex(1)`, `tnfdump(1)`, `TNF_PROBE(3TNF)`, `tnf_process_disable(3TNF)`, `attributes(5)`

NOTES

It is possible to make a *tnf_type* definition be recursive or mutually recursive e.g. a structure that uses the "next" field to point to itself (a linked list). If such a structure is sent in to a `TNF_PROBE(3TNF)`, then the entire linked list will be logged to the trace file (until the "next" field is `NULL`). But, if the list is circular, it will result in an infinite loop. To break the recursion, either don't include the "next" field in the *tnf_type*, or define the type of the "next" member as `tnf_opaque`.

NAME TNF_PROBE, TNF_PROBE_0, TNF_PROBE_1, TNF_PROBE_2, TNF_PROBE_3, TNF_PROBE_4, TNF_PROBE_5, TNF_PROBE_0_DEBUG, TNF_PROBE_1_DEBUG, TNF_PROBE_2_DEBUG, TNF_PROBE_3_DEBUG, TNF_PROBE_4_DEBUG, TNF_PROBE_5_DEBUG, TNF_DEBUG – probe insertion interface

SYNOPSIS

```
cc [ flag ... ] [-DTNF_DEBUG] file ... [-ltnfprobe] [ library ... ]
#include <tnf/probe.h>
TNF_PROBE_0(name, keys, detail);

TNF_PROBE_1(name, keys, detail, arg_type_1, arg_name_1, arg_value_1);

TNF_PROBE_2(name, keys, detail, arg_type_1, arg_name_1, arg_value_1, arg_type_2,
arg_name_2, arg_value_2);

TNF_PROBE_3(name, keys, detail, arg_type_1, arg_name_1, arg_value_1, arg_type_2,
arg_name_2, arg_value_2, arg_type_3, arg_name_3, arg_value_3);

TNF_PROBE_4(name, keys, detail, arg_type_1, arg_name_1, arg_value_1, arg_type_2,
arg_name_2, arg_value_2, arg_type_3, arg_name_3, arg_value_3, arg_type_4, arg_name_4,
arg_value_4);

TNF_PROBE_5(name, keys, detail, arg_type_1, arg_name_1, arg_value_1, arg_type_2,
arg_name_2, arg_value_2, arg_type_3, arg_name_3, arg_value_3, arg_type_4, arg_name_4,
arg_value_4, arg_type_5, arg_name_5, arg_value_5);

TNF_PROBE_0_DEBUG(name, keys, detail);

TNF_PROBE_1_DEBUG(name, keys, detail, arg_type_1, arg_name_1, arg_value_1);

TNF_PROBE_2_DEBUG(name, keys, detail, arg_type_1, arg_name_1, arg_value_1, arg_type_2,
arg_name_2, arg_value_2);

TNF_PROBE_3_DEBUG(name, keys, detail, arg_type_1, arg_name_1, arg_value_1, arg_type_2,
arg_name_2, arg_value_2, arg_type_3, arg_name_3, arg_value_3);

TNF_PROBE_4_DEBUG(name, keys, detail, arg_type_1, arg_name_1, arg_value_1, arg_type_2,
arg_name_2, arg_value_2, arg_type_3, arg_name_3, arg_value_3, arg_type_4, arg_name_4,
arg_value_4);

TNF_PROBE_5_DEBUG(name, keys, detail, arg_type_1, arg_name_1, arg_value_1, arg_type_2,
arg_name_2, arg_value_2, arg_type_3, arg_name_3, arg_value_3, arg_type_4, arg_name_4,
arg_value_4, arg_type_5, arg_name_5, arg_value_5);
```

DESCRIPTION

This macro interface is used to insert probes into C or C++ code for tracing. See `tracing(3TNF)` for a discussion of the Solaris tracing architecture, including example source code that uses it.

You can place probes anywhere in C and C++ programs including `.init` sections, `.fini` sections, multi-threaded code, shared objects, and shared objects opened by

`dlopen(3DL)`. Use probes to generate trace data for performance analysis or to write debugging output to `stderr`. Probes are controlled at runtime by `prex(1)`.

The trace data is logged to a trace file in Trace Normal Form (TNF). The interface for the user to specify the name and size of the trace file is described in `prex(1)`. Think of the trace file as the least recently used circular buffer. Once the file has been filled, newer events will overwrite the older ones.

Use `TNF_PROBE_0` through `TNF_PROBE_5` to create production probes. These probes are compiled in by default. Developers are encouraged to embed such probes strategically, and to leave them compiled within production software. Such probes facilitate on-site analysis of the software.

Use `TNF_PROBE_0_DEBUG` through `TNF_PROBE_5_DEBUG` to create debug probes. These probes are compiled out by default. If you compile the program with the preprocessor option `-DTNF_DEBUG` (see `cc(1B)`), or with the preprocessor control statement `#define TNF_DEBUG` ahead of the `#include <tnf/probe.h>` statement, the debug probes will be compiled into the program. When compiled in, debug probes differ in only one way from the equivalent production probes. They contain an additional "debug" attribute which may be used to distinguish them from production probes at runtime, for example, when using `prex()`. Developers are encouraged to embed any number of probes for debugging purposes. Disabled probes have such a small runtime overhead that even large numbers of them do not make a significant impact.

If you compile with the preprocessor option `-DNPROBE` (see `cc(1B)`), or place the preprocessor control statement `#define NPROBE` ahead of the `#include <tnf/probe.h>` statement, no probes will be compiled into the program.

name The *name* of the probe should follow the syntax guidelines for identifiers in ANSI C. The use of *name* declares it, hence no separate declaration is necessary. This is a block scope declaration, so it does not affect the name space of the program.

keys *keys* is a string of space-separated keywords that specify the groups that the probe belongs to. Semicolons, single quotation marks, and the equal character (=) are not allowed in this string. If any of the groups are enabled, the probe is enabled. *keys* cannot be a variable. It must be a string constant.

detail *detail* is a string that consists of `<attribute> <value>` pairs that are each separated by a semicolon. The first word (up to the space) is considered to be the attribute and the rest of the string (up to the semicolon) is considered the value. Single quotation marks are used to denote a string value. Besides quotation marks, spaces separate multiple values. The value is optional. Although semicolons or single quotation marks generally are not allowed within either the attribute or the value, when text with embedded spaces is meant to denote a single value, use single quotes surrounding this text.

Use *detail* for one of two reasons. First, use *detail* to supply an attribute that a user can type into `prex(1)` to select probes. For example, if a user defines an attribute called `color`, then `prex(1)` can select probes based on the value of `color`. Second, use *detail* to annotate a probe with a string that is written out to a trace file only once. `prex(1)` uses spaces to tokenize the value when searching for a match. Spaces around the semicolon delimiter are allowed. *detail* cannot be a variable; it must be a string constant. For example, the *detail* string:

```
"XYZ%debug 'entering function A'; XYZ%exception 'no file';
XYZ%func_entry; XYZ%color red blue"
```

consists of 4 units:

| Attribute | Value | Values that <code>prex</code> matches on |
|----------------|-----------------------|--|
| XYZ%debug | 'entering function A' | 'entering function A' |
| XYZ%exception | 'no file' | 'no file' |
| XYZ%func_entry | ./.* | (regular expression) |
| XYZ%color | red blue | red <or> blue |

Attribute names must be prefixed by the vendor stock symbol followed by the '%' character. This avoids conflicts in the attribute name space. All attributes that do not have a '%' character are reserved. The following attributes are predefined:

| Attribute | Semantics |
|-----------|--|
| name | name of probe |
| keys | keys of the probe (value is space-separated tokens) |
| file | file name of the probe |
| line | line number of the probe |
| slots | slot names of the probe event (<i>arg_name_n</i>) |
| object | the executable or shared object that this probe is in. |
| debug | distinguishes debug probes from production probes |

arg_type_n

This is the type of the *n*th argument. The following are predefined TNF types:

| tnf Type | Associated C type (and semantics) |
|---------------|---|
| tnf_int | int |
| tnf_uint | unsigned int |
| tnf_long | long |
| tnf_ulong | unsigned long |
| tnf_longlong | long long (if implemented in compilation system) |
| tnf_ulonglong | unsigned long long (if implemented in compilation system) |
| tnf_float | float |
| tnf_double | double |
| tnf_string | char * |
| tnf_opaque | void * |

To define new TNF types that are records consisting of the predefined TNF types or references to other user defined types, use the interface specified in `TNF_DECLARE_RECORD(3TNF)`.

arg_name_n *arg_name_n* is the name that the user associates with the *n*th argument. Do not place quotation marks around *arg_name_n*. Follow the syntax guidelines for identifiers in ANSI C. The string version of *arg_name_n* is stored for every probe and can be accessed as the attribute "slots".

arg_value_n *arg_value_n* is evaluated to yield a value to be included in the trace file. A read access is done on any variables that are mentioned in *arg_value_n*. In a multi-threaded program, it is the user's responsibility to place locks around the `TNF_PROBE` macro if *arg_value_n* contains a variable that should be read protected.

EXAMPLES **EXAMPLE 1** `tracing(3TNF)`.

See `tracing(3TNF)` for complete examples showing debug and production probes in source code.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfd |
| MT Level | MT-Safe |

SEE ALSO

`cc(1B)`, `ld(1)`, `prex(1)`, `tnfdump(1)`, `libthread(3THR)`, `libtnfctl(3TNF)`, `TNF_DECLARE_RECORD(3TNF)`, `dlopen(3DL)`, `libtnfctl(3TNF)`, `tnf_process_disable(3TNF)`, `tracing(3TNF)`, `attributes(5)`

NOTES

If attaching to a running program with `prex(1)` to control the probes, compile the program with `-ltnfprobe` or start the program with the environment variable `LD_PRELOAD` set to `libtnfprobe.so.1`. See `ld(1)`. If `libtnfprobe` is explicitly linked into the program, it must be before `libthread` on the link line.

