



---

## ChorusOS man pages section 2K: Kernel System Calls

---

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

Part No: 806-3326  
December 10, 1999

Copyright 1999 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, ChorusOS, 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 1999 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, ChorusOS, 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 27**

Intro(2K) 33

acap(2K) 52

aconf(2K) 53

acreate(2K) 54

acred(2K) 58

actorCreate(2K) 59

actorDelete(2K) 61

actorName(2K) 62

actorPi(2K) 63

actorPrivilege(2K) 64

actorSelf(2K) 65

actorStop(2K) 66

actorStart(2K) 66

actorStat(2K) 68

actorStop(2K) 69

actorStart(2K) 69

afexec(2K) 71

afexecl(2K) 71

afexecv(2K) 71  
afexecl(2K) 71  
afexecve(2K) 71  
afexeclp(2K) 71  
afexecvp(2K) 71  
afexec(2K) 75  
afexecl(2K) 75  
afexecv(2K) 75  
afexecl(2K) 75  
afexecve(2K) 75  
afexeclp(2K) 75  
afexecvp(2K) 75  
afexec(2K) 79  
afexecl(2K) 79  
afexecv(2K) 79  
afexecl(2K) 79  
afexecve(2K) 79  
afexeclp(2K) 79  
afexecvp(2K) 79  
afexec(2K) 83  
afexecl(2K) 83  
afexecv(2K) 83  
afexecl(2K) 83  
afexecve(2K) 83  
afexeclp(2K) 83  
afexecvp(2K) 83  
afexec(2K) 87  
afexecl(2K) 87

afexecv(2K)	87
afexecl(2K)	87
afexecve(2K)	87
afexeclp(2K)	87
afexecvp(2K)	87
afexec(2K)	91
afexecl(2K)	91
afexecv(2K)	91
afexecl(2K)	91
afexecve(2K)	91
afexeclp(2K)	91
afexecvp(2K)	91
afexec(2K)	95
afexecl(2K)	95
afexecv(2K)	95
afexecl(2K)	95
afexecve(2K)	95
afexeclp(2K)	95
afexecvp(2K)	95
aload(2K)	99
alParamBuild(2K)	99
alParamUnpack(2K)	99
agetalparam(2K)	99
agetId(2K)	102
akill(2K)	103
aload(2K)	104
alParamBuild(2K)	104
alParamUnpack(2K)	104

agetalparam(2K) 104  
aload(2K) 107  
alParamBuild(2K) 107  
alParamUnpack(2K) 107  
agetalparam(2K) 107  
aload(2K) 110  
alParamBuild(2K) 110  
alParamUnpack(2K) 110  
agetalparam(2K) 110  
astart(2K) 113  
astat(2K) 114  
atrace(2K) 115  
await(2K) 118  
awaits(2K) 118  
await(2K) 120  
awaits(2K) 120  
dladdr(2K) 122  
dlclose(2K) 124  
dlerror(2K) 125  
dlopen(2K) 126  
dlsym(2K) 129  
ethIpcStackAttach(2K) 131  
ethOsiStackAttach(2K) 132  
eventInit(2K) 134  
eventClear(2K) 134  
eventPost(2K) 134  
eventWait(2K) 134  
eventInit(2K) 136

eventClear(2K)	136
eventPost(2K)	136
eventWait(2K)	136
eventInit(2K)	138
eventClear(2K)	138
eventPost(2K)	138
eventWait(2K)	138
eventInit(2K)	140
eventClear(2K)	140
eventPost(2K)	140
eventWait(2K)	140
_exit(2K)	142
grpAllocate(2K)	143
grpPortInsert(2K)	145
grpPortRemove(2K)	145
grpPortInsert(2K)	146
grpPortRemove(2K)	146
ipcCall(2K)	147
ipcGetData(2K)	149
ipcReceive(2K)	150
ipcReply(2K)	153
ipcSave(2K)	154
ipcRestore(2K)	154
ipcSave(2K)	155
ipcRestore(2K)	155
ipcSend(2K)	156
ipcSysInfo(2K)	159
ipcTarget(2K)	161

svLapCreate(2K) 163  
lapDescZero(2K) 163  
lapDescIsZero(2K) 163  
lapDescDup(2K) 163  
svLapCreate(2K) 165  
lapDescZero(2K) 165  
lapDescIsZero(2K) 165  
lapDescDup(2K) 165  
svLapCreate(2K) 167  
lapDescZero(2K) 167  
lapDescIsZero(2K) 167  
lapDescDup(2K) 167  
lapInvoke(2K) 169  
svLapBind(2K) 170  
svLapUnbind(2K) 170  
lapResolve(2K) 170  
msgAllocate(2K) 172  
msgFree(2K) 174  
msgGet(2K) 175  
msgPut(2K) 177  
msgRemove(2K) 178  
msgSpaceCreate(2K) 179  
msgSpaceOpen(2K) 181  
mutexInit(2K) 182  
mutexGet(2K) 182  
mutexRel(2K) 182  
mutexTry(2K) 182  
mutexInit(2K) 184

mutexGet(2K)	184
mutexRel(2K)	184
mutexTry(2K)	184
mutexInit(2K)	186
mutexGet(2K)	186
mutexRel(2K)	186
mutexTry(2K)	186
mutexInit(2K)	188
mutexGet(2K)	188
mutexRel(2K)	188
mutexTry(2K)	188
padGet(2K)	190
padKeyCreate(2K)	191
padKeyDelete(2K)	193
padSet(2K)	194
portCreate(2K)	195
portDeclare(2K)	195
portCreate(2K)	197
portDeclare(2K)	197
portDelete(2K)	199
portEnable(2K)	200
portMigrate(2K)	201
portGetSeqNum(2K)	201
portUi(2K)	203
portLi(2K)	203
portMigrate(2K)	204
portGetSeqNum(2K)	204
portPi(2K)	206

portUi(2K) 208  
portLi(2K) 208  
ptdErrnoAddr(2K) 209  
ptdGet(2K) 210  
ptdKeyCreate(2K) 211  
ptdKeyDelete(2K) 213  
ptdRemoteGet(2K) 214  
ptdRemoteSet(2K) 215  
ptdSet(2K) 216  
ptdThreadDelete(2K) 217  
ptdThreadId(2K) 218  
rgnAllocate(2K) 219  
rgnDup(2K) 223  
rgnFree(2K) 225  
rgnInitFromActor(2K) 227  
rgnMapFromActor(2K) 229  
rgnPhysMap(2K) 231  
rgnSetInherit(2K) 233  
rgnSetPaging(2K) 233  
rgnSetOpaque(2K) 233  
rgnSetInherit(2K) 235  
rgnSetPaging(2K) 235  
rgnSetOpaque(2K) 235  
rgnSetInherit(2K) 237  
rgnSetPaging(2K) 237  
rgnSetOpaque(2K) 237  
rgnSetProtect(2K) 239  
rgnStat(2K) 241

rtMutexInit(2K)	244
rtMutexGet(2K)	244
rtMutexRel(2K)	244
rtMutexTry(2K)	244
rtMutexInit(2K)	246
rtMutexGet(2K)	246
rtMutexRel(2K)	246
rtMutexTry(2K)	246
rtMutexInit(2K)	248
rtMutexGet(2K)	248
rtMutexRel(2K)	248
rtMutexTry(2K)	248
rtMutexInit(2K)	250
rtMutexGet(2K)	250
rtMutexRel(2K)	250
rtMutexTry(2K)	250
schedAdmin(2K)	252
semInit(2K)	253
semP(2K)	253
semV(2K)	253
semInit(2K)	255
semP(2K)	255
semV(2K)	255
semInit(2K)	257
semP(2K)	257
semV(2K)	257
svExcHandler(2K)	259
svAbortHandler(2K)	259

svActorAbortHandler(2K) 261  
svActorAbortHandlerConnect(2K) 261  
svActorAbortHandlerDisconnect(2K) 261  
svActorAbortHandlerGetConnected(2K) 261  
svActorAbortHandler(2K) 263  
svActorAbortHandlerConnect(2K) 263  
svActorAbortHandlerDisconnect(2K) 263  
svActorAbortHandlerGetConnected(2K) 263  
svActorAbortHandler(2K) 265  
svActorAbortHandlerConnect(2K) 265  
svActorAbortHandlerDisconnect(2K) 265  
svActorAbortHandlerGetConnected(2K) 265  
svActorAbortHandler(2K) 267  
svActorAbortHandlerConnect(2K) 267  
svActorAbortHandlerDisconnect(2K) 267  
svActorAbortHandlerGetConnected(2K) 267  
svActorExcHandler(2K) 269  
svActorExcHandlerConnect(2K) 269  
svActorExcHandlerDisconnect(2K) 269  
svActorExcHandlerGetConnected(2K) 269  
svActorExcHandler(2K) 272  
svActorExcHandlerConnect(2K) 272  
svActorExcHandlerDisconnect(2K) 272  
svActorExcHandlerGetConnected(2K) 272  
svActorExcHandler(2K) 275  
svActorExcHandlerConnect(2K) 275  
svActorExcHandlerDisconnect(2K) 275  
svActorExcHandlerGetConnected(2K) 275

svActorExcHandler(2K) 278  
svActorExcHandlerConnect(2K) 278  
svActorExcHandlerDisconnect(2K) 278  
svActorExcHandlerGetConnected(2K) 278  
svActorStopHandler(2K) 281  
svActorStopHandlerConnect(2K) 281  
svActorStopHandlerDisconnect(2K) 281  
svActorStopHandlerGetConnected(2K) 281  
svActorStopHandler(2K) 284  
svActorStopHandlerConnect(2K) 284  
svActorStopHandlerDisconnect(2K) 284  
svActorStopHandlerGetConnected(2K) 284  
svActorStopHandler(2K) 287  
svActorStopHandlerConnect(2K) 287  
svActorStopHandlerDisconnect(2K) 287  
svActorStopHandlerGetConnected(2K) 287  
svActorStopHandler(2K) 290  
svActorStopHandlerConnect(2K) 290  
svActorStopHandlerDisconnect(2K) 290  
svActorStopHandlerGetConnected(2K) 290  
svActorVirtualTimeout(2K) 293  
svActorVirtualTimeoutSet(2K) 293  
svActorVirtualTimeoutCancel(2K) 293  
svActorVirtualTimeout(2K) 295  
svActorVirtualTimeoutSet(2K) 295  
svActorVirtualTimeoutCancel(2K) 295  
svActorVirtualTimeout(2K) 297  
svActorVirtualTimeoutSet(2K) 297

svActorVirtualTimeoutCancel(2K) 297  
svCopyIn(2K) 299  
svCopyInString(2K) 299  
svCopyOut(2K) 299  
svCopyIn(2K) 301  
svCopyInString(2K) 301  
svCopyOut(2K) 301  
svCopyIn(2K) 303  
svCopyInString(2K) 303  
svCopyOut(2K) 303  
svExcHandler(2K) 305  
svAbortHandler(2K) 305  
svGetInvoker(2K) 307  
svLapBind(2K) 308  
svLapUnbind(2K) 308  
lapResolve(2K) 308  
svLapCreate(2K) 310  
lapDescZero(2K) 310  
lapDescIsZero(2K) 310  
lapDescDup(2K) 310  
svLapDelete(2K) 312  
svLapBind(2K) 313  
svLapUnbind(2K) 313  
lapResolve(2K) 313  
svMaskAll(2K) 315  
svUnmaskAll(2K) 315  
svUnmask(2K) 315  
svMaskedLockInit(2K) 316

svMaskedLockGet(2K)	316
svMaskedLockTry(2K)	316
svMaskedLockRel(2K)	316
svSpinLockInit(2K)	316
svSpinLockGet(2K)	316
svSpinLockTry(2K)	316
svSpinLockRel(2K)	316
svMaskedLockInit(2K)	319
svMaskedLockGet(2K)	319
svMaskedLockTry(2K)	319
svMaskedLockRel(2K)	319
svSpinLockInit(2K)	319
svSpinLockGet(2K)	319
svSpinLockTry(2K)	319
svSpinLockRel(2K)	319
svMaskedLockInit(2K)	322
svMaskedLockGet(2K)	322
svMaskedLockTry(2K)	322
svMaskedLockRel(2K)	322
svSpinLockInit(2K)	322
svSpinLockGet(2K)	322
svSpinLockTry(2K)	322
svSpinLockRel(2K)	322
svMaskedLockInit(2K)	325
svMaskedLockGet(2K)	325
svMaskedLockTry(2K)	325
svMaskedLockRel(2K)	325
svSpinLockInit(2K)	325

svSpinLockGet(2K) 325  
svSpinLockTry(2K) 325  
svSpinLockRel(2K) 325  
svMemRead(2K) 328  
svMemWrite(2K) 328  
svMemRead(2K) 329  
svMemWrite(2K) 329  
svMsgHandler(2K) 330  
svMsgHdlReply(2K) 330  
svMsgHandler(2K) 333  
svMsgHdlReply(2K) 333  
svPagesAllocate(2K) 336  
svPagesFree(2K) 336  
svPagesAllocate(2K) 338  
svPagesFree(2K) 338  
svMaskedLockInit(2K) 340  
svMaskedLockGet(2K) 340  
svMaskedLockTry(2K) 340  
svMaskedLockRel(2K) 340  
svSpinLockInit(2K) 340  
svSpinLockGet(2K) 340  
svSpinLockTry(2K) 340  
svSpinLockRel(2K) 340  
svMaskedLockInit(2K) 343  
svMaskedLockGet(2K) 343  
svMaskedLockTry(2K) 343  
svMaskedLockRel(2K) 343  
svSpinLockInit(2K) 343

svSpinLockGet(2K)	343
svSpinLockTry(2K)	343
svSpinLockRel(2K)	343
svMaskedLockInit(2K)	346
svMaskedLockGet(2K)	346
svMaskedLockTry(2K)	346
svMaskedLockRel(2K)	346
svSpinLockInit(2K)	346
svSpinLockGet(2K)	346
svSpinLockTry(2K)	346
svSpinLockRel(2K)	346
svMaskedLockInit(2K)	349
svMaskedLockGet(2K)	349
svMaskedLockTry(2K)	349
svMaskedLockRel(2K)	349
svSpinLockInit(2K)	349
svSpinLockGet(2K)	349
svSpinLockTry(2K)	349
svSpinLockRel(2K)	349
svSysCtx(2K)	352
svSysPanic(2K)	353
svSysTimeout(2K)	354
svSysTimeoutSet(2K)	354
svSysTimeoutCancel(2K)	354
svTimeoutGetRes(2K)	354
svSysTimeout(2K)	356
svSysTimeoutSet(2K)	356
svSysTimeoutCancel(2K)	356

svTimeoutGetRes(2K) 356  
svSysTimeout(2K) 358  
svSysTimeoutSet(2K) 358  
svSysTimeoutCancel(2K) 358  
svTimeoutGetRes(2K) 358  
svSysTrapHandler(2K) 360  
svSysTrapHandlerConnect(2K) 360  
svSysTrapHandlerDisconnect(2K) 360  
svSysTrapHandlerGetConnected(2K) 360  
svSysTrapHandler(2K) 362  
svSysTrapHandlerConnect(2K) 362  
svSysTrapHandlerDisconnect(2K) 362  
svSysTrapHandlerGetConnected(2K) 362  
svSysTrapHandler(2K) 364  
svSysTrapHandlerConnect(2K) 364  
svSysTrapHandlerDisconnect(2K) 364  
svSysTrapHandlerGetConnected(2K) 364  
svSysTrapHandler(2K) 366  
svSysTrapHandlerConnect(2K) 366  
svSysTrapHandlerDisconnect(2K) 366  
svSysTrapHandlerGetConnected(2K) 366  
svThreadVirtualTimeout(2K) 368  
svThreadVirtualTimeoutSet(2K) 368  
svThreadVirtualTimeoutCancel(2K) 368  
svThreadVirtualTimeout(2K) 370  
svThreadVirtualTimeoutSet(2K) 370  
svThreadVirtualTimeoutCancel(2K) 370  
svThreadVirtualTimeout(2K) 372

svThreadVirtualTimeoutSet(2K) 372  
svThreadVirtualTimeoutCancel(2K) 372  
svSysTimeout(2K) 374  
svSysTimeoutSet(2K) 374  
svSysTimeoutCancel(2K) 374  
svTimeoutGetRes(2K) 374  
svTrapConnect(2K) 376  
svTrapDisConnect(2K) 376  
svTrapConnect(2K) 378  
svTrapDisConnect(2K) 378  
svMaskAll(2K) 380  
svUnmaskAll(2K) 380  
svUnmask(2K) 380  
svMaskAll(2K) 381  
svUnmaskAll(2K) 381  
svUnmask(2K) 381  
svVirtualTimeoutSet(2K) 382  
svVirtualTimeoutCancel(2K) 382  
svVirtualTimeoutSet(2K) 384  
svVirtualTimeoutCancel(2K) 384  
sysBench(2K) 386  
sysGetConf(2K) 388  
sysGetEnv(2K) 389  
sysLog(2K) 390  
sysRead(2K) 391  
sysWrite(2K) 391  
sysPoll(2K) 391  
sysRead(2K) 392

sysWrite(2K) 392  
sysPoll(2K) 392  
sysReboot(2K) 393  
sysSetEnv(2K) 394  
sysShutdown(2K) 395  
sysTime(2K) 396  
sysTimeGetRes(2K) 396  
sysTime(2K) 397  
sysTimeGetRes(2K) 397  
sysTimer(2K) 398  
sysTimerStartPeriodic(2K) 398  
sysTimerStartFreerun(2K) 398  
sysTimerReadCounter(2K) 398  
sysTimerGetCounterPeriod(2K) 398  
sysTimerGetCounterFrequency(2K) 398  
sysTimerStop(2K) 398  
sysTimer(2K) 400  
sysTimerStartPeriodic(2K) 400  
sysTimerStartFreerun(2K) 400  
sysTimerReadCounter(2K) 400  
sysTimerGetCounterPeriod(2K) 400  
sysTimerGetCounterFrequency(2K) 400  
sysTimerStop(2K) 400  
sysTimer(2K) 402  
sysTimerStartPeriodic(2K) 402  
sysTimerStartFreerun(2K) 402  
sysTimerReadCounter(2K) 402  
sysTimerGetCounterPeriod(2K) 402

sysTimerGetCounterFrequency(2K) 402  
sysTimerStop(2K) 402  
sysTimer(2K) 404  
sysTimerStartPeriodic(2K) 404  
sysTimerStartFreerun(2K) 404  
sysTimerReadCounter(2K) 404  
sysTimerGetCounterPeriod(2K) 404  
sysTimerGetCounterFrequency(2K) 404  
sysTimerStop(2K) 404  
sysTimer(2K) 406  
sysTimerStartPeriodic(2K) 406  
sysTimerStartFreerun(2K) 406  
sysTimerReadCounter(2K) 406  
sysTimerGetCounterPeriod(2K) 406  
sysTimerGetCounterFrequency(2K) 406  
sysTimerStop(2K) 406  
sysTimer(2K) 408  
sysTimerStartPeriodic(2K) 408  
sysTimerStartFreerun(2K) 408  
sysTimerReadCounter(2K) 408  
sysTimerGetCounterPeriod(2K) 408  
sysTimerGetCounterFrequency(2K) 408  
sysTimerStop(2K) 408  
sysTimer(2K) 410  
sysTimerStartPeriodic(2K) 410  
sysTimerStartFreerun(2K) 410  
sysTimerReadCounter(2K) 410  
sysTimerGetCounterPeriod(2K) 410

sysTimerGetCounterFrequency(2K) 410  
sysTimerStop(2K) 410  
sysUnsetEnv(2K) 412  
sysRead(2K) 413  
sysWrite(2K) 413  
sysPoll(2K) 413  
threadAbort(2K) 414  
threadAborted(2K) 414  
threadAbort(2K) 416  
threadAborted(2K) 416  
threadActivate(2K) 418  
threadBind(2K) 419  
threadContext(2K) 421  
threadCreate(2K) 423  
threadDelay(2K) 426  
threadDelete(2K) 427  
threadLoadR(2K) 428  
threadStoreR(2K) 428  
threadName(2K) 430  
threadSuspend(2K) 431  
threadResume(2K) 431  
threadScheduler(2K) 433  
threadSelf(2K) 438  
threadSemInit(2K) 439  
threadSemPost(2K) 439  
threadSemWait(2K) 439  
threadSemInit(2K) 441  
threadSemPost(2K) 441