NAME | tnf_process_disable, tnf_process_enable, tnf_thread_disable, tnf_thread_enable – probe control internal interface

SYNOPSIS

```
cc [ flag ... ] file ... -ltnfprobe [ library ... ]
#include <tnf/probe.h>
void tnf_process_disable(void);

void tnf_process_enable(void);

void tnf_thread_disable(void);

void tnf_thread_enable(void);
```

DESCRIPTION | There are three levels of granularity for controlling tracing and probe functions (called probing from here on) – probing for the entire process, a particular thread, and the probe itself can be disabled/enabled. The first two (process and thread) are controlled by this interface. The probe is controlled via the application `prex(1)`.

`tnf_process_disable()` turns off probing for the process. The default process state is to have probing enabled. `tnf_process_enable()` turns on probing for the process.

`tnf_thread_disable()` turns off probing for the currently running thread. Threads are "born" or created with this state enabled. `tnf_thread_enable()` turns on probing for the currently running thread. If the program is a non-threaded program, these two thread interfaces disable or enable probing for the process.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfd |
| MT-Level | MT-Safe |

SEE ALSO | `prex(1)`, `tnfdump(1)`, `TNF_DECLARE_RECORD(3TNF)`, `TNF_PROBE(3TNF)`, `attributes(5)`

NOTES | A probe is considered enabled only if:

- `prex(1)` has enabled the probe AND
- the process has probing enabled – which is the default or could be set via `tnf_process_enable()` AND
- the thread that hits the probe has probing enabled – which is every thread's default or could be set via `tnf_thread_enable()`.

There is a run time cost associated with determining that the probe is disabled. To reduce the performance effect of probes, this cost should be minimized. The quickest way that a probe can be determined to be disabled is by the enable control that `prex(1)` uses. Therefore, to disable all the probes in a process use the `disable` command in `prex(1)` rather than `tnf_process_disable()`.

`tnf_process_disable()` and `tnf_process_enable()` should only be used to toggle probing based on some internal program condition. `tnf_thread_disable()` should be used to turn off probing for threads that are uninteresting.

| NAME | tracing – overview of tnf tracing system |
|------------------------|---|
| DESCRIPTION | <p>tnf tracing is a set of programs and API's that can be used to present a high-level view of the performance of an executable, a library, or part of the kernel. tracing is used to analyze a program's performance and identify the conditions that produced a bug.</p> <p>The core elements of tracing are:</p> |
| TNF_PROBE_*() | <p>The TNF_PROBE_*() macros define "probes" to be placed in code which, when enabled and executed, cause information to be added to a trace file. See TNF_PROBE(3TNF). If there are insufficient TNF_PROBE_* macros to store all the data of interest for a probe, data may be grouped into records. See TNF_DECLARE_RECORD(3TNF).</p> |
| prex | <p>Displays and controls probes in running software. See prex(1).</p> |
| kernel probes | <p>A set of probes built into the Solaris kernel which capture information about system calls, multithreading, page faults, swapping, memory management, and I/O. You can use these probes to obtain detailed traces of kernel activity under your application workloads. See tnf_kernel_probes(4).</p> |
| tnfextract | <p>A program that extracts the trace data from the kernel's in-memory buffer into a file. See tnfextract(1).</p> |
| tnfdump | <p>A program that displays the information from a trace file. See tnfdump(1).</p> |
| libtnfctl | <p>A library of interfaces that controls probes in a process. See libtnfctl(3TNF). prex(1) also utilizes this library. Other tools and processes use the libtnfctl interfaces to exercise fine control over their own probes.</p> |
| tnf_process_enable() | <p>A routine called by a process to turn on tracing and probe functions for the current process. See tnf_process_enable(3TNF).</p> |
| tnf_process_disable() | <p>A routine called by a process to turn off tracing and probe functions for the current process. See tnf_process_disable(3TNF).</p> |

`tnf_thread_enable()` A routine called by a process to turn on tracing and probe functions for the currently running thread. See `tnf_thread_enable(3TNF)`.

`tnf_thread_disable()` A routine called by a process to turn off tracing and probe functions for the currently running thread. See `tnf_thread_disable(3TNF)`.

EXAMPLES**EXAMPLE 1** Tracing a Process

The following function in some daemon process accepts job requests of various types, queuing them for later execution. There are two "debug probes" and one "production probe." Note that probes which are intended for debugging will not be compiled into the final version of the code; however, production probes are compiled into the final product.

```

/*
 * To compile in all probes (for development):
 * cc -DTNF_DEBUG ...
 *
 * To compile in only production probes (for release):
 * cc ...
 *
 * To compile in no probes at all:
 * cc -DNPROBE ...
 */
#include <tnf/probe.h>
void work(long, char *);
enum work_request_type { READ, WRITE, ERASE, UPDATE };
static char *work_request_name[] = {"read", "write", "erase", "update"};
main()
{
    long i;
    for (i = READ; i <= UPDATE; i++)
        work(i, work_request_name[i]);
}
void work(long request_type, char *request_name)
{
    static long q_length;
    TNF_PROBE_2_DEBUG(work_start, "work",
"XYZ%debug 'in function work'",
tnf_long, request_type_arg, request_type,
tnf_string, request_name_arg, request_name);
    /* assume work request is queued for later processing */
    q_length++;
    TNF_PROBE_1(work_queue, "work queue",
"XYZ%work_load heavy",
tnf_long, queue_length, q_length);
    TNF_PROBE_0_DEBUG(work_end, "work", "");
}

```

The production probe "work_queue," which remains compiled in the code, will, when enabled, log the length of the work queue each time a request is received.

The debug probes "work_start" and "work_end," which are compiled only during the development phase, track entry to and exit from the `work()` function and measure how much time is spent executing it. Additionally, the debug probe "work_start" logs the value of the two incoming arguments `request_type` and `request_name`. The runtime overhead for disabled probes is low enough that one can liberally embed them in the code with little impact on performance.

For debugging, the developer would compile with `-DTNF_DEBUG`, run the program under control of `prex(1)`, enable the probes of interest (in this case, all probes), continue the program until exit, and dump the trace file:

```
% cc
-DTNF_DEBUG -o daemon daemon.c # compile in all probes
% prex daemon # run program under prex control
Target process stopped
Type "continue" to resume the target, "help" for help ...
prex> list probes $all # list all probes in program
<probe list output here>
prex> enable $all # enable all probes
prex> continue # let target process execute
<program output here>
prex: target process finished
% ls /tmp/trace-* # trace output is in trace-<pid>
/tmp/trace-4194
% tnfdump /tmp/trace-4194 # get ascii output of trace file
<trace records output here>
```

For the production version of the system, the developer simply compiles without `-DTNF_DEBUG`.

EXAMPLE 2 Tracing the Kernel

Kernel tracing is similar to tracing a process; however, there are some differences. For instance, to trace the kernel, you need superuser privileges. The following example uses `prex(1)` and traces the probes in the kernel that capture system call information.

```
Allocate kernel
trace buffer and capture trace data:
root# prex -k
Type "help" for help ...
prex> buffer alloc 2m # allocate kernel trace buffer
Buffer of size 2097152 bytes allocated
prex> list probes $all # list all kernel probes
<probe list output here>
prex> list probes syscall # list syscall probes
# (keys=syscall)
<syscall probes list output here>
```

```

prex> enable syscall # enable only syscall probes
prex> ktrace on # turn on kernel tracing
<Run your application in another window at this point>
prex> ktrace off # turn off kernel tracing
prex> quit # exit prex
Extract the kernel's trace buffer into a file:
root# tnfxtract /tmp/ktrace # extract kernel trace buffer
Reset kernel tracing:
root# prex -k
prex> disable $all # disable all probes
prex> untrace $all # untrace all probes
prex> buffer dealloc # deallocate kernel trace buffer
prex> quit

```

CAUTION: Do not deallocate the trace buffer until you have extracted it into a trace file. Otherwise, you will lose the trace data that you collected from your experiment!

Examine the kernel trace file:

```

root# tnfdump /tmp/ktrace # get ascii dump of trace file
<trace records output here>

```

prex can also attach to a running process, list probes, and perform a variety of other tasks.

ATTRIBUTES

See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability | SUNWtnfd |
| MT Level | MT-Safe |

SEE ALSO

[prex\(1\)](#), [tnfdump\(1\)](#), [tnfxtract\(1\)](#), [TNF_DECLARE_RECORD\(3TNF\)](#), [TNF_PROBE\(3TNF\)](#), [libtnfctl\(3TNF\)](#), [tnf_process_disable\(3TNF\)](#), [tnf_kernel_probes\(4\)](#), [attributes\(5\)](#)