threadSemWait(2K)	441
threadSemInit(2K)	443
threadSemPost(2K)	443
threadSemWait(2K)	443
threadStart(2K)	445
threadStop(2K)	445
threadStat(2K)	447
threadStart(2K)	449
threadStop(2K)	449
threadLoadR(2K)	451
threadStoreR(2K)	451
threadSuspend(2K)	453
threadResume(2K)	453
threadTimes(2K)	455
timerCreate(2K)	456
timerDelete(2K)	457
timerGetRes(2K)	458
timerSet(2K)	459
timerThreadPoolInit(2K)	460
timerThreadPoolWait(2K)	461
uiBuild(2K)	463
uiClear(2K)	463
uiEqual(2K)	463
uiGetSite(2K)	463
uiIsLocal(2K)	463
uiSite(2K)	463
uiValid(2K)	463
uiBuild(2K)	465

uiClear(2K) 465  
uiEqual(2K) 465  
uiGetSite(2K) 465  
uiIsLocal(2K) 465  
uiSite(2K) 465  
uiValid(2K) 465  
uiBuild(2K) 467  
uiClear(2K) 467  
uiEqual(2K) 467  
uiGetSite(2K) 467  
uiIsLocal(2K) 467  
uiSite(2K) 467  
uiValid(2K) 467  
uiBuild(2K) 469  
uiClear(2K) 469  
uiEqual(2K) 469  
uiGetSite(2K) 469  
uiIsLocal(2K) 469  
uiSite(2K) 469  
uiValid(2K) 469  
uiBuild(2K) 471  
uiClear(2K) 471  
uiEqual(2K) 471  
uiGetSite(2K) 471  
uiIsLocal(2K) 471  
uiSite(2K) 471  
uiValid(2K) 471  
uiLocalSite(2K) 473

uiBuild(2K) 474  
uiClear(2K) 474  
uiEqual(2K) 474  
uiGetSite(2K) 474  
uiIsLocal(2K) 474  
uiSite(2K) 474  
uiValid(2K) 474  
uiBuild(2K) 476  
uiClear(2K) 476  
uiEqual(2K) 476  
uiGetSite(2K) 476  
uiIsLocal(2K) 476  
uiSite(2K) 476  
uiValid(2K) 476  
univTime(2K) 478  
univTimeSet(2K) 478  
univTimeAdjust(2K) 478  
univTime(2K) 480  
univTimeSet(2K) 480  
univTimeAdjust(2K) 480  
univTime(2K) 482  
univTimeSet(2K) 482  
univTimeAdjust(2K) 482  
virtualTimeGetRes(2K) 484  
vmCopy(2K) 485  
vmFree(2K) 486  
vmLock(2K) 487  
vmUnLock(2K) 487

vmPageSize(2K)	489
vmPhysAddr(2K)	490
vmSetPar(2K)	491
vmStat(2K)	493
vmLock(2K)	494
vmUnLock(2K)	494
<b>Index</b>	<b>495</b>

# PREFACE

---

---

## Overview

A man page is provided for both the naive user, and sophisticated user who is familiar with the ChorusOS™ operating system and is in need of on-line information. 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.

The following is a list of sections in the ChorusOS man pages and the information it references:

- *Section 1CC: User Utilities; Host and Target Utilities*
- *Section 1M: System Management Utilities*
- *Section 2DL: System Calls; Data Link Services*
- *Section 2K: System Calls; Kernel Services*
- *Section 2MON: System Calls; Monitoring Services*
- *Section 2POSIX: System Calls; POSIX System Calls*
- *Section 2RESTART: System Calls; Hot Restart and Persistent Memory*
- *Section 2SEG: System Calls; Virtual Memory Segment Services*
- *Section 3FTPD: Libraries; FTP Daemon*
- *Section 3M: Libraries; Mathematical Libraries*
- *Section 3POSIX: Libraries; POSIX Library Functions*
- *Section 3RPC: Libraries; RPC Services*
- *Section 3STDC: Libraries; Standard C Library Functions*
- *Section 3TELD: Libraries; Telnet Services*
- *Section 4CC: Files*

- *Section 5FEA: ChorusOS Features and APIs*
- *Section 7P: Protocols*
- *Section 7S: Services*
- *Section 9DDI: Device Driver Interfaces*
- *Section 9DKI: Driver to Kernel Interface*
- *Section 9DRV: Driver Implementations*

ChorusOS man pages are grouped in Reference Manuals, with one reference manual per section.

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 pathname 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"> <li>[ ]     The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.</li> <li>. . .    Ellipses. Several values may be provided for the previous argument, or the previous argument can be specified multiple times, for example, 'filename . . . '.</li> <li>         Separator. Only one of the arguments separated by this character can be specified at time.</li> <li>{ }     Braces. The options and/or arguments enclosed within braces are</li> </ul>

interdependent, such that everything enclosed must be treated as a unit.

FEATURES	This section provides the list of features which offer an interface. An API may be associated with one or more system features. The interface will be available if one of the associated features has been configured.
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, functions and such, are described under USAGE.
OPTIONS	This 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 is provided as a guidance on use. This section lists special rules, features and commands that require in-depth explanations. The subsections listed below are used to explain built-in functionality:
EXAMPLES	<p>Commands Modifiers Variables Expressions Input Grammar</p> <p>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 <code>example%</code> or if the user must be superuser, <code>example#</code>. Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS and USAGE sections.</p>
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 filenames 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.
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.

# System Calls

<b>NAME</b>	Intro – introduction to ChorusOS error codes and system calls																																				
<b>DESCRIPTION</b>	This manual describes the principal ChorusOS APIs. All system calls and other interfaces provided for use by applications are documented here. In this context “applications” covers a wide range of functions including device drivers, OS personality emulation servers, real time control programs, and so forth.																																				
<b>FEATURES</b>	<p>Each API function is associated with one or more system features. A given interface is available if, and only if, one of its associated features was configured on the target system when that system was built.</p> <p>Each manual page has a FEATURES section, where the list of features offering this interface is provided.</p> <p>The features available in ChorusOS are introduced here. For each feature, there is a man page in the 5FEA section which gives an description of the feature and a list of all interfaces associated with the feature.</p> <table border="0"> <tr> <td>DATE</td> <td>time of day service</td> </tr> <tr> <td>EVENT</td> <td>event flag sets</td> </tr> <tr> <td>HOT_RESTART</td> <td>management of restartable actors and persistent memory</td> </tr> <tr> <td>IPC</td> <td>provide OSI stack entry points</td> </tr> <tr> <td>IPC_REMOTE</td> <td>inter-process communication</td> </tr> <tr> <td>IPC_REMOTE_COMM</td> <td>IPC remote communication</td> </tr> <tr> <td>LAPBIND</td> <td>built-in nameserver for local access points</td> </tr> <tr> <td>LAPSAFE</td> <td>safe mode for LAP invocations</td> </tr> <tr> <td>LOG</td> <td>system logging</td> </tr> <tr> <td>VIRTUAL_ADDRESS_SPACE</td> <td>memory management features</td> </tr> <tr> <td>ON_DEMAND_PAGING</td> <td>memory management features</td> </tr> <tr> <td>MIPC</td> <td>message queues</td> </tr> <tr> <td>MON</td> <td>system monitoring</td> </tr> <tr> <td>PERF</td> <td>performance support</td> </tr> <tr> <td>RTC</td> <td>realtime clock</td> </tr> <tr> <td>RTMUTEX</td> <td>real-time mutexes</td> </tr> <tr> <td>ROUND_ROBIN</td> <td>scheduler features</td> </tr> <tr> <td>TIMER</td> <td>system time</td> </tr> </table>	DATE	time of day service	EVENT	event flag sets	HOT_RESTART	management of restartable actors and persistent memory	IPC	provide OSI stack entry points	IPC_REMOTE	inter-process communication	IPC_REMOTE_COMM	IPC remote communication	LAPBIND	built-in nameserver for local access points	LAPSAFE	safe mode for LAP invocations	LOG	system logging	VIRTUAL_ADDRESS_SPACE	memory management features	ON_DEMAND_PAGING	memory management features	MIPC	message queues	MON	system monitoring	PERF	performance support	RTC	realtime clock	RTMUTEX	real-time mutexes	ROUND_ROBIN	scheduler features	TIMER	system time
DATE	time of day service																																				
EVENT	event flag sets																																				
HOT_RESTART	management of restartable actors and persistent memory																																				
IPC	provide OSI stack entry points																																				
IPC_REMOTE	inter-process communication																																				
IPC_REMOTE_COMM	IPC remote communication																																				
LAPBIND	built-in nameserver for local access points																																				
LAPSAFE	safe mode for LAP invocations																																				
LOG	system logging																																				
VIRTUAL_ADDRESS_SPACE	memory management features																																				
ON_DEMAND_PAGING	memory management features																																				
MIPC	message queues																																				
MON	system monitoring																																				
PERF	performance support																																				
RTC	realtime clock																																				
RTMUTEX	real-time mutexes																																				
ROUND_ROBIN	scheduler features																																				
TIMER	system time																																				

**POSIX-compliant  
API features**

SEM	semaphores
USER_MODE	support for user actors
VTIMER	thread execution timing
The following features provide POSIX-compatible functions (1003.1, 1003.1b, 1003.1c).	
POSIX-THREADS	POSIX 1003.1c <code>pthread</code> functions.
POSIX-TIMERS	POSIX 1003.1b real-time clock and timer functions.
POSIX_MQ	POSIX 1003.1b real-time message queues.
POSIX_SHM	POSIX 1003.1b real-time shared memory objects.
POSIX_SOCKETS	POSIX 1003.1j (standard BSD compatible) socket functions.

---

The POSIX 1003.1c `pthread` functions and 1003.1b real-time clock and timer functions are delivered into the standard `libcxx.a` library.

---

**ERRORS**

**Warning:** For reasons of compatibility with existing standards, different system calls use distinct error return conventions. Always refer to individual man pages when interpreting error codes.

Most ChorusOS system calls are of the integer type and return a direct value which is zero or positive when the call is successful, and negative to indicate an error. The possible error return values are listed below.

The PRIVATE-DATA calls follow a similar convention (error code returned directly by the function) but the error codes are positive, not negative. PRIVATE-DATA error values are indicated below.