| | | | | | |
|----------------------|---|---------------------|---|--------------------|--|
| NAME | volmgt_acquire – reserve removable media device | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lvolmgt [library ...] #include <sys/types.h> #include <volmgt.h> int volmgt_acquire(char *dev, char *id, int ovr, char **err, pid_t *pidp);</pre> | | | | |
| DESCRIPTION | <p>The <code>volmgt_acquire()</code> routine reserves the removable media device specified as <code>dev</code>. <code>volmgt_acquire()</code> operates in two different modes, depending on whether or not Volume Management is running. See <code>vold(1M)</code>.</p> <p>If Volume Management <i>is</i> running, <code>volmgt_acquire()</code> attempts to reserve the removable media device specified as <code>dev</code>. Specify <code>dev</code> as <i>either</i> a symbolic device name (for example, <code>floppy0</code>) or a physical device pathname (for example, <code>/vol/dsk/unnamed_floppy</code>).</p> <p>If Volume Management <i>is not</i> running, <code>volmgt_acquire()</code> requires callers to specify a physical device pathname for <code>dev</code>. Specifying <code>dev</code> as a symbolic device name is <i>not</i> acceptable. In this mode, <code>volmgt_acquire()</code> relies entirely on the major and minor numbers of the device to determine whether or not the device is reserved.</p> <p>If <code>dev</code> is free, <code>volmgt_acquire()</code> updates the internal device reservation database with the caller's process id (<code>pid</code>) and the specified <code>id</code> string.</p> <p>If <code>dev</code> is reserved by another process, the reservation attempt fails and <code>volmgt_acquire()</code>:</p> <ul style="list-style-type: none"> ■ sets <code>errno</code> to <code>EBUSY</code> ■ fills the caller's <code>id</code> value in the array pointed to by <code>err</code> ■ fills in the <code>pid</code> to which the pointer <code>pidp</code> points with the <code>pid</code> of the process which holds the reservation, if the supplied <code>pidp</code> is non-zero <p>If the override <code>ovr</code> is non-zero, the call overrides the device reservation.</p> | | | | |
| RETURN VALUES | <p>Upon successful completion, <code>volmgt_acquire()</code> returns a non-zero value.</p> <p>Upon failure, <code>volmgt_acquire()</code> returns 0. If the return value is 0, and <code>errno</code> is set to <code>EBUSY</code>, the address pointed to by <code>err</code> contains the string that was specified as <code>id</code> (when the device was reserved by the process holding the reservation).</p> | | | | |
| ERRORS | <p>The <code>volmgt_acquire()</code> routine fails if one or more of the following are true:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;"><code>EINVAL</code></td> <td>One of the specified arguments is invalid or missing.</td> </tr> <tr> <td><code>EBUSY</code></td> <td><code>dev</code> is already reserved by another process (and <code>ovr</code> was not set to a non-zero value)</td> </tr> </table> | <code>EINVAL</code> | One of the specified arguments is invalid or missing. | <code>EBUSY</code> | <code>dev</code> is already reserved by another process (and <code>ovr</code> was not set to a non-zero value) |
| <code>EINVAL</code> | One of the specified arguments is invalid or missing. | | | | |
| <code>EBUSY</code> | <code>dev</code> is already reserved by another process (and <code>ovr</code> was not set to a non-zero value) | | | | |

EXAMPLES**EXAMPLE 1** Using volmgt_acquire()

In the following example, Volume Management is running and the first floppy drive is reserved, accessed and released.

```
#include <volmgt.h>
char *errp;
if (!volmgt_acquire("floppy0", "FileMgr", 0, NULL,
    &errp, NULL)) {
    /* handle error case */
    ...
}
/* floppy acquired - now access it */
if (!volmgt_release("floppy0")) {
    /* handle error case */
    ...
}
}
```

EXAMPLE 2 Using volmgt_acquire() To Override A Lock On Another Process

The following example shows how callers can override a lock on another process using volmgt_acquire().

```
char *errp, buf[20];
int override = 0;
pid_t pid;
if (!volmgt_acquire("floppy0", "FileMgr", 0, &errp,
    &pid)) {
    if (errno == EBUSY) {
        (void) printf("override %s (pid=%ld)?\n",
            errp, pid); {
            (void) fgets(buf, 20, stdin);
            if (buf[0] == 'y') {
                override++;
            }
        }
    } else {
        /* handle other errors */
        ...
    }
}
if (override) {
    if (!volmgt_acquire("floppy0", "FileMgr", 1,
        &errp, NULL)) {
        /* really give up this time! */
        ...
    }
}
}
```

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

vold(1M), free(3C), malloc(3C), volmgt_release(3VOLMGT), attributes(5)

NOTES

When returning a string through *err*, `volmgt_acquire()` allocates a memory area using `malloc(3C)`. Use `free(3C)` to release the memory area when no longer needed.

The *ovr* argument is intended to allow callers to override the current device reservation. It is assumed that the calling application has determined that the current reservation can safely be cleared. See `EXAMPLES`.

NAME | volmgt_check – have Volume Management check for media

SYNOPSIS | `cc [flag ...] file ... -lvolmgt [library ...]`
`#include <volmgt.h>`

DESCRIPTION | `int volmgt_check(char *pathname);`
 This routine asks Volume Management to check the specified *pathname* and determine if new media has been inserted in that drive.
 If a null pointer is passed in, then Volume Management will check each device it is managing that can be checked.
 If new media is found, `volmgt_check()` tells Volume Management to initiate any "actions" specified in `/etc/vold.conf` (see `vold.conf(4)`).

RETURN VALUES | This routine returns 0 if no media was found, and a non-zero value if any media was found.

ERRORS | This routine can fail, returning 0, if a `stat(2)` or `open(2)` of the supplied *pathname* fails, or if any of the following is true:
ENXIO | Volume Management is not running.
EINTR | An interrupt signal was detected while checking for media.

EXAMPLES | **EXAMPLE 1** Checking If Any New Media Is Inserted
 To check if any drive managed by Volume Management has any new media inserted in it:

```
if (volmgt_check(NULL)) {
    (void) printf("Volume Management found media\n");
}
```

 This would also request Volume Management to take whatever action was specified in `/etc/vold.conf` for any media found.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | `cc(1B)`, `volcheck(1)`, `vold(1M)`, `open(2)`, `stat(2)`, `volmgt_inuse(3VOLMGT)`, `volmgt_running(3VOLMGT)`, `vold.conf(4)`, `attributes(5)`, `volfs(7FS)`

NOTES | Volume Management must be running for this routine to work.

Since `volmgt_check()` returns 0 for two different cases (both when no media is found, and when an error occurs), it is up to the user to check *errno* to differentiate the two, and to ensure that Volume Management is running.

| NAME | volmgt_feature_enabled – check whether specific Volume Management features are enabled | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>cc [flag ...] file ... -l volmgt [library ...] #include <volmgt.h> int volmgt_feature_enabled(char *feat_str);</pre> | | | | |
| DESCRIPTION | <p>The volmgt_feature_enabled() routine checks whether specific Volume Management features are enabled. volmgt_feature_enabled() checks for the Volume Management features passed in to it by the feat_str parameter.</p> <p>Currently, the only supported feature string that volmgt_feature_enabled() checks for is floppy-summit-interfaces. The floppy-summit-interfaces feature string checks for the presence of the libvolmgt routines volmgt_acquire() and volmgt_release().</p> <p>The list of features that volmgt_feature_enabled() checks for is expected to expand in the future.</p> | | | | |
| RETURN VALUES | 0 is returned if the specified feature is not currently available. A non-zero value indicates that the specified feature is currently available. | | | | |
| EXAMPLES | <p>EXAMPLE 1 A sample of the volmgt_feature_enabled() function.</p> <p>In the following example, volmgt_feature_enabled() checks whether the floppy-summit-interfaces feature is enabled.</p> <pre>if (volmgt_feature_enabled("floppy-summit-interfaces")) { (void) printf("Media Sharing Routines ARE present\n"); } else { (void) printf("Media Sharing Routines are NOT present\n"); }</pre> | | | | |
| ATTRIBUTES | <p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | volmgt_acquire(3VOLMGT), volmgt_release(3VOLMGT), attributes(5) | | | | |

NAME | volmgt_inuse – check whether or not Volume Management is managing a *pathname*

SYNOPSIS |

```
cc [ flag ... ] file ... -lvolmgt [ library ... ]
#include <volmgt.h>

int volmgt_inuse(char *pathname);
```

DESCRIPTION | volmgt_inuse() checks whether Volume Management is managing the specified *pathname*.

RETURN VALUES | A non-zero value is returned if Volume Management is managing the specified *pathname*, otherwise 0 is returned.

ERRORS | This routine can fail, returning 0, if a *stat(2)* of the supplied *pathname* or an *open(2)* of */dev/volctl* fails, or if any of the following is true:

ENXIO | Volume Management is not running.

EINTR | An interrupt signal was detected while checking for the supplied *pathname* for use.

EXAMPLES | **EXAMPLE 1** Using volmgt_inuse()

To see if Volume Management is managing the first floppy disk:

```
if (volmgt_inuse("/dev/rdiskette0") != 0) {
    (void) printf("volmgt is managing diskette 0\n");
} else {
    (void) printf("volmgt is NOT managing diskette 0\n");
}
```

ATTRIBUTES | See *attributes(5)* for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | *cc(1B)*, *vold(1M)*, *open(2)*, *stat(2)*, *errno(3C)*, *volmgt_check(3VOLMGT)*, *volmgt_running(3VOLMGT)*, *attributes(5)*, *volfs(7FS)*

NOTES | This routine requires Volume Management to be running.

Since *volmgt_inuse()* returns 0 for two different cases (both when a volume is not in use, and when an error occurs), it is up to the user to check *errno* to differentiate the two, and to ensure that Volume Management is running.