The POSIX-THREADS functions whose names begin with `pthread_` follow a similar convention, but with POSIX error return values.

All other POSIX features (POSIX-TIMERS, POSIX\_MQ, POSIX\_SHM, POSIX\_SOCKETS, remaining POSIX-THREADS calls) and in addition the ChorusOS ACTOR\_EXTENDED\_MNGT, return `-1` in case of error and set the `errno` variable. Note that `errno` is not cleared on successive calls, so it should be tested only after an error has been indicated via a function return code. In multi-threaded actors, `errno` is implemented not as a global variable but as a macro which refers to a per-thread `errno` value. A complete list of the POSIX error codes (possible `errno` values) can be found in the file `<errno.h>`.

**Standard ChorusOS  
error codes**

The list of error codes is given below, with general meanings. However, this general definition may be overwritten by certain system calls. The meaning should be interpreted according to the type and the circumstances of the call.

When a thread that executes in supervisor address space performs a system call, the arguments of the call are not fully checked by the system. This applies to threads executing in supervisor actors and to threads executing in user actors when the VIRTUAL\_ADDRESS\_SPACE feature is disabled. In general, when such a thread invokes a system call with wrong arguments the result may be unpredictable.

The implication is that when a man page states that a system call will return a given error code in a given error situation, this will be fully enforced by the system only for threads executing in user address space. For threads executing in supervisor address space, the error situations may not be detected by the system. In particular this is the case for the K\_EFAULT error situations.

- 0 K\_OK No error  
The call returned successfully.
- 1 K\_EINVAL Invalid argument  
An invalid argument has been provided.
- 2 K\_ENOMEM Out of resources  
The kernel has been unable to perform the call due to a lack of resources.
- 3 K\_ETOOMUCH Message too big  
Attempted to send a message body longer than K\_CMSGsizEMAX bytes.
- 4 K\_ENOPORT Port invalid  
The port local identifier is invalid (the port was never created, or was deleted, or has migrated before or during the call).
- 5 K\_EFAULT Bad address  
The kernel encountered a hardware access fault in attempting to access one of the arguments of the call.
- 6 K\_ENOTIMP Not implemented  
The call, or one of its options, is not implemented in the current version.
- 7 K\_EUNKNOWN Unreachable destination  
Attempt to send an RPC request to an unreachable port.
- 8 K\_ETIMEOUT Time-out occurred  
A system call failed because the time-out occurred.
- 9 K\_EFAIL Transaction failed  
A system call failed because the actor which received the request failed.
- 10 K\_EABORT Aborted wait

- Abortable calls (for example. *ipcReceive*, *ipcCall*, *threadSuspend*) have been interrupted because the waiting thread has been aborted.
- 11 **K\_EFULL** System queues saturated  
An *ipcSend*, *ipcReply* or *ipcCall* failed because the local destination port queue is full, or the local communication system is saturated.
  - 12 **K\_EROUND** Alignment error  
Attempted to allocate a memory region whose starting address is not aligned to a page boundary.
  - 13 **K\_EADDR** The address of an argument is invalid  
Attempted to allocate a memory region overlapping another.
  - 14 **K\_EPRIV** Privilege violation  
Attempted to perform a forbidden operation.
  - 15 **K\_EBADMODE** Invalid addressing mode  
Attempted to send a message with an invalid addressing mode.
  - 16 **K\_ESIZE** Size error  
The size given in the parameter is incompatible with other parameters.
  - 17 **K\_EBUSY** Memory locking conflict  
Attempted to flush a part of a local cache locked in memory.
  - 18 **K\_EABORTRPC** Aborted RPC  
Aborted a remote IPC notification.
  - 19 **K\_EMAPPER** Mapper access error  
Unable to access a mapper.
  - 20 **K\_ENOEMPTY** Non empty address space  
Attempted to duplicate regions within a non-empty address space.
  - 21 **K\_EOFFSET** Bad segment offset  
Attempted to map an impossible part of a segment.
  - 22 **K\_EOVERLAP** Region overlapping  
Attempted to allocate a memory region that would overlap with another one.
  - 23 **K\_EPROT** Region access rights violation  
Attempted to write in a read-only region.
  - 24 **K\_ESPACE** Address space violation  
Attempted to allocate a region outside the valid range of virtual addresses for the address space.
  - 25 **K\_EUNDEF** Undefined default mapper  
The system default mapper is not defined.

- 26 **K\_EARGS** Arguments inconsistent  
The combination of a number of the call arguments is inconsistent.
- 28 **K\_EBUSYPORT** Port in use  
Attempted to migrate a port to which a message handler is attached.
- 29 **K\_ELINKFAILURE** Link failure  
A remote communication protocol failure has been detected during remote IPC.
- 30 **K\_EBADMSG** Invalid message identifier  
The message identifier parameter is not a valid message address, or the corresponding message has already been consumed.
- 31 **K\_ENOTAVAILABLE** System call not available  
Attempted to invoke a system call that is not available within this configuration.

**PRIVATE-DATA error codes**

- 0 **K\_OK** No error  
The call returned successfully.
- 1 **PD\_ENOKEY** No more keys available  
The key table is full, no more keys can be created.
- 2 **PD\_ENOMEM** Out of resources  
The kernel has been unable to perform the call because of a lack of resources.
- 3 **PD\_EINVAL** Invalid argument  
An invalid argument has been given.
- 4 **PD\_ESERVER** Cannot connect to server  
The Private Data Manager is not accessible.
- 5 **PD\_EUSER** Called from user mode

**DEFINITIONS AND ARGUMENT TYPES**

All system call arguments are integer or pointer sized. In other words, if an argument is described as "a structure ...", it is actually "a pointer to a structure ...". The following argument types may occur in system calls exported by the Core Executive or by any other module, with the semantics indicated and constant option values.

*KnCap\* actorcap*

Usually, *actorcap* is a pointer to a user memory location which either contains an actor capability (input parameter) or will hold an actor capability (output parameter). As an input parameter,

*actorcap* may contain one of the following special values instead:

K\_MYACTOR      Designates the home actor of the calling thread.

K\_SVACTOR      Designates the supervisor address space.

*KnThreadLid threadli*

Usually, *threadli* is a thread local identifier. Instead, *threadli* may contain the special value K\_MYSELF, which designates the calling thread.

*KnTimeVal\* waitLimit*

This argument is used in system calls that make the calling thread block until an explicit event or action occurs. It allows the caller to control the maximum waiting time and abortability of the wait. Usually, *waitLimit* is a pointer to a user memory location containing a timeout interval (expressed as a (second,nanosecond) pair, see *sysTime(2K)*). In this case, if the call blocks, it is ABORTABLE. Instead, *waitLimit* may contain one of the following special values:

K\_NOBLOCK

Specifies a polling call. If the relevant condition is not satisfied immediately, the call returns with an error code and does not wait. The call is NONABORTABLE in the sense that it neither consumes the abort state (if the current thread was previously aborted) nor returns K\_EABORT. For compatibility with previous versions of the system K\_NOBLOCK has the value 0.

K\_NOTIMEOUT

Specifies an unlimited wait. The call is ABORTABLE. For compatibility with previous versions of the system K\_NOTIMEOUT has the value -1.

K\_NOTIMEOUT\_NOABORT

Also specifies an unlimited wait. The call is NONABORTABLE.

## Site number

The Chorus kernel has the concept of "Site Number" which has the following properties

- There is a unique Site Number per site.
- The site number is a 32 bit unsigned integer
- The site Number is embedded into all structures of the type "KnUniqueId" which are associated to the following kernel entities:
  - IPC ports
  - IPC Group Capabilities (KnCap)
  - Actor Capabilities (KnCap)

The site number of the target is provided to the kernel when you boot the system and can be provided in two different ways:

1. By the boot program within the "siteNumber" field of the "bootConf" structure; in this case, it is the responsibility of the boot program to determine the site number of the target. The way the site number is found by the boot is fully "boot dependent", and may be specific to the target board. For instance, it may be stored in the "NVRAM" memory, or computed more dynamically from some (unique) board identifier. When the target is booted with the standard ChorusOS network boot monitor, it is the whole IP address used by the boot monitor which is provided as the site number of the target.
2. Within the boot image through the "chorusSiteId" kernel tunable; in this case, the site number is statically fixed within the boot image. This approach is less flexible than the case above, as the same boot image cannot be booted on multiple similar target boards. For this reason, it should only be used when there is no way for the site number of the target board to be determined when you boot the system.

---

**Note -**

- the value fixed through the tunable takes precedence over the value provided by the boot program.
  - by default, the site number is set to zero. If the REMOTE IPC feature of the kernel has been configured, a WARNING message is displayed on the console and the REMOTE IPC feature is not activated. In other words, only local ICP communications are enabled when the site number has not been set.
- 

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

Name	Description
<code>_exit(2K)</code>	terminate a <code>c_actor</code>
<code>acap(2K)</code>	get a <code>c_actor</code> capability
<code>aconf(2K)</code>	get configurable system variables
<code>acreate(2K)</code>	creates a <code>c_actor</code>
<code>acred(2K)</code>	get/set <code>c_actor</code> credentials
<code>actorCreate(2K)</code>	create an actor
<code>actorDelete(2K)</code>	delete an actor
<code>actorName(2K)</code>	get and/or set the symbolic name of an actor
<code>actorPi(2K)</code>	get and/or set the protection identifier of an actor
<code>actorPrivilege(2K)</code>	get and/or set the privilege of an actor
<code>actorSelf(2K)</code>	get the current actor capability

actorStart(2K)	See actorStop(2K)
actorStat(2K)	get the status of all actors on a site
actorStop(2K)	Stop an actor; Start an actor
afexec(2K)	create a new c_actor
afexecl(2K)	See afexec(2K)
afexecle(2K)	See afexec(2K)
afexeclp(2K)	See afexec(2K)
afexecv(2K)	See afexec(2K)
afexecve(2K)	See afexec(2K)
afexecvp(2K)	See afexec(2K)
agetId(2K)	get c_actor's ID
agetalparam(2K)	See aload(2K)
akill(2K)	kill or restart an actor
alParamBuild(2K)	See aload(2K)
alParamUnpack(2K)	See aload(2K)
aload(2K)	Manage the loading of an actor
astart(2K)	activates a c_actor
astat(2K)	list all active c_actors
atrace(2K)	c_actor trace
await(2K)	wait for c_actor to terminate or stop
awaits(2K)	See await(2K)
dladdr(2K)	translate address to symbolic information
dlclose(2K)	close a dynamic object
dlerror(2K)	get diagnostic information
dlopen(2K)	gain access to a dynamic object file
dlsym(2K)	get the address of a symbol in a dynamic object
ethIpcStackAttach(2K)	attach an