| NAME | volmgt_ownspath – check Volume Management name space for path | | | | | | |
|----------------------|--|----------------|-----------------|----------|------|------------------|--------|
| SYNOPSIS | <pre>cc [flag ...] file ...-lvolgmt [library ...] #include <volmgt.h> int volmgt_ownspath(char *path);</pre> | | | | | | |
| DESCRIPTION | volmgt_ownspath() checks to see if a given <i>path</i> is contained in the Volume Management name space. This is achieved by comparing the beginning of the supplied path name with the output from volmgt_root(3VOLMGT) | | | | | | |
| PARAMETERS | <i>path</i> A string containing the path. | | | | | | |
| RETURN VALUES | <pre>non-zero The path is owned by Volume Management. 0 volgmt() does not have path in its name space, or Volume Management is not running.</pre> | | | | | | |
| EXAMPLES | <p>EXAMPLE 1 Using volmgt_ownspath()</p> <p>The following example first checks if volmgt() is running, then checks the Volume Management name space for <i>path</i>, and then returns the <i>id</i> for the piece of media.</p> <pre>char *path; ... if (volmgt_running()) { if (volmgt_ownspath(path)) { (void) printf("id of %s is %lld\n", path, media_getid(path)); } }</pre> | | | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT Level</td> <td>Safe</td> </tr> <tr> <td>Commitment Level</td> <td>Public</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT Level | Safe | Commitment Level | Public |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | | | |
| MT Level | Safe | | | | | | |
| Commitment Level | Public | | | | | | |
| SEE ALSO | volmgt_root(3VOLMGT), volmgt_running(3VOLMGT) attributes(5) | | | | | | |

| | |
|----------------------|---|
| NAME | volmgt_release – release removable media device reservation |
| SYNOPSIS | <pre>cc [flag ...] file ... -lvolmgt [library ...] #include <volmgt.h></pre> |
| DESCRIPTION | <p>The <code>volmgt_release()</code> routine releases the removable media device reservation specified as <i>dev</i>. See <code>volmgt_acquire(3VOLMGT)</code> for a description of <i>dev</i>.</p> <p>If <i>dev</i> is reserved by the caller, <code>volmgt_release()</code> updates the internal device reservation database to indicate that the device is no longer reserved. If the requested device is reserved by another process, the release attempt fails and <code>errno</code> is set to 0.</p> |
| RETURN VALUES | <p>Upon successful completion, <code>volmgt_release</code> returns a non-zero value. Upon failure, 0 is returned.</p> |
| ERRORS | <p>On failure, <code>volmgt_release()</code> returns 0, and sets <code>errno</code> for one of the following conditions:</p> <p><code>EINVAL</code> <i>dev</i> was invalid or missing.</p> <p><code>EBUSY</code> <i>dev</i> was not reserved by the caller.</p> |
| EXAMPLES | <p>EXAMPLE 1 Using <code>volmgt_release()</code></p> <p>In the following example, Volume Management is running, and the first floppy drive is reserved, accessed and released.</p> <pre>#include <volmgt.h> char *errp; if (!volmgt_acquire("floppy0", "FileMgr", 0, &errp, NULL)) { /* handle error case */ ... } /* floppy acquired - now access it */ if (!volmgt_release("floppy0")) { /* handle error case */ ... }</pre> |
| ATTRIBUTES | See <code>attributes(5)</code> for descriptions of the following attributes: |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT-Level | MT-Safe |
| Interface Stability | Stable |

SEE ALSO

vold(1M), volmgt_acquire(3VOLMGT), attributes(5)

| NAME | volmgt_root – return the Volume Management root directory | | | | |
|----------------------|---|----------------|-----------------|----------|---------|
| SYNOPSIS | <pre>cc [flag ...] file ... -lvolmgt [library ...] #include <volmgt.h> char *volmgt_root(void);</pre> | | | | |
| DESCRIPTION | volmgt_root() returns the current Volume Management root directory, which by default is /vol but can be configured to be in a different location. | | | | |
| RETURN VALUES | A pointer to a static string containing the root directory for Volume Management is returned. | | | | |
| ERRORS | This routine may fail if an open() of /dev/volctl fails. If this occurs a pointer to the default Volume Management root directory is returned. | | | | |
| EXAMPLES | <p>EXAMPLE 1 Finding the Volume Management root directory.</p> <p>To find out where the Volume Management root directory is:</p> <pre>if ((path = volmgt_root()) != NULL) { (void) printf("Volume Management root dir=%s\n", path); } else { (void) printf("can't find Volume Management root dir\n"); }</pre> | | | | |
| FILES | /vol Default location for the Volume Management root directory | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | cc(1B), vold(1M), open(2), volmgt_check(3VOLMGT), volmgt_inuse(3VOLMGT), volmgt_running (3VOLMGT), attributes(5), volfs(7FS) | | | | |
| NOTES | This routine will return the default root directory location even when Volume Management is not running. | | | | |

NAME | volmgt_running – return whether or not Volume Management is running

SYNOPSIS | `cc [flag ...] file ... -lvolmgt [library ...]`
`#include <volmgt.h>`

DESCRIPTION | `int volmgt_running(void);`
 volmgt_running() tells whether or not Volume Management is running.

RETURN VALUES | A non-zero value is returned if Volume Management is running, else 0 is returned.

ERRORS | volmgt_running() will fail, returning 0, if a stat(2) or open(2) of /dev/volctl fails, or if any of the following is true:
 ENXIO Volume Management is not running.
 EINTR An interrupt signal was detected while checking to see if Volume Management was running.

EXAMPLES | **EXAMPLE 1** Using volmgt_running()
 To see if Volume Management is running:

```

if (volmgt_running() != 0) {
    (void) printf("Volume Management is running\n");
} else {
    (void) printf("Volume Management is NOT running\n");
}
    
```

ATTRIBUTES | See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO | cc(1B), vold(1M), open(2), stat(2), volmgt_check(3VOLMGT), volmgt_inuse (3VOLMGT), attributes(5), volfs(7FS)

NOTES | Volume Management must be running for many of the Volume Management library routines to work.

| | | | | | | | | | |
|----------------------|--|-------|-----------------------------------|-------|---|-------|-----------------------------------|-------|--|
| NAME | volmgt_symname, volmgt_symdev – convert between Volume Management symbolic names, and the devices that correspond to them | | | | | | | | |
| SYNOPSIS | <pre>cc [flag ...] file ... -lvolmgt [library ...] #include <volmgt.h> char *volmgt_symname(char *pathname); char *volmgt_symdev(char *symname);</pre> | | | | | | | | |
| DESCRIPTION | <p>These two routines compliment each other, translating between Volume Management's symbolic name for a device, called a <i>symname</i>, and the <i>/dev pathname</i> for that same device.</p> <p>volmgt_symname () converts a supplied <i>/dev pathname</i> to a <i>symname</i>, Volume Management's idea of that device's symbolic name (see volfs(7FS) for a description of Volume Management symbolic names).</p> <p>volmgt_symdev () does the opposite conversion, converting between a <i>symname</i>, Volume Management's idea of a device's symbolic name for a volume, to the <i>/dev pathname</i> for that device.</p> | | | | | | | | |
| RETURN VALUES | <p>volmgt_symname () returns the symbolic name for the device pathname supplied, and volmgt_symdev () returns the device pathname for the supplied symbolic name.</p> <p>These strings are allocated upon success, and therefore must be freed by the caller when they are no longer needed (see free(3C)).</p> | | | | | | | | |
| ERRORS | <p>volmgt_symname () can fail, returning a null string pointer, if a stat(2) of the supplied <i>pathname</i> fails, or if an open(2) of <i>/dev/volctl</i> fails, or if any of the following is true:</p> <table border="0"> <tr> <td style="padding-right: 20px;">ENXIO</td> <td>Volume Management is not running.</td> </tr> <tr> <td>EINTR</td> <td>An interrupt signal was detected while trying to convert the supplied <i>pathname</i> to a <i>symname</i>.</td> </tr> </table> <p>volmgt_symdev () can fail if an open(2) of <i>/dev/volctl</i> fails, or if any of the following is true:</p> <table border="0"> <tr> <td style="padding-right: 20px;">ENXIO</td> <td>Volume Management is not running.</td> </tr> <tr> <td>EINTR</td> <td>An interrupt signal was detected while trying to convert the supplied <i>symname</i> to a <i>/dev pathname</i>.</td> </tr> </table> | ENXIO | Volume Management is not running. | EINTR | An interrupt signal was detected while trying to convert the supplied <i>pathname</i> to a <i>symname</i> . | ENXIO | Volume Management is not running. | EINTR | An interrupt signal was detected while trying to convert the supplied <i>symname</i> to a <i>/dev pathname</i> . |
| ENXIO | Volume Management is not running. | | | | | | | | |
| EINTR | An interrupt signal was detected while trying to convert the supplied <i>pathname</i> to a <i>symname</i> . | | | | | | | | |
| ENXIO | Volume Management is not running. | | | | | | | | |
| EINTR | An interrupt signal was detected while trying to convert the supplied <i>symname</i> to a <i>/dev pathname</i> . | | | | | | | | |
| EXAMPLES | <p>EXAMPLE 1 Testing Floppies</p> <p>The following tests how many floppies Volume Management currently sees in floppy drives (up to 10):</p> <pre>for (i=0; i < 10; i++) { (void) sprintf(path, "floppy%d", i); if (volmgt_symdev(path) != NULL) {</pre> | | | | | | | | |

```

        (void) printf("volume %s is in drive %d\
",
        path, i);
    }
}

```

EXAMPLE 2 Finding The Symbolic Name

This code finds out what symbolic name (if any) Volume Management has for /dev/rdisk/c0t6d0s2 :

```

if ((nm = volmgt_symname("/dev/rdisk/c0t6d0s2")) == NULL) {
    (void) printf("path not managed\
");
} else {
    (void) printf("path managed as %s\
", nm);
}

```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level | MT-Safe |

SEE ALSO

`cc(1B)`, `vold(1M)`, `open(2)`, `stat(2)`, `free(3C)`, `malloc(3C)`, `volmgt_check(3VOLMGT)`, `volmgt_inuse(3VOLMGT)`, `volmgt_running(3VOLMGT)`, `attributes(5)`, `volfs(7FS)`

NOTES

These routines only work when Volume Management is running.

BUGS

There should be a straightforward way to query Volume Management for a list of all media types it's managing, and how many of each type are being managed.

| NAME | y0, y1, yn – Bessel functions of the second kind | | | | |
|----------------------|--|----------------|-----------------|----------|---------|
| SYNOPSIS | cc [<i>flag ...</i>] file ... -lm [<i>library ...</i>] double y0(double x); double y1(double x); double yn(int n, double x); | | | | |
| DESCRIPTION | The y0(), y1() and yn() functions compute Bessel functions of x of the second kind of orders 0, 1 and n respectively. The value of x must be positive. | | | | |
| RETURN VALUES | Upon successful completion, y0(), y1() and yn() will return the relevant Bessel value of x of the second kind. If x is NaN, NaN is returned. If the x argument to y0(), y1() or yn() is negative, -HUGE_VAL or NaN is returned, and errno may be set to EDOM. If x is 0.0, -HUGE_VAL is returned and errno may be set to ERANGE or EDOM. If the correct result would cause overflow, -HUGE_VAL is returned and errno may be set to ERANGE. For exceptional cases, matherr(3M) tabulates the values to be returned as dictated by Standards other than XPG4. | | | | |
| ERRORS | The y0(), y1() and yn() functions may fail if: EDOM The value of x is negative. ERANGE The value of x is too large in magnitude, or x is 0.0, or the correct result would cause overflow. | | | | |
| USAGE | An application wishing to check for error situations should set errno to 0 before calling y0(), y1() or yn(). If errno is non-zero on return, or the return value is NaN, an error has occurred. | | | | |
| ATTRIBUTES | See attributes(5) for descriptions of the following attributes: | | | | |
| | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table> | ATTRIBUTE TYPE | ATTRIBUTE VALUE | MT-Level | MT-Safe |
| ATTRIBUTE TYPE | ATTRIBUTE VALUE | | | | |
| MT-Level | MT-Safe | | | | |
| SEE ALSO | isnan(3M), j0(3M), matherr(3M), attributes(5), standards(5) | | | | |

Index

A

- absolute value function — fabs 198
- access CPU performance counters
 - in other processes -
cpc_pctx_bind_event 85
- aclcheck — check the validity of an ACL 29
- aclfrommode - convert an ACL to or from permission bits 32
- aclfromtext - convert internal representation to or from external representation 33
- aclsort — sort an ACL 31
- acltomode - convert an ACL to or from permission bits 32
- acltotext - convert internal representation to or from external representation 33
- acos — arc cosine function 35
- acosh - inverse hyperbolic functions 36
- advance - regular expression compile and match routines 370
- allocate or deallocate a buffer for trace data
 - tnftcl_buffer_alloc 392
 - tnftcl_buffer_dealloc 392
- annotate source code with info for tools
 - NOTE 311
 - _NOTE 311
- arc cosine function — acos 35
- arc sine function — asin 37
- arc tangent function — atan2 38-39
- asin — arc sine function 37
- asinh - inverse hyperbolic functions 36
- associate callbacks with process events —
pctx_set_events 361
- atan — arc tangent function 39
- atan2 — arc tangent function 38
- atanh - inverse hyperbolic functions 36
- au_close - construct audit records 40
- au_open - construct audit records 40
- au_preselect — preselect an audit record 41
- au_to - create audit record tokens 44
- au_to_arg - create audit record tokens 44
- au_to_attr - create audit record tokens 44
- au_to_data - create audit record tokens 44
- au_to_groups - create audit record tokens 44
- au_to_in_addr - create audit record tokens 44
- au_to_ipc - create audit record tokens 44
- au_to_ipc_perm - create audit record tokens 44
- au_to_important - create audit record tokens 44
- au_to_me - create audit record tokens 44
- au_to_new_in_addr - create audit record tokens 44
- au_to_new_process - create audit record tokens 44
- au_to_new_socket - create audit record tokens 44
- au_to_new_subject - create audit record tokens 44
- au_to_opaque - create audit record tokens 44
- au_to_path - create audit record tokens 44
- au_to_process - create audit record tokens 44
- au_to_return - create audit record tokens 44
- au_to_socket - create audit record tokens 44

au_to_subject – create audit record tokens 44
 au_to_text – create audit record tokens 44
 au_user_mask — get user’s binary preselection mask 47
 au_write – write audit records 40
 audit control file information
 – endac 207
 – getacdir 207
 – getacflg 207
 – getacinfo 207
 – getacmin 207
 – getacna 207
 – setac 207
 audit record tokens, manipulating
 – au_close 40
 – au_open 40
 — au_preselect 41
 – au_write 40
 authentication information routines for PAM
 – pam_get_item 336
 – pam_set_item 336
 authentication transaction routines for PAM
 – pam_end 353
 – pam_start 353

B

base 10 logarithm function — log10 267
 Basic Security Module functions
 – au_close 40
 – au_open 40
 — au_preselect 41
 — au_user_mask 47
 – au_write 40
 Bessel functions of the first kind
 – j0 233
 – j1 233
 – jn 233
 Bessel functions of the second kind
 – y0 451
 – y1 451
 – yn 451
 bgets — read stream up to next delimiter 49
 buffer
 split into fields — bufsplit 50

C

cbrt — cube root function 51
 ceil — ceiling value function 52
 ceiling value function — ceil 52
 change or add a value to the PAM environment
 — pam_putenv 330
 check the validity of an ACL — aclcheck 29
 check whether or not Volume Management
 is managing a pathname —
 volmgt_inuse 443
 check whether specific Volume Management
 features are enabled —
 volmgt_feature_enabled 442
 chkauthattr – verify user authorization 215
 class-dependent data translation
 – elf32_xlatetof 155
 – elf32_xlatetom 155
 – elf64_xlatetof 155
 – elf64_xlatetom 155
 close a tnfcntl handle — tnfcntl_close 394
 commands
 open, close to and from a command –
 p2open, p2close 313
 compile – regular expression compile and
 match routines 370
 compute natural logarithm — log1p 268
 computes exponential functions — expm1 197
 config_admin – configuration administration
 interface 54
 config_ap_id_cmp – configuration
 administration interface 54
 config_change_state – configuration
 administration interface 54
 config_list – configuration administration
 interface 54
 config_list_ext – configuration administration
 interface 54
 config_private_func – configuration
 administration interface 54
 config_stat – configuration administration
 interface 54
 config_strerror – configuration administration
 interface 54
 config_test – configuration administration
 interface 54
 config_unload – configuration administration
 interface 54

configuration administration interface –
 config_list_ext 54

connect to a DMI service provider
 — ConnectToServer 62, 119

control kernel tracing and process filtering
 – tnftcl_filter_list_add 421
 – tnftcl_filter_list_delete 421
 – tnftcl_filter_list_get 421
 – tnftcl_filter_state_set 421
 – tnftcl_trace_state_set 421

control probes of another process where caller
 provides /proc functionality
 – tnftcl_check_libs 396
 – tnftcl_indirect_open 396

convert an ACL to or from permission bits –
 actlmode 32–33

convert a supplied name into an absolute
 pathname that can be used to
 access removable media —
 media_findname 284

convert between Volume Management symbolic
 names, and the devices that
 correspond to them
 – volmgt_symdev 449
 – volmgt_symname 449

coordinate CPC library and application versions
 — cpc_version 93

copysign — return magnitude of first
 argument and sign of second
 argument 64

cos — cosine function 65

cosh — hyperbolic cosine function 66

cosine function — cos 65

cpc — hardware performance counters 67

cpc_access — test access CPU performance
 counters 70

cpc_bind_event – use CPU performance
 counters on lwps 71

cpc_count_sys_events – enable and disable
 performance counters 77

cpc_count_usr_events – enable and disable
 performance counters 77

cpc_event — data structure to describe CPU
 performance counters 79

cpc_event_accum – simple difference and
 accumulate operations 81

cpc_event_diff – simple difference and
 accumulate operations 81

cpc_eventtostr – translate strings to and from
 events 90

cpc_getcciname – determine CPU performance
 counter configuration 83

cpc_getcpuref – determine CPU performance
 counter configuration 83

cpc_getcpuver – determine CPU performance
 counter configuration 83

cpc_getnpic – determine CPU performance
 counter configuration 83

cpc_getusage – determine CPU performance
 counter configuration 83

cpc_pctx_bind_event – access CPU
 performance counters
 in other processes 85

cpc_pctx_invalidate – access CPU
 performance counters
 in other processes 85

cpc_pctx_rele – access CPU performance
 counters in other
 processes 85

cpc_pctx_take_sample – access CPU
 performance counters in other
 processes 85

cpc_rele – use CPU performance counters on
 lwps 71

cpc_strtoevent – translate strings to and from
 events 90

cpc_take_sample – use CPU performance
 counters on lwps 71

cpc_version — coordinate CPC library and
 application versions 93

cpc_walk_names – determine CPU performance
 counter configuration 83

cplus_demangle – decode a C++ encoded
 symbol name 94

create audit record tokens – au_to_in_addr 44

create DmiOctetString in dynamic memory
 — newDmiOctetString 307

create DmiString in dynamic memory
 — newDmiString 308

create handle for internal process probe control
 — tnftcl_internal_open 399

create handle for kernel probe control —
 tnftcl_kernel_open 401

cube root function — `cbrt` 51

D

data structure to describe CPU performance counters — `cpc_event` 79

decode a C++ encoded symbol name

– `cplus_demangle` 94

– `demangle` 94

`demangle` – decode a C++ encoded symbol name 94

destroy a layout object —

`m_destroy_layout` 283

determine CPU performance counter

configuration –

`cpc_getusage` 83

device ID interfaces for user applications –

`devid_compare` 95

`devid_compare` – device ID interfaces for user applications 95

`devid_deviceid_to_nmlist` – device ID interfaces for user applications 95

`devid_free` – device ID interfaces for user applications 95

`devid_free_nmlist` – device ID interfaces for user applications 95

`devid_get` – device ID interfaces for user applications 95

`devid_get_minor_name` – device ID interfaces for user applications 95

`devid_sizeof` – device ID interfaces for user applications 95

directories

create, remove them in a path – `mkdirp`, `rmdirp` 292

`dmi_error` — print error in string form 134

`DmiAddComponent` – Management Interface database administration functions 124

`DmiAddGroup` – Management Interface database administration functions 124

`DmiAddLanguage` – Management Interface database administration functions 124

`DmiAddRow` – Management Interface operation functions 129

`DmiDeleteComponent` – Management Interface database administration functions 124