ethOsiStackAttach(2K)	attach an
eventClear(2K)	See eventInit(2K)
eventInit(2K)	initialize an event flag set; Clear events in an event flag set; Post events in an event flag set; Wait for events in an event flag set
eventPost(2K)	See eventInit(2K)
eventWait(2K)	See eventInit(2K)
grpAllocate(2K)	allocate a port group capability
grpPortInsert(2K)	insert a port into a port group; remove a port from a port group
grpPortRemove(2K)	See grpPortInsert(2K)
ipcCall(2K)	send an RPC request and wait for the reply
ipcGetData(2K)	get the current message body
ipcReceive(2K)	receive a message
ipcReply(2K)	reply to the current message
ipcRestore(2K)	See ipcSave(2K)
ipcSave(2K)	Save the current message; Restore a saved message as the current message
ipcSend(2K)	send a message
ipcSysInfo(2K)	get system information about the current message
ipcTarget(2K)	build a message target
lapDescDup(2K)	See svLapCreate(2K)
lapDescIsZero(2K)	See svLapCreate(2K)
lapDescZero(2K)	See svLapCreate(2K)
lapInvoke(2K)	invoke a lap handler
lapResolve(2K)	See svLapBind(2K)
msgAllocate(2K)	allocates a message from a message space

msgFree(2K)	free a message of a message space
msgGet(2K)	retrieves the first message of a message queue
msgPut(2K)	post a message to a message queue
msgRemove(2K)	remove a message from a message queue
msgSpaceCreate(2K)	create a message space
msgSpaceOpen(2K)	open a message space
mutexGet(2K)	See mutexInit(2K)
mutexInit(2K)	initialize a mutex; acquire a mutex; release a mutex; try to acquire a mutex
mutexRel(2K)	See mutexInit(2K)
mutexTry(2K)	See mutexInit(2K)
padGet(2K)	return actor-specific values associated with keys
padKeyCreate(2K)	create a private key for an actor
padKeyDelete(2K)	delete an actor private key
padSet(2K)	set the actor's key to a specific value
portCreate(2K)	create a port; declare a port
portDeclare(2K)	See portCreate(2K)
portDelete(2K)	delete a port
portEnable(2K)	Enable a port portDisable; Disable a port
portGetSeqNum(2K)	See portMigrate(2K)
portLi(2K)	See portUi(2K)
portMigrate(2K)	Migrate a port; Get a port sequence number
portPi(2K)	get and/or set the protection identifier of a port
portUi(2K)	Get the unique identifier of a port, given its local identifier; Get the local

	identifier of a port, given its unique identifier
ptdErrnoAddr(2K)	return a thread-specific errno address
ptdGet(2K)	return thread-specific value associated with key
ptdKeyCreate(2K)	create a thread-specific data key
ptdKeyDelete(2K)	delete a thread-specific data key
ptdRemoteGet(2K)	return a thread-specific data value for another thread
ptdRemoteSet(2K)	set a thread-specific data value for another thread
ptdSet(2K)	set a thread-specific value
ptdThreadDelete(2K)	delete all thread-specific values and call destructors
ptdThreadId(2K)	return the thread ID
rgnAllocate(2K)	allocate a region in an actor address space
rgnDup(2K)	duplicate an actor address space
rgnFree(2K)	deallocate regions of an actor address space
rgnInitFromActor(2K)	allocate a region in an actor address space and initialise it from another region
rgnMapFromActor(2K)	create a region in an actor address space and map another region to it
rgnPhysMap(2K)	create a region in an actor address space and map (on demand) to it physical memory specified by the caller
rgnSetInherit(2K)	Change inheritance options associated with a region; Change paging options associated with a region; Change inheritance options associated with a region; Change

	opaque values associated with a region
rgnSetOpaque(2K)	See rgnSetInherit(2K)
rgnSetPaging(2K)	See rgnSetInherit(2K)
rgnSetProtect(2K)	change protection options associated with a region
rgnStat(2K)	get the statistics of a region of an actor address space
rtMutexGet(2K)	See rtMutexInit(2K)
rtMutexInit(2K)	Initialize a realtime mutex; Acquire a realtime mutex; Release a realtime mutex; Try to acquire a realtime mutex
rtMutexRel(2K)	See rtMutexInit(2K)
rtMutexTry(2K)	See rtMutexInit(2K)
schedAdmin(2K)	scheduling classes administration
semInit(2K)	initialize a semaphore; wait on a semaphore; signal a semaphore
semP(2K)	See semInit(2K)
semV(2K)	See semInit(2K)
svAbortHandler(2K)	See svExcHandler(2K)
svActorAbortHandler(2K)	Connect an actor abort handler; Disconnect an actor abort handler; Get an actor abort handler
svActorAbortHandlerConnect(2K)	See svActorAbortHandler(2K)
svActorAbortHandlerDisconnect(2K)	See svActorAbortHandler(2K)
svActorAbortHandlerGetConnected(2K)	See svActorAbortHandler(2K)
svActorExcHandler(2K)	Connect an actor exception handler; Disconnect an actor exception handler; Get an actor exception handler

svActorExcHandlerConnect(2K)	See svActorExcHandler(2K)
svActorExcHandlerDisconnect(2K)	See svActorExcHandler(2K)
svActorExcHandlerGetConnected(2K)	See svActorExcHandler(2K)
svActorStopHandler(2K)	Actor stop handler management: Connect an actor stop handler; Disconnect an actor stop handler; Get an actor stop handler
svActorStopHandlerConnect(2K)	See svActorStopHandler(2K)
svActorStopHandlerDisconnect(2K)	See svActorStopHandler(2K)
svActorStopHandlerGetConnected(2K)	See svActorStopHandler(2K)
svActorVirtualTimeout(2K)	Set an actor's virtual timeout; Cancel an actor's virtual timeout
svActorVirtualTimeoutCancel(2K)	See svActorVirtualTimeout(2K)
svActorVirtualTimeoutSet(2K)	See svActorVirtualTimeout(2K)
svCopyIn(2K)	Copy from trap caller space; Copy string from trap caller space; Copy to trap caller space
svCopyInString(2K)	See svCopyIn(2K)
svCopyOut(2K)	See svCopyIn(2K)
svExcHandler(2K)	Define an exception handler; Define an abort handler
svGetInvoker(2K)	get handler invoker
svLapBind(2K)	bind a symbolic name to a lap descriptor; unbind the symbolic name bound to a lap descriptor; get a lap descriptor from a lap symbolic name
svLapCreate(2K)	create a lap; reset a lap descriptor; test if a lap descriptor has been initialized; duplicate a lap descriptor

svLapDelete(2K)	delete a lap
svLapUnbind(2K)	See svLapBind(2K)
svMaskAll(2K)	Disable interrupt processing; Enable interrupt processing; Reenable interrupt processing
svMaskedLockGet(2K)	See svMaskedLockInit(2K)
svMaskedLockInit(2K)	Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock
svMaskedLockRel(2K)	See svMaskedLockInit(2K)
svMaskedLockTry(2K)	See svMaskedLockInit(2K)
svMemRead(2K)	Copy from supervisor caller space; Copy to supervisor caller space
svMemWrite(2K)	See svMemRead(2K)
svMsgHandler(2K)	Connect/disconnect a message handler; Prepare a reply to a handled message
svMsgHdlReply(2K)	See svMsgHandler(2K)
svPagesAllocate(2K)	supervisor address space memory allocator
svPagesFree(2K)	See svPagesAllocate(2K)
svSpinLockGet(2K)	See svMaskedLockInit(2K)
svSpinLockInit(2K)	See svMaskedLockInit(2K)
svSpinLockRel(2K)	See svMaskedLockInit(2K)
svSpinLockTry(2K)	See svMaskedLockInit(2K)
svSysCtx(2K)	get system context table address
svSysPanic(2K)	trigger the invocation of the panic handler

<code>svSysTimeout(2K)</code>	Request a timeout; Cancel a timeout; Get timeout resolution
<code>svSysTimeoutCancel(2K)</code>	See <code>svSysTimeout(2K)</code>
<code>svSysTimeoutSet(2K)</code>	See <code>svSysTimeout(2K)</code>
<code>svSysTrapHandler(2K)</code>	Connect a trap handler; Disconnect a trap handler; Get a trap handler
<code>svSysTrapHandlerConnect(2K)</code>	See <code>svSysTrapHandler(2K)</code>
<code>svSysTrapHandlerDisconnect(2K)</code>	See <code>svSysTrapHandler(2K)</code>
<code>svSysTrapHandlerGetConnected(2K)</code>	See <code>svSysTrapHandler(2K)</code>
<code>svThreadVirtualTimeout(2K)</code>	Set a thread's virtual timeout; Cancel a thread's virtual timeout
<code>svThreadVirtualTimeoutCancel(2K)</code>	See <code>svThreadVirtualTimeout(2K)</code>
<code>svThreadVirtualTimeoutSet(2K)</code>	See <code>svThreadVirtualTimeout(2K)</code>
<code>svTimeoutGetRes(2K)</code>	See <code>svSysTimeout(2K)</code>
<code>svTrapConnect(2K)</code>	Connect a trap handler; Disconnect a trap handler
<code>svTrapDisConnect(2K)</code>	See <code>svTrapConnect(2K)</code>
<code>svUnmask(2K)</code>	See <code>svMaskAll(2K)</code>
<code>svUnmaskAll(2K)</code>	See <code>svMaskAll(2K)</code>
<code>svVirtualTimeoutCancel(2K)</code>	See <code>svVirtualTimeoutSet(2K)</code>
<code>svVirtualTimeoutSet(2K)</code>	Set a virtual timeout; Cancel a virtual timeout
<code>sysBench(2K)</code>	kernel benchmark utility
<code>sysGetConf(2K)</code>	Get Chorus module configuration value
<code>sysGetEnv(2K)</code>	Get a value from the Chorus configuration environment
<code>sysLog(2K)</code>	log a message in the kernel's cyclical buffer

<code>sysPoll(2K)</code>	See <code>sysRead(2K)</code>
<code>sysRead(2K)</code>	Read characters from the system console; Write characters to the system console; Poll characters from the system console
<code>sysReboot(2K)</code>	request a reboot of the local site
<code>sysSetEnv(2K)</code>	Set a value in the ChorusOS configuration environment
<code>sysShutdown(2K)</code>	shut down or restart the system
<code>sysTime(2K)</code>	get system time; get system time resolution
<code>sysTimeGetRes(2K)</code>	See <code>sysTime(2K)</code>
<code>sysTimer(2K)</code>	system timer management
	See <code>sysTimer(2K)</code>
<code>sysTimerGetCounterFrequency(2K)</code>	
<code>sysTimerGetCounterPeriod(2K)</code>	See <code>sysTimer(2K)</code>
<code>sysTimerReadCounter(2K)</code>	See <code>sysTimer(2K)</code>
<code>sysTimerStartFreerun(2K)</code>	See <code>sysTimer(2K)</code>
<code>sysTimerStartPeriodic(2K)</code>	See <code>sysTimer(2K)</code>
<code>sysTimerStop(2K)</code>	See <code>sysTimer(2K)</code>
<code>sysUnsetEnv(2K)</code>	delete a value from the ChorusOS configuration environment
<code>sysWrite(2K)</code>	See <code>sysRead(2K)</code>
<code>threadAbort(2K)</code>	Abort a thread; Check whether the current thread has been aborted
<code>threadAborted(2K)</code>	See <code>threadAbort(2K)</code>
<code>threadActivate(2K)</code>	make a thread active
<code>threadBind(2K)</code>	bind a thread to a processor
<code>threadContext(2K)</code>	get and/or set the context of a thread
<code>threadCreate(2K)</code>	create a thread
<code>threadDelay(2K)</code>	delay the current thread

threadDelete(2K)	delete a thread
threadLoadR(2K)	Get the current thread's valid soft register value; Reset the current thread's valid soft register value
threadName(2K)	get and/or set the symbolic name of a thread
threadResume(2K)	See threadSuspend(2K)
threadScheduler(2K)	get and/or set thread scheduling parameters
threadSelf(2K)	get the current thread local identifier
threadSemInit(2K)	Initialize a thread semaphore; Signal a thread semaphore; Wait on a thread semaphore
threadSemPost(2K)	See threadSemInit(2K)
threadSemWait(2K)	See threadSemInit(2K)
threadStart(2K)	Stop a thread; Start a thread
threadStat(2K)	obtain the descriptions of the threads running in an actor
threadStop(2K)	See threadStart(2K)
threadStoreR(2K)	See threadLoadR(2K)
threadSuspend(2K)	Suspend a thread; Resume a thread
threadTimes(2K)	get thread execution time
timerCreate(2K)	create a timer
timerDelete(2K)	delete a timer
timerGetRes(2K)	get the timer resolution
timerSet(2K)	start, cancel or query a timer
timerThreadPoolInit(2K)	initialize a timer thread pool
timerThreadPoolWait(2K)	wait for a timer expiration event
uiBuild(2K)	Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a

	unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared
uiClear(2K)	See uiBuild(2K)
uiEqual(2K)	See uiBuild(2K)
uiGetSite(2K)	See uiBuild(2K)
uiIsLocal(2K)	See uiBuild(2K)
uiLocalSite(2K)	get the local site number
uiSite(2K)	See uiBuild(2K)
uiValid(2K)	See uiBuild(2K)
univTime(2K)	Get time-of-day; Set time-of-day; Adjust time-of-day univTimeGetRes; Get time-of-day resolution
univTimeAdjust(2K)	See univTime(2K)
univTimeSet(2K)	See univTime(2K)
virtualTimeGetRes(2K)	get virtual time resolution
vmCopy(2K)	copy data between actor address spaces
vmFree(2K)	free memory
vmLock(2K)	Fix data in memory; Data in memory
vmPageSize(2K)	get the minimum allocatable memory block size
vmPhysAddr(2K)	get the physical address corresponding to a virtual address
vmSetPar(2K)	set the memory management parameters
vmStat(2K)	get memory management statistics
vmUnLock(2K)	See vmLock(2K)

**NAME** | acap – get a c\_actor capability

**SYNOPSIS** | #include <am/afexec.h>  
 int **acap**(int *aid*, KnCap \**cactorcap*);

**FEATURES** | ACTOR\_EXTENDED\_MNGT

**DESCRIPTION** | *acap*( ) copies, at the location pointed to by *cactorcap*, the capability of the *c\_actor* designated by the identifier *aid*. See *intro*(2K).

**RETURN VALUES** | Upon successful completion, *acap*( ) returns 0. Otherwise it returns -1 and sets *errno* to indicate one of the following error conditions:  
 [EINVAL] | *aid* is zero.  
 [ESRCH] | *aid* does not exist.  
 [EPERM] | the calling *c\_actor* is not a trusted *c\_actor*.  
 [EFAULT] | *cactorcap* points outside the allocated address space of the *c\_actor*.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | *afexec*(2K), *agetId*(2K), *intro*(2K)

<b>NAME</b>	aconf – get configurable system variables						
<b>SYNOPSIS</b>	<pre>#include &lt;am/aconf.h&gt; int aconf(int name);</pre>						
<b>FEATURES</b>	ACTOR_EXTENDED_MGMT						
<b>DESCRIPTION</b>	<p>The <code>aconf()</code> function provides a method for the application to determine the current value of configurable system limits, modules or options (variables) defined by the ChorusOS 4.0 AM.</p> <p>The <code>name</code> argument represents the system variable/module to be queried. The symbolic constants used for <code>name</code> are defined in <code>&lt;am/aconf.h&gt;</code> and appear in the table that follows.</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Variable name</th> <th style="text-align: left;">Value</th> </tr> </thead> <tbody> <tr> <td>AMDS</td> <td><code>_AC_M_AMDS</code> Dump module.</td> </tr> <tr> <td>ARG_MAX</td> <td><code>_AC_C_ARG_MAX</code> Maximum length of arguments for the exec functions, in bytes, including environment.</td> </tr> </tbody> </table> <p>Other values may be obtained by using <code>sysctl(3POSIX)</code> for values defined by the IOM, or by using <code>sysconf(3POSIX)</code>.</p>	Variable name	Value	AMDS	<code>_AC_M_AMDS</code> Dump module.	ARG_MAX	<code>_AC_C_ARG_MAX</code> Maximum length of arguments for the exec functions, in bytes, including environment.
Variable name	Value						
AMDS	<code>_AC_M_AMDS</code> Dump module.						
ARG_MAX	<code>_AC_C_ARG_MAX</code> Maximum length of arguments for the exec functions, in bytes, including environment.						
<b>RETURN VALUE</b>	<p>If <code>name</code> is an invalid value, <code>aconf()</code> returns <code>-1</code>.</p> <p>Otherwise, <code>aconf()</code> returns the current variable value on the system. The value does not change during the lifetime of the calling actor. If <code>name</code> corresponds to an AM module, <code>aconf()</code> returns <code>1</code> if the module is present, or <code>0</code> if the module is not present.</p>						
<b>ERRORS</b>	[EINVAL]                      The <code>name</code> argument is invalid.						
<b>ATTRIBUTES</b>	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: left;">ATTRIBUTE TYPE</th> <th style="text-align: left;">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						
<b>SEE ALSO</b>	<code>sysconf(3POSIX)</code>						

<b>NAME</b>	acreate – creates a <code>c_actor</code>
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int acreate(KnCap *cactorcap, const AcParam *param);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p>The <i>acreate</i> call creates a <code>c_actor</code> according to the arguments specified by <i>param</i>.</p> <p>The <i>param</i> argument points to an <i>AcParam</i> structure which should contain the specification of the new <code>c_actor</code>. An <i>AcParam</i> structure has the following members:</p> <pre>int      acSite ;           /* new c_actor execution site ID */ int      acFlags ;         /* c_actor options */ char*    acStdin ;         /* standard input */ char*    acStdout ;        /* standard output */ char*    acStderr ;        /* standard error output */ char*    acCurdir ;       /* current directory */ char*    acRootdir ;       /* root directory */ cx_cred_t* acCred ;        /* c_actor credentials (This member is ignored.) */</pre> <p>The new <code>c_actor</code> runs on the local site. <i>acSite</i> is ignored.</p> <p>The <i>acFlags</i> values are constructed by using a combination of the following flags. Note that either <code>AFX_USER_SPACE</code> or <code>AFX_SUPERVISOR_SPACE</code>, but not both, can be used.</p> <p><code>AFX_USER_SPACE</code>      If <code>AFX_TRUSTED</code> is set, the new <code>c_actor</code> is implemented as a system actor (see <code>intro(2K)</code>) and is a Trusted <code>c_actor</code> (see <code>intro(2K)</code>), otherwise the new <code>c_actor</code> is implemented as a user actor and is not Trusted.</p> <p><code>AFX_SUPERVISOR_SPACE</code>      The new <code>c_actor</code> is implemented as a supervisor actor (see <code>intro(2K)</code>) and is a Trusted <code>c_actor</code> (see <code>intro(2K)</code>).</p> <p><code>AFX_ANY_SPACE</code>      <code>AFX_ANY_SPACE</code> is meaningful only with <i>afexec</i>; it is not applicable with <i>acreate</i>.</p> <p>                            The new <code>c_actor</code> privilege (<code>user</code> or <code>supervisor</code>) will be deduced from binary at load time. This flag takes precedence over the two previous ones. If the <code>AFX_NON_TRUSTED</code> flag is set, the privilege deduced from binary will be the <code>user</code> privilege.</p> <p><code>AFX_DEBUGGEE</code>      The new <code>c_actor</code> will be started in debug mode. If <code>AFX_WAIT_ATTACH</code> is also set, the new <code>c_actor</code> will be stopped to wait for a debugger to perform an <code>ATRACE_ATTACH</code> command through the (see</p>

	<code>atrace(2K)</code> system call; otherwise, the calling <code>c_actor</code> is considered as its own debugger.
<code>AFX_EXECUTE_IN_PLACE</code>	The new <code>c_actor</code> will be executed in place, meaning the code will not be copied, but instead executed where found. Data will be copied. This requires the actor to be fully linked.
<code>AFX_KDB_SYMBOLS</code>	A description of the new <code>c_actor</code> symbols will be loaded in memory. This symbol description will enable symbolic debugging of the new <code>c_actor</code> using the kernel debugger. This flag is only meaningful for a <code>supervisor c_actor</code> .
<code>AFX_DEBUGMODE</code>	The new <code>c_actor</code> will be put into the <code>DEBUGMODE</code> mode (see <code>actorCreate(2K)</code> ) That means that more information than usual about thread execution will be gathered by the system. Not to be confused with <code>AFX_DEBUGGEE</code> .
<code>AFX_STOPPED</code>	The new <code>c_actor</code> will be put into the <code>STOPPED</code> state. Not to be confused with <code>AFX_WAIT_ATTACH</code> , as no <code>ATRACE_ATTACH</code> command is expected.