`DmiDeleteGroup` – Management Interface database administration functions 124

`DmiDeleteLanguage` – Management Interface database administration functions 124

`DmiDeleteRow` – Management Interface operation functions 129

`DmiGetAttribute` – Management Interface operation functions 129

`DmiGetConfig` – Management Interface initialization functions 135

`DmiGetMultiple` – Management Interface operation functions 129

`DmiGetVersion` – Management Interface initialization functions 135

`DmiListAttributes` – Management Interface listing functions 139

`DmiListClassNames` – Management Interface listing functions 139

`DmiListComponents` – Management Interface listing functions 139

`DmiListComponentsByClass` – Management Interface listing functions 139

`DmiListGroup` – Management Interface listing functions 139

`DmiListLanguages` – Management Interface listing functions 139

`DmiOriginateEvent` – Service Provider functions for components 145

`DmiRegister` – Management Interface initialization functions 135

`DmiRegisterCi` – Service Provider functions for components 145

`DmiSetAttribute` – Management Interface operation functions 129

`DmiSetConfig` – Management Interface initialization functions 135

`DmiSetMultiple` – Management Interface operation functions 129

`DmiUnregister` – Management Interface initialization functions 135

DmiUnRegisterCi – Service Provider functions
for components 145

E

elf – object file access library 157
 get entries from name list — nlist 310
elf_begin – process ELF object files 163
elf_cntl – control an elf file descriptor 168
elf_end – process ELF object files 163
elf_errmsg – error handling 170
elf_errno – error handling 170
elf_fill – set fill byte 171
elf_flagdata – manipulate flags 172
elf_flagehdr – manipulate flags 172
elf_flagelf – manipulate flags 172
elf_flagphdr – manipulate flags 172
elf_flagphdr – manipulate flags 172
elf_flagshdr – manipulate flags 172
elf_getarhdr – retrieve archive member
 header 174
elf_getarsym – retrieve archive symbol
 table 176
elf_getbase – get the base offset for an object
 file 177
elf_getdata – get section data 178
elf_getident – retrieve file identification
 data 182
elf_getscn – get section information 184
elf_hash – compute hash value 186
elf_kind – determine file type 187
elf_memory – process ELF object files 163
elf_ndxscn – get section information 184
elf_newdata – get section data 178
elf_newscn – get section information 184
elf_next – process ELF object files 163
elf_nextscn – get section information 184
elf_rand – process ELF object files 163
elf_rawdata – get section data 178
elf_rawfile – retrieve uninterpreted file
 contents 188
elf_strptr – make a string pointer 189
elf_update – update an ELF descriptor 190
elf_version – coordinate ELF library and
 application versions 194
elf32_checksum – return the checksum of an
 elf image

 – elf64_checksum 147
elf32_fsize – return the size of an object file
 type 148
elf32_getehdr – retrieve class-dependent object
 file header 149
elf32_getphdr – retrieve class-dependent
 program header table 151
elf32_getshdr – retrieve class-dependent section
 header 153
elf32_newehdr – retrieve class-dependent object
 file header 149
elf32_newphdr – retrieve class-dependent
 program header table 151
elf32_xlatetof – class-dependent data
 translation 155
elf32_xlatetom – class-dependent data
 translation 155
elf64_checksum – return the checksum of an
 elf image
 – elf32_checksum 147
elf64_fsize – return the size of an object file
 type 148
elf64_getehdr – retrieve class-dependent object
 file header 149
elf64_getphdr – retrieve class-dependent
 program header table 151
elf64_getshdr – retrieve class-dependent section
 header 153
elf64_newehdr – retrieve class-dependent object
 file header 149
elf64_newphdr – retrieve class-dependent
 program header table 151
elf64_xlatetof – class-dependent data
 translation 155
elf64_xlatetom – class-dependent data
 translation 155
enable and disable performance counters –
 cpc_count_usr_events 77
encryption
 determine whether a buffer of characters is
 encrypted — isencrypt 231
endac – get audit control file information 207
endauclass – close audit_class database
 file 209
endauevent – close audit_event database
 file 212

endauthattr – get authorization database entry 215
 endauuser – get audit_user database entry 218
 endexecattr – get execution attribute database entry 220
 endprofattr – get profile description database entry 224
 enduserattr – get user_attr entry 226
 erf – error and complementary error functions 195
 erfc – error and complementary error functions 195
 error and complementary error functions
 – erf 195
 – erfc 195
 Euclidean distance function — hypot 229
 Executable and Linking Formatelf
 exp — exponential function 196
 expm1 — computes exponential functions 197
 exponential function — exp 196

F

fabs — absolute value function 198
 files
 search for named file in named directories
 — pathfind 357
 floating-point remainder value function —
 fmod 200
 floor — floor function 199
 floor function — floor 199
 fmod — floating-point remainder value function 200
 free dynamic memory allocated for input DmiString structure
 — freeDmiString 201
 free_authattr – release memory 215
 free_execattr – release memory 220
 free_profattr – get profile description database entry 224
 free_userattr – get user_attr entry 226
 freeDmiString— free dynamic memory allocated for input DmiString structure 201
 functions to manage lockfile(s) for user’s mailbox
 – maillock 271

– mailunlock 271
 – touchlock 271

G

gamma – log gamma function 257
 gamma_r – log gamma function 257
 get user_attr entry – getuserid 226
 get and set media attributes
 – media_getattr 287
 – media_setattr 287
 get the trace attributes from a tnftcl handle —
 tnftcl_trace_attrs_get 418
 getacdir – get audit control file information 207
 getacflg – get audit control file information 207
 getacinfo – get audit control file information 207
 getacmin – get audit control file information 207
 getacna – get audit control file information 207
 getaclassent – get audit_class database entry 209
 getaclassent_r – get audit_class database entry 209
 getaclassnam – get audit_class database entry 209
 getaclassnam_r – get audit_class database entry 209
 getauditflags() — generate process audit state 223
 getauditflagsbin() – convert audit flag specifications 211
 getauditflagschar() – convert audit flag specifications 211
 getaevent – get audit_event database entry 212
 getaevent_r – get audit_event database entry 212
 getaevnam – get audit_event database entry 212
 getaevnam_r – get audit_event database entry 212
 getaevnonam – get audit_event database entry 212
 getaevnum – get audit_event database entry 212

getauevnum_r – get audit_event database entry 212
 getauthattr – get authorization database entry 215
 getauthnam – get authorization database entry 215
 getauuserent – get audit_user database entry 218
 getauuserent_r – get audit_user database entry 218
 getauusernam – get audit_user database entry 218
 getauusernam_r – get audit_user database entry 218
 getexecattr – get execution attribute database entry 220
 getexecuser – get many entries 220
 getprofattr – get profile description database entry 224
 getprofnam – get profile description database entry 224
 getuserattr – get user_attr entry 226
 getusernam – get user_attr entry 226
 getuserid – get user_attr entry 226
 gmatch — shell global pattern matching 228

H

hardware performance counters — cpc 67
 have Volume Management check for media — volmgt_check 440
 hyperbolic cosine function — cosh 66
 hyperbolic sine function — sinh 379
 hyperbolic tangent function — tanh 391
 hypot — Euclidean distance function 229

I

ilogb — returns an unbiased exponent 230
 initialize a layout object — m_create_layout 279
 initialize kernel statistics facility
 – kstat_close 245
 – kstat_open 245
 interfaces for direct probe and process control for another process
 – tnftcl_continue 402
 – tnftcl_exec_open 402

– tnftcl_pid_open 402
 interfaces to query and to change the state of a probe
 – tnftcl_probe_connect 411
 – tnftcl_probe_disable 411
 – tnftcl_probe_disconnect_all 411
 – tnftcl_probe_enable 411
 – tnftcl_probe_state_get 411
 – tnftcl_probe_trace 411
 – tnftcl_probe_untrace 411
 inverse hyperbolic functions
 – acosh 36
 – asinh 36
 – atanh 36
 isencrypt — determine whether a buffer of characters is encrypted 231
 isnan — test for NaN 232
 iterate over probes
 – tnftcl_probe_apply 408
 – tnftcl_probe_apply_ids 408

J

j0 – Bessel functions of the first kind 233
 j1 – Bessel functions of the first kind 233
 jn – Bessel functions of the first kind 233

K

kernel virtual memory functions
 copy data from kernel image or running system – kvm_read, kvm_kread, kvm_uread 255
 get u-area for process – kvm_getu 248
 get entries from kernel symbol table — kvm_nlist 252
 kstat — kernel statistics facility 237
 kstat_chain_update — update the kstat header chain 243
 kstat_close – initialize kernel statistics facility 245
 kstat_data_lookup – find a kstat by name 244
 kstat_lookup – find a kstat by name 244
 kstat_open – initialize kernel statistics facility 245
 kstat_read – read or write kstat data 246
 kstat_write – read or write kstat data 246

- specify a kernel to examine – `kvm_open`,
`kvm_close` 253
- `kstat` — kernel statistics facility 237
- `kstat_chain_update` — update the `kstat` header
chain 243
- `kstat_close` – initialize kernel statistics
facility 245
- `kstat_data_lookup` – find a `kstat` by name 244
- `kstat_lookup` – find a `kstat` by name 244
- `kstat_open` – initialize kernel statistics
facility 245
- `kstat_read` – read or write `kstat` data 246
- `kstat_write` – read or write `kstat` data 246
- `kva_match` — look up a key in a key-value
array 247
- `kvm_close` – specify kernel to examine 253
- `kvm_getcmd` – get invocation arguments for
process 248
- `kvm_getproc` – read system process
structures 250
- `kvm_getu` – get u-area for process 248
- `kvm_kread` – copy data from a kernel image or
running system 255
- `kvm_kwrite` – copy data to a kernel image or
running system 255
- `kvm_nextproc` – read system process
structures 250
- `kvm_nlist` — get entries from kernel symbol
table 252
- `kvm_open` – specify kernel to examine 253
- `kvm_read` – copy data from kernel image or
running system 255
- `kvm_setproc` – read system process
structures 250
- `kvm_uread` – copy data from a kernel image or
running system 255
- `kvm_uwrite` – copy data to a kernel image or
running system 255
- `kvm_write` – copy data to kernel image or
running system 255

L

- layout transformation —
`m_transform_layout` 296

- layout transformation for wide
character strings —
`m_wtransform_layout` 301
- `lgamma` – log gamma function 257
- `lgamma_r` – log gamma function 257
- `libdevinfo` — library of device information
functions 259
- library for TNF probe control in a process or the
kernel — `libtnfctl` 262
- library of device information functions —
`libdevinfo` 259
- `libtnfctl` — library for TNF probe control in a
process or the kernel 262
- load exponent of a radix-independent
floating-point number —
`scalb` 375–376
- `log` — natural logarithm function 269
 - `gamma` 257
 - `gamma_r` 257
 - `lgamma` 257
 - `lgamma_r` 257
- `log10` — base 10 logarithm function 267
- `log1p` — compute natural logarithm 268
- `logb` — radix-independent exponent 270

M

- `m_create_layout` — initialize a layout
object 279
- `m_destroy_layout` — destroy a layout
object 283
- `m_getvalues_layout` — query layout values of a
LayoutObject 291
- `m_setvalues_layout` — set layout values of a
LayoutObject 295
- `m_transform_layout` — layout
transformation 296
- `m_wtransform_layout` — layout transformation
for wide character strings 301
- `maillock` – functions to manage lockfile(s) for
user’s mailbox 271
- `mailunlock` – functions to manage lockfile(s) for
user’s mailbox 271
- Management Interface database administration
functions
 - `DmiAddComponent` 124
 - `DmiAddGroup` 124

- DmiAddLanguage 124
- DmiDeleteComponent 124
- DmiDeleteGroup 124
- DmiDeleteLanguage 124
- Management Interface initialization functions
 - DmiGetConfig 135
 - DmiGetVersion 135
 - DmiRegister 135
 - DmiSetConfig 135
 - DmiUnregister 135
- Management Interface listing functions
 - DmiListAttributes 139
 - DmiListClassNames 139
 - DmiListComponents 139
 - DmiListComponentsByClass 139
 - DmiListGroups 139
 - DmiListLanguages 139
- Management Interface operation functions
 - DmiAddRow 129
 - DmiDeleteRow 129
 - DmiGetAttribute 129
 - DmiGetMultiple 129
 - DmiSetAttribute 129
 - DmiSetMultiple 129
- map a tnftcl error code to a string —
 - tnftcl_strerror 417
- getexecuser - find one entry 220
- math library exception-handling —
 - matherr 273
- mathematical functions
 - gamma 257
 - gamma_r 257
 - lgamma 257
 - lgamma_r 257
- matherr — math library
 - exception-handling 273
- md5 - MD5 hashing functions 281
- MD5 hashing functions - MD5Init 281
- md5_calc - MD5 hashing functions 281
- MD5Final - MD5 hashing functions 281
- MD5Init - MD5 hashing functions 281
- MD5Update - MD5 hashing functions 281
- media_findname — convert a supplied name
 - into an absolute pathname
 - that can be used to access
 - removable media 284
- media_getattr - get and set media
 - attributes 287
- media_setattr - get and set media
 - attributes 287
- memory management
 - copy a file into memory — copylist 63
- mkdirp - create directories in a path 292
- modify/delete user credentials for an
 - authentication service —
 - pam_setcred 332
- mp - multiple precision integer arithmetic 293
- mp_gcd - multiple precision integer
 - arithmetic 293
- mp_itom - multiple precision integer
 - arithmetic 293
- mp_madd - multiple precision integer
 - arithmetic 293
- mp_mcmp - multiple precision integer
 - arithmetic 293
- mp_mdiv - multiple precision integer
 - arithmetic 293
- mp_mfree - multiple precision integer
 - arithmetic 293
- mp_min - multiple precision integer
 - arithmetic 293
- mp_mout - multiple precision integer
 - arithmetic 293
- mp_msub - multiple precision integer
 - arithmetic 293
- mp_mtox - multiple precision integer
 - arithmetic 293
- mp_mult - multiple precision integer
 - arithmetic 293
- mp_pow - multiple precision integer
 - arithmetic 293
- mp_rpow - multiple precision integer
 - arithmetic 293
- mp_xtom - multiple precision integer
 - arithmetic 293
- multiple precision integer arithmetic
 - mp 293
 - mp_gcd 293
 - mp_itom 293
 - mp_madd 293
 - mp_mcmp 293
 - mp_mdiv 293
 - mp_mfree 293

- mp_min 293
- mp_mout 293
- mp_msub 293
- mp_mtox 293
- mp_mult 293
- mp_pow 293
- mp_rpow 293
- mp_xtom 293

N

- natural logarithm function — log 269
- newDmiOctetString — create DmiOctetString in dynamic memory 307
- newDmiString — create DmiString in dynamic memory 308
- next representable double-precision floating-point number — nextafter 309
- nextafter — next representable double-precision floating-point number 309
- NOTE - annotate source code with info for tools 311
 - NOTE vs _NOTE 312
 - NoteInfo Argument 312
- _NOTE - annotate source code with info for tools 311

P

- p2close - close pipes to and from a command 313
- p2open - open pipes to and from a command 313
- PAM — Pluggable Authentication Module 315, 338
 - Administrative Interface 317
 - Interface Overview 315
 - Stacking Multiple Schemes 317
 - Stateful Interface 316
- PAM error messages
 - get string — pam_strerror 356
- PAM routines to maintain module specific state
 - pam_get_data 334
 - pam_set_data 334
- PAM Service Module APIs
 - PAM 338

- pam_acct_mgmt — perform PAM account validation procedures 318
- pam_authenticate — perform authentication within the PAM framework 320
- pam_chauthtok — perform password related functions within the PAM framework 322
- pam_close_session - perform PAM session creation and termination operations 328
- pam_end - authentication transaction routines for PAM 353
- pam_get_data - PAM routines to maintain module specific state 334
- pam_get_item - authentication information routines for PAM 336
- pam_getenv — returns the value for a PAM environment name 324
- pam_getenvlist — returns a list of all the PAM environment variables 325
- pam_open_session - perform PAM session creation and termination operations 328
- pam_putenv — change or add a value to the PAM environment 330
- pam_set_data - PAM routines to maintain module specific state 334
- pam_set_item - authentication information routines for PAM 336
- pam_setcred — modify/delete user credentials for an authentication service 332
- pam_sm — PAM Service Module APIs
 - Interaction with the User 339
 - Interface Overview 338
 - Stateful Interface 339
- pam_sm_acct_mgmt — service provider implementation for pam_acct_mgmt 342
- pam_sm_authenticate — service provider implementation for pam_authenticate 344
- pam_sm_chauthtok — service provider implementation for pam_chauthtok 346

pam_sm_close_session – Service provider implementation for pam_open_session and pam_close_session 349
 pam_sm_open_session – Service provider implementation for pam_open_session and pam_close_session 349
 pam_sm_setcred — service provider implementation for pam_setcred 351
 pam_start – authentication transaction routines for PAM 353
 pathfind — search for named file in named directories 357
 pctx_capture – process context library 359
 pctx_create – process context library 359
 pctx_release – process context library 359
 pctx_run – process context library 359
 pctx_set_events — associate callbacks with process events 361
 perform authentication within the PAM framework — pam_authenticate 320
 perform PAM account validation procedures — pam_acct_mgmt 318
 perform PAM session creation and termination operations
 – pam_close_session 328
 – pam_open_session 328
 perform password related functions within the PAM framework — pam_chauthtok 322
 pipes
 open, close to and from a command – p2open, p2close 313
 Pluggable Authentication Module
 – PAM 315
 pow — power function 364
 power function — pow 364
 print a DmiString
 – printDmiString 367
 print data in DmiAttributeValues list
 – printDmiAttributeValues 365
 print data in input data union
 – printDmiDataUnion 366
 print error in string form

 – dmi_error 134
 printDmiAttributeValues— print data in DmiAttributeValues list 365
 printDmiDataUnion— print data in input data union 366
 printDmiString— print a DmiString 367
 probe insertion interface
 – TNF_DEBUG 426
 – TNF_PROBE_0 426
 – TNF_PROBE_0_DEBUG 426
 – TNF_PROBE_1 426
 – TNF_PROBE_1_DEBUG 426
 – TNF_PROBE_2 426
 – TNF_PROBE_2_DEBUG 426
 – TNF_PROBE_3 426
 – TNF_PROBE_3_DEBUG 426
 – TNF_PROBE_4 426
 – TNF_PROBE_4_DEBUG 426
 – TNF_PROBE_5 426
 – TNF_PROBE_5_DEBUG 426
 process context library – pctx_capture 359
 provide a transient program number
 – reg_ci_callback 369

Q

query layout values of a LayoutObject — m_getvalues_layout 291

R

radix-independent exponent — logb 270
 read and write a disk's VTOC – read_vtoc 368
 – kvm_getproc 250
 – kvm_nextproc 250
 – kvm_setproc 250
 write_vtoc 368
 read or write kstat data
 – kstat_read 246
 – kstat_write 246
 read_vtoc – read and write a disk's VTOC 368
 regexpr – regular expression compile and match routines 370
 register callbacks for probe creation and destruction — tnftcl_register_funcs 416
 regular expression compile and match routines

- advance 370
- compile 370
- regexpr 370
- step 370
- release removable media device reservation —
 - volmgt_release 445
- remainder — remainder function 373
- remainder function — remainder 373
- reserve removable media device —
 - volmgt_acquire 437
- retrieve archive symbol table —
 - elf_getarsym 176
- retrieve class-dependent object file header
 - elf32_getehdr 149
 - elf32_newehdr 149
 - elf64_getehdr 149
 - elf64_newehdr 149
- retrieve class-dependent program header table
 - elf32_getphdr 151
 - elf32_newphdr 151
 - elf64_getphdr 151
 - elf64_newphdr 151
- retrieve class-dependent section header
 - elf32_getshdr 153
 - elf64_getshdr 153
- returns a list of all the PAM
 - environment variables
 - pam_getenvlist 325
- return magnitude of first argument and
 - sign of second argument —
 - copysign 64
- return the size of an object file type
 - elf32_fsize 148
 - elf64_fsize 148
- returns the value for a PAM environment name
 - pam_getenv 324
- return the Volume Management root directory
 - volmgt_root 447
- return whether or not Volume
 - Management is running
 - volmgt_running 448
- returns an unbiased exponent — ilogb 230
- rint — round-to-nearest integral value 374
- rmdirp - remove directories in a path 292
- round-to-nearest integral value — rint 374