The `acStdin`, `acStdout`, and `acStderr` fields specify the path name for the standard input, the standard output and the standard error output files for the new `c_actor`.

The `acCurdir` field allows you to specify the pathname for the current directory of the new `c_actor`.

The `acRootdir` field allows you to specify the pathname for the root directory of the new `c_actor`.

The `acCred` field allows you to set the credentials of the new `c_actor`. This field may only be set by a `trusted c_actor`.

Unlike `afexec`, `acreate` does not load the new `c_actor` or activate it.

After `acreate` has been created, the new `c_actor` exists, a corresponding actor has been created according to the `acFlags` specified, the standard I/O files and directories specified are open, and the new `c_actor` has the specified credentials. However, the new `c_actor` does not have any regions or threads. In addition, the new `c_actor` is not `mature`, which means that if the calling `c_actor` is deleted, the new `c_actor` will be deleted.

If one of the pointer-type members of `*param` is `NULL`, the corresponding parameter is considered `un-specified`. If a parameter is `un-specified`,

the new `c_actor` will inherit the corresponding running attributes of the calling `c_actor`.

File descriptors with the close-on-afexec flag set (see `afexec(2K)`) will not be inherited by the new `c_actor`.

**LIMITATIONS**

Correct functioning of a supervisor actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) is not guaranteed. Even if the supervisor actor functions normally, only one supervisor actor can be executed in place at any one time from a given binary file. Attempting to execute more than one supervisor actor in place from the same binary file will be rejected by the system.

Correct functioning of any actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) when in flat memory mode is not guaranteed. Even if the actor functions normally when in flat memory mode, only one actor can be executed in place at any one time from a given binary file. Attempting to execute more than one actor in place in this situation will be rejected by the system.

**RETURN VALUES**

Upon successful completion, `acreate` returns a non-negative integer that is the *aid* of the new `c_actor`. Otherwise `acreate` returns `-1` and sets `errno` to indicate one of the following error conditions:

- [EPERM]                      The calling `c_actor` is not trusted and is either trying to create a trusted `c_actor` or to set the new `c_actor`'s credentials.
- [ENOENT]                    One or more components of the new `c_actor`'s I/O files' path names do not exist.
- [ENOTDIR]                   A component of the new `c_actor`'s I/O files' path prefix is not a directory.
- [EACCES]                    Search permission is denied for a directory listed in one of the new `c_actor`'s I/O files' path prefix.
- [EINVAL]                    Inconsistent attribute flags were provided.
- [EINVAL]                    *parame* was `NULL`.
- [ENAMETOOLONG]            The length of a component of one of the I/O file's path exceeds `NAME_MAX` characters, or the length of the complete pathname exceeds `PATH_MAX` characters.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

afexec(2K), akill(2K), aload(2K), astart(2K), hrfexec(2RESTART)

**NAME** | `acred` – get/set `c_actor` credentials

**SYNOPSIS** | `#include <cx/cred.h>`  
`int acred(const KnCap *cactorcap, cx_cred_t *oldCred, const cx_cred_t *newCred);`

**FEATURES** | `ACTOR_EXTENDED_MNGT`

**DESCRIPTION** | If *oldCred* is not the `NULL` pointer, a copy of the credential structure of the `c_actor` designated by its capability *cactorcap* will be copied at the location pointed to by *oldCred*.  
 If *newCred* is not the `NULL` pointer, the credentials of the `c_actor` designated by its capability *cactorcap* will be set to the value of the credentials pointed to by *newCred*.  
 A structure `cx_cred_t` includes the following members:  

```
uid_t          cr_uid;          /* c_actor's user ID */
gid_t          cr_gid;          /* c_actor's group ID */
unsigned short cr_ngroups;     /* number of groups in cr_groups */
gid_t          cr_groups[];     /* supplementary group list */
```

**RETURN VALUE** | Upon successful completion, *acred* returns 0; otherwise it returns the following error conditions:  
`[ESRCH]` | *cactorcap* doesn't designate a valid `c_actor`.  
`[EPERM]` | *newCred* is not the `NULL` pointer and the caller is not a trusted `c_actor`.  
`[EFAULT]` | *cactorcap*, *oldCred* or *newCred* points outside the allocated address space of the calling `c_actor`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | `afexec(2K)`

<b>NAME</b>	actorCreate – create an actor
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int actorCreate(KnCap *actorinit, KnCap *actorcap, KnActorPrivilege privilege, KnActorStatus status);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>actorCreate</i> call creates a new actor. A capability for this actor is returned in <i>actorcap</i>, which is a pointer to a <i>KnCap</i> structure whose members are the following:</p> <pre>KnUniqueId    ui ;           /* entity name */ KnKey         key ;         /* modification key */</pre> <p>Where <i>ui</i> is the unique identifier of the actor, and <i>key</i> is the key needed to modify the actor (for instance, creating or deleting threads or other resources, or stopping the actor).</p> <p>The new actor inherits a number of attributes from its init actor, whose capability is given by <i>actorinit</i>. If <i>actorinit</i> is <i>K_MYACTOR</i>, the current actor is taken as the init actor. The CORE module attributes inherited from the init actor are: the exception handler (see <i>svExcHandler(2K)</i>) and the abort handler (see <i>svAbortHandler(2K)</i>). Furthermore, other kernel modules, such as communication modules, may define attributes which are inherited by the new actor.</p> <p>The <i>privilege</i> field gives the actor's privilege. Two main privileges may be given to an actor: <i>USER</i> or <i>SYSTEM</i>. A system actor's thread has access to some privileged kernel calls even if the thread is not running with the <i>SUPERVISOR</i> privilege level (see <i>threadCreate(2K)</i>). A <i>SYSTEM</i> actor may only be created by a thread running with the <i>SUPERVISOR</i> privilege or if its actor is a <i>SYSTEM</i> actor.</p> <p>In addition, <i>SYSTEM</i> actors may be created as <i>SUPERVISOR</i> actors: a <i>SUPERVISOR</i> actor shares the kernel address space, and no user address space is allocated to it. <i>SUPERVISOR</i> actors may only contain <i>SUPERVISOR</i> threads.</p> <p>A <i>USER</i> actor is created with <i>privilege</i> equal to <i>K_USERACTOR</i>; a <i>SYSTEM</i> actor is created with <i>privilege</i> equal to <i>K_SYSTEMACTOR</i>; a <i>SUPERVISOR</i> actor is created with <i>privilege</i> equal to <i>K_SUPACTOR</i>.</p> <p>When created, the new actor contains neither thread nor memory region.</p> <p>The <i>status</i> parameter is the initial actor status: if the <i>K_STOPPED</i> flag is set in the <i>status</i> parameter, the new actor is created in the <i>STOPPED</i> state. That means that even if threads are created in the actor (using <i>threadCreate(2K)</i>), none of these threads are able to run until the actor is put into the <i>ACTIVE</i> state using <i>actorStart(2K)</i>.</p>

If the `K_STOPPED` flag is not set, threads that are created in the *ACTIVE* state are able to run immediately.

if the `K_DEBUGMODE` flag is set in the *status* parameter, the new actor will be put into the *DEBUGMODE* mode. That means that more information than usual about thread execution will be gathered by the system. This helps to support some source level debuggers.

The maximum number of actors that may be created on a site is tunable parameter.

**RETURN VALUE**

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

- [K\_EINVAL] *actorinit* is an inconsistent actor capability.
- [K\_EUNKNOWN] *actorinit* does not specify a reachable actor.
- [K\_EFAULT] Some of the data provided are outside the current actor's address space.
- [K\_ENOMEM] The system is out of resources.
- [K\_EPRIV] The current thread is not allowed to create a *SYSTEM* actor.

**RESTRICTIONS**

The init actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`actorDelete(2K)`, `actorStart(2K)`, `actorStop(2K)`, `actorCapability(2K)`, `threadCreate(2K)`

<b>NAME</b>	actorDelete – delete an actor								
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int actorDelete(KnCap *actorcap);</pre>								
<b>FEATURES</b>	CORE								
<b>DESCRIPTION</b>	<p>The <i>actorDelete</i> call deletes the actor whose capability is given by <i>actorcap</i> .</p> <p>If <i>actorcap</i> is K_MYACTOR, the current actor is deleted.</p> <p>When the actor is deleted, all the actor's resources are deleted. For example, threads are deleted just as if <i>threadDelete(2K)</i> had been called.</p>								
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.								
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EINVAL]</td> <td><i>actorcap</i> is an inconsistent actor capability.</td> </tr> <tr> <td>[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td>[K_EFAULT]</td> <td>Some of the data provided are outside the current actor's address space.</td> </tr> <tr> <td>[K_EBUSY]</td> <td>The actor is already in the process of being deleted.</td> </tr> </table>	[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.	[K_EFAULT]	Some of the data provided are outside the current actor's address space.	[K_EBUSY]	The actor is already in the process of being deleted.
[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.								
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.								
[K_EFAULT]	Some of the data provided are outside the current actor's address space.								
[K_EBUSY]	The actor is already in the process of being deleted.								
<b>RESTRICTIONS</b>	<p>The deleted actor and the current actor must be located on the same site.</p> <p>Interrupt, trap, exception or time-out handlers connected by the actor are not disconnected on actor deletion in this version. When deleting <i>SUPERVISOR</i> actors, these handlers must be disconnected explicitly before invoking <i>actorDelete</i>. Otherwise, unpredictable results may occur.</p>								
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Interface Stability	Evolving								
<b>SEE ALSO</b>	<i>actorCreate(2K)</i> , <i>threadDelete(2K)</i>								

**NAME** actorName – get and/or set the symbolic name of an actor

**SYNOPSIS** #include <exec/chExec.h>  
int actorName(KnCap \*actorcap, char \*oldname, char \*newname);

**FEATURES** CORE

**DESCRIPTION** The *actorName* call gets and/or sets the symbolic name of the actor whose capability is given by *actorcap* (see *actorCreate(2K)*). If *actorcap* is K\_MYACTOR, the current actor is affected.