S

- scalb — load exponent of a radix-independent
 - floating-point number 375
- scalbn — load exponent of a radix-independent
 - floating-point number 376
- Service Provider functions for components
 - DmiOriginateEvent 145
 - DmiRegisterCi 145
 - DmiUnRegisterCi 145
- service provider implementation for
 - pam_acct_mgmt —
 - pam_sm_acct_mgmt 342
- service provider implementation for
 - pam_authenticate —
 - pam_sm_authenticate 344
- service provider implementation for
 - pam_chauthtok —
 - pam_sm_chauthtok 346
- Service provider implementation for
 - pam_open_session and
 - pam_close_session
 - pam_sm_close_session 349
 - pam_sm_open_session 349
- service provider implementation
 - for pam_setcred —
 - pam_sm_setcred 351
- set layout values of a LayoutObject —
 - m_setvalues_layout 295
- setac - get audit control file information 207
- setauclass - rewind audit_class database
 - file 209
- setauuser - rewind audit_event database
 - file 212
- setauthattr - get authorization database
 - entry 215
- setauuser - get audit_user database entry 218
- setexecattr - get execution attribute database
 - entry 220
- setprofattr - get profile description database
 - entry 224
- setuserattr - get user_attr entry 226
- shell global pattern matching — gmatch 228
- significand — significand function 377
- significand function — significand 377
- simple difference and accumulate operations -
 - cpc_event_diff 81
- sin — sine function 378

sine function — sin 378
 sinh — hyperbolic sine function 379
 sort an ACL — aclsort 31
 sqrt — square root function 380
 square root function — sqrt 380
 SSAAgentIsAlive — Sun Solstice Enterprise Agent registration and communication helper functions 381
 SSAGetTrapPort — Sun Solstice Enterprise Agent registration and communication helper functions 381
 SSAOidCmp — Sun Solstice Enterprise Agent OID helper functions 384
 SSAOidCpy — Sun Solstice Enterprise Agent OID helper functions 384
 SSAOidDup — Sun Solstice Enterprise Agent OID helper functions 384
 SSAOidFree — Sun Solstice Enterprise Agent OID helper functions 384
 SSAOidInit — Sun Solstice Enterprise Agent OID helper functions 384
 SSAOidNew — Sun Solstice Enterprise Agent OID helper functions 384
 SSAOidString — Sun Solstice Enterprise Agent OID helper functions 384
 SSAOidStrToOid — Sun Solstice Enterprise Agent OID helper functions 384
 SSAOidZero — Sun Solstice Enterprise Agent OID helper functions 384
 SSAREgSubagent — Sun Solstice Enterprise Agent registration and communication helper functions 381
 SSAREgSubtable — Sun Solstice Enterprise Agent registration and communication helper functions 381
 SSAREgSubtree — Sun Solstice Enterprise Agent registration and communication helper functions 381
 SSASendTrap — Sun Solstice Enterprise Agent registration and communication helper functions 381
 SSAStrCpy — Sun Solstice Enterprise Agent string helper functions 386
 SSAStrInit — Sun Solstice Enterprise Agent string helper functions 386
 SSAStrToChar — Sun Solstice Enterprise Agent string helper functions 386
 SSAStrZero — Sun Solstice Enterprise Agent string helper functions 386
 SSASubagentOpen — Sun Solstice Enterprise Agent registration and communication helper functions 381
 step — regular expression compile and match routines 370
 strfind — string manipulations 389
 strcadd — copy strings, compressing or expanding C language escape codes 387
 strccpy — copy strings, compressing or expanding C language escape codes 387
 streadd — copy strings, compressing or expanding C language escape codes 387
 STREAMS
 determine whether a buffer of characters is encrypted — isencrypt 231
 read stream up to next delimiter — bgets 49
 split buffer into fields — bufsplit 50
 strecpy — copy strings, compressing or expanding C language escape codes 387
 strfind — string manipulations 389
 string manipulations — strfind 389
 strrspn 389
 strtrns 389
 string manipulations
 — strfind 389
 — strrspn 389
 — strtrns 389
 string operation
 get PAM error message string — pam_strerror 356

strings
 copy, compressing or expanding C
 language escape codes 387

strfind – string manipulations 389

strfind – string manipulations 389

Sun Solstice Enterprise Agent OID helper
 functions

- SSAOidCmp 384
- SSAOidCpy 384
- SSAOidDup 384
- SSAOidFree 384
- SSAOidInit 384
- SSAOidNew 384
- SSAOidString 384
- SSAOidStrToOid 384
- SSAOidZero 384

Sun Solstice Enterprise Agent registration
 and communication helper
 functions

- SSAAgentIsAlive 381
- SSAGetTrapPort 381
- SSARegSubagent 381
- SSARegSubtable 381
- SSARegSubtree 381
- SSASendTrap 381
- SSASubagentOpen 381

Sun Solstice Enterprise Agent string helper
 functions

- SSAStringCpy 386
- SSAStringInit 386
- SSAStringToChar 386
- SSAStringZero 386

Sun::Solaris::Kstat — tied hash interface to the
 kstat facility 234

T

tan — tangent function 390

tangent function — tan 390

tanh — hyperbolic tangent function 391

test access CPU performance counters —
 cpc_access 70

test for NaN — isnan 232

tied hash interface to the kstat facility —
 Sun::Solaris::Kstat 234

TNF_DEBUG – probe insertion interface 426

TNF_PROBE – probe insertion interface

arg_name_n 429

arg_type_n 428

arg_value_n 429

detail 427

keys 427

name 427

TNF_PROBE_0 – probe insertion interface 426

TNF_PROBE_0_DEBUG – probe insertion
 interface 426

TNF_PROBE_1 – probe insertion interface 426

TNF_PROBE_1_DEBUG – probe insertion
 interface 426

TNF_PROBE_2 – probe insertion interface 426

TNF_PROBE_2_DEBUG – probe insertion
 interface 426

TNF_PROBE_3 – probe insertion interface 426

TNF_PROBE_3_DEBUG – probe insertion
 interface 426

TNF_PROBE_4 – probe insertion interface 426

TNF_PROBE_4_DEBUG – probe insertion
 interface 426

TNF_PROBE_5 – probe insertion interface 426

TNF_PROBE_5_DEBUG – probe insertion
 interface 426

tnf_process_disable() – disables probing for the
 process 431

tnf_process_enable() – enables probing for the
 process 431

tnf_thread_disable() – disables probing for the
 calling thread 431

tnf_thread_enable() – enables probing for the
 calling thread 431

tnfctl_buffer_alloc – allocate or deallocate a
 buffer for trace data 392

tnfctl_buffer_dealloc – allocate or deallocate a
 buffer for trace data 392

tnfctl_check_libs – control probes of another
 process where caller provides
 /proc functionality 396

tnfctl_close — close a tnfctl handle 394

tnfctl_continue – interfaces for direct probe and
 process control for another
 process 402

tnfctl_exec_open – interfaces for direct probe
 and process control for
 another process 402

tnftcl_filter_list_add – control kernel tracing and process filtering 421
 tnftcl_filter_list_delete – control kernel tracing and process filtering 421
 tnftcl_filter_list_get – control kernel tracing and process filtering 421
 tnftcl_filter_state_set – control kernel tracing and process filtering 421
 tnftcl_indirect_open – control probes of another process where caller provides /proc functionality 396
 tnftcl_internal_open — create handle for internal process probe control 399
 tnftcl_kernel_open — create handle for kernel probe control 401
 tnftcl_pid_open – interfaces for direct probe and process control for another process 402
 tnftcl_probe_apply – iterate over probes 408
 tnftcl_probe_apply_ids – iterate over probes 408
 tnftcl_probe_connect – interfaces to query and to change the state of a probe 411
 tnftcl_probe_disable – interfaces to query and to change the state of a probe 411
 tnftcl_probe_disconnect_all – interfaces to query and to change the state of a probe 411
 tnftcl_probe_enable – interfaces to query and to change the state of a probe 411
 tnftcl_probe_state_get – interfaces to query and to change the state of a probe 411
 tnftcl_probe_trace – interfaces to query and to change the state of a probe 411
 tnftcl_probe_untrace – interfaces to query and to change the state of a probe 411
 tnftcl_register_funcs — register callbacks for probe creation and destruction 416

tnftcl_strerror — map a tnftcl error code to a string 417
 tnftcl_trace_attrs_get — get the trace attributes from a tnftcl handle 418
 tnftcl_trace_state_set – control kernel tracing and process filtering 421
 touchlock – functions to manage lockfile(s) for user’s mailbox 271
 translate strings to and from events – cpc_eventtostr 90

U

use CPU performance counters on lwps – cpc_rele 71

V

volmgt_acquire — reserve removable media device 437
 volmgt_check — have Volume Management check for media 440
 volmgt_feature_enabled — check whether specific Volume Management features are enabled 442
 volmgt_inuse — check whether or not Volume Management is managing a pathname 443
 volmgt_release — release removable media device reservation 445
 volmgt_root — return the Volume Management root directory 447
 volmgt_running — return whether or not Volume Management is running 448
 volmgt_symdev – convert between Volume Management symbolic names, and the devices that correspond to them 449
 volmgt_symname – convert between Volume Management symbolic names, and the devices that correspond to them 449
 VTOC, disk’s
 read a disk’s VTOC – read_vtoc 368
 write a disk’s VTOC – write_vtoc 368

W

write_vtoc – read and write a disk's VTOC 368

y1 – Bessel functions of the second kind 451

yn – Bessel functions of the second kind 451

Y

y0 – Bessel functions of the second kind 451