If *oldname* is not NULL, the actor’s symbolic name is copied to the caller address space at the location specified by *oldname*. An actor’s symbolic name has a maximum size of K\_ACTORNAMEMAX (including the NULL character). If *newname* is not NULL, the actor’s symbolic name is set to the new name pointed to by *newname* in the caller address space. The new name is truncated to a maximum size of K\_ACTORNAMEMAX (including the NULL character).

**RETURN VALUE** Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EFAULT]	Some of the data provided are outside the current actor’s address space.

**RESTRICTIONS** The target actor and the current actor must be located on the same site.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *actorCreate(2K)*

<b>NAME</b>	actorPi – get and/or set the protection identifier of an actor								
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; #include &lt;ipc/chId.h&gt; #include &lt;ipc/chIpc.h&gt; int actorPi(KnCap *actorcap, KnProtId *oldpi, KnProtId *newpi);</pre>								
<b>FEATURES</b>	IPC								
<b>DESCRIPTION</b>	<p>The <i>actorPi</i> call gets and/or sets the protection identifier of the actor whose capability is given by <i>actorcap</i>. If <i>actorcap</i> is <code>K_MYACTOR</code>, the operation is applied to the current actor.</p> <p>The <i>oldpi</i> and <i>newpi</i> pointers indicate <i>KnProtId</i> structures. The structure of <i>KnProtId</i> is described in <i>portPi(2K)</i>. The corresponding actor's protection identifier is copied to the structure pointed to by <i>oldpi</i> in the client address space. The new corresponding context is copied from the structure pointed to by <i>newpi</i>.</p> <p>If <i>oldpi</i> is a NULL pointer, the actor's protection identifier is not copied in the caller's address space. If <i>newpi</i> is a NULL pointer, the actor's protection identifier is not modified. The calling thread must be a <i>SUPERVISOR</i> thread (see <i>threadCreate(2K)</i>) or must belong to a <i>SYSTEM</i> actor (see <i>actorCreate(2K)</i>).</p>								
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.								
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td><i>actorcap</i> is an inconsistent actor capability.</td> </tr> <tr> <td style="vertical-align: top;">[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the data provided are outside the current actor's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_EPRIV]</td> <td>The current thread is neither a supervisor thread nor a thread of a system actor.</td> </tr> </table>	[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.	[K_EFAULT]	Some of the data provided are outside the current actor's address space.	[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor.
[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.								
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.								
[K_EFAULT]	Some of the data provided are outside the current actor's address space.								
[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor.								
<b>RESTRICTIONS</b>	The target actor and the current actor must be located on the same site.								
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Interface Stability	Evolving								
<b>SEE ALSO</b>	<i>actorCreate(2K)</i> , <i>threadCreate(2K)</i>								

**NAME** actorPrivilege – get and/or set the privilege of an actor

**SYNOPSIS**

```
#include <exec/chExec.h>
int actorPrivilege(KnCap *actorcap, KnActorPrivilege *oldpriv, KnActorPrivilege *newpriv);
```

**FEATURES** CORE

**DESCRIPTION** *actorPrivilege* gets and/or sets the privilege of the actor the capability of which is given by *actorcap*. If *actorcap* is `K_MYACTOR`, the operation is applied to the current actor.

*oldpriv* and *newpriv* are pointers to *KnActorPrivilege* data. *KnActorPrivilege* represents the actor privilege, as described in *actorCreate* (2K).

If *oldpriv* is a NULL pointer, the actor’s privilege is not copied in the caller’s address space. If *newpriv* is a NULL pointer, the actor’s privilege is not modified. The calling thread must be a *SUPERVISOR* thread (see *threadCreate*(2K)) or must belong to a *SYSTEM* actor (see *actorCreate*(2K)).

The only allowed privilege modifications are from USER to SYSTEM or from SYSTEM to USER.

**RETURN VALUE** Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EFAULT]	Some of the provided data are outside the current actor’s address space.
[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor.

**RESTRICTIONS** The target actor and the current actor must be located on the same site.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *actorCreate*(2K), *threadCreate*(2K)

<b>NAME</b>	actorSelf – get the current actor capability					
<b>SYNOPSIS</b>	#include <exec/chExec.h> int actorSelf(KnCap *actorcap);					
<b>FEATURES</b>	CORE					
<b>DESCRIPTION</b>	<i>actorSelf</i> returns the capability of the current actor into the structure pointed to by <i>actorcap</i> .					
<b>RETURN VALUE</b>	Upon successful completion, a value of 0 is returned. Otherwise, a negative error code is returned.					
<b>ERRORS</b>	[K_EFAULT]	Some of the provided data are outside the current actor's address space.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Interface Stability	Evolving					
<b>SEE ALSO</b>	actorCreate(2K)					

<b>NAME</b>	actorStop, actorStart – Stop an actor; Start an actor
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int actorStart(KnCap * actorcap);  int actorStop(KnCap * actorcap);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>actorStop</i> call stops the actor whose capability is given by <i>actorcap</i> (see <i>actorCreate</i> (2K)).</p> <p>If <i>actorcap</i> is K_MYACTOR, the current actor is stopped.</p> <p>The effect of <i>actorStop</i> is to prevent all the threads from running until the actor is restarted using <i>actorStart</i> (2K) .</p> <p>Performing <i>actorStop</i> is equivalent to performing <i>threadStop</i> on all the actor's threads.</p> <p>All other threads executing in the actor and having entered it through safe LAPs, and any thread attempting to enter a stopped actor through a safe LAP will also be stopped.</p> <p>The effect of <i>actorStop</i> is not instantaneous on threads executing a system call that is implemented via a trap, but it guarantees that a thread performing this type of system call will not return from that call.</p> <p>When an actor has been stopped, the internal status of its threads may, however, change (if messages are delivered or resumed, for example).</p> <p>It is possible to resume individual threads using <i>threadStart</i> (2K).</p> <p>The <i>actorStart</i> call starts the actor whose capability is given by <i>actorcap</i>. It may be applied to a newly created actor, or to an actor previously stopped using <i>actorStop</i> (2K).</p> <p>If <i>actorcap</i> is K_MYACTOR, the current actor is started.</p> <p>After <i>actorStart</i> , any of the runnable actor's threads may run user code.</p> <p>Performing <i>actorStart</i> is equivalent to performing <i>threadStart</i> on all the actor's threads and to resuming all other threads executing in the actor or attempting to enter it through safe LAPs.</p>
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	[K_EINVAL] <i>actorcap</i> is an inconsistent actor capability.

[K\_EUNKNOWN]

*actorcap* does not specify a reachable actor.

[K\_EFAULT]

Some of the data provided is out of the current actor's address space.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**`actorCreate(2K)` , `threadContext(2K)` , `threadStart(2K)` ,  
`threadStop(2K)`

**NAME** actorStat – get the status of all actors on a site

**SYNOPSIS**

```
#include <exec/chExec.h>
int actorStat(unsigned int options, KnActorStat *stat, unsigned int buffsize);
```

**FEATURES** CORE

**DESCRIPTION** The *actorStat* call obtains the status of all actors created on the current site. The *stat* argument points to a buffer in which the descriptions of existing actors will be returned. The *buffsize* argument gives the size of the buffer. The buffer is considered part of the caller (current user or supervisor) address space. The *options* argument is reserved for future use and must be set to 0. On successful return the buffer will contain an array of descriptors, each one describing a single actor. Each descriptor is a *KnActorStat* structure with the following fields:

```
KnCap asCap ;
KnActorStatus asStatus ;
```

The *asCap* field is the actor capability. The *asStatus* field is the actor status (K\_STOPPED or K\_ACTIVE, see *actorCreate(2K)* for details). The *actorStat* call is restricted to threads running with *SUPERVISOR* privilege or whose actor is a *SYSTEM* actor.

**RETURN VALUE** If successful, *actorStat* returns the number of actors created on the current site. Otherwise a negative error code is returned.

**ERRORS**

[K_EFAULT]	Some of the arguments provided are outside the caller's address space.
[K_EPRIV]	The current thread is not allowed to perform this call.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *actorCreate(2K)*

<b>NAME</b>	actorStop, actorStart – Stop an actor; Start an actor
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int actorStart(KnCap * actorcap);  int actorStop(KnCap * actorcap);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>actorStop</i> call stops the actor whose capability is given by <i>actorcap</i> (see <i>actorCreate</i> (2K)).</p> <p>If <i>actorcap</i> is K_MYACTOR, the current actor is stopped.</p> <p>The effect of <i>actorStop</i> is to prevent all the threads from running until the actor is restarted using <i>actorStart</i> (2K) .</p> <p>Performing <i>actorStop</i> is equivalent to performing <i>threadStop</i> on all the actor's threads.</p> <p>All other threads executing in the actor and having entered it through safe LAPs, and any thread attempting to enter a stopped actor through a safe LAP will also be stopped.</p> <p>The effect of <i>actorStop</i> is not instantaneous on threads executing a system call that is implemented via a trap, but it guarantees that a thread performing this type of system call will not return from that call.</p> <p>When an actor has been stopped, the internal status of its threads may, however, change (if messages are delivered or resumed, for example).</p> <p>It is possible to resume individual threads using <i>threadStart</i> (2K).</p> <p>The <i>actorStart</i> call starts the actor whose capability is given by <i>actorcap</i>. It may be applied to a newly created actor, or to an actor previously stopped using <i>actorStop</i> (2K).</p> <p>If <i>actorcap</i> is K_MYACTOR, the current actor is started.</p> <p>After <i>actorStart</i> , any of the runnable actor's threads may run user code.</p> <p>Performing <i>actorStart</i> is equivalent to performing <i>threadStart</i> on all the actor's threads and to resuming all other threads executing in the actor or attempting to enter it through safe LAPs.</p>
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	[K_EINVAL] <i>actorcap</i> is an inconsistent actor capability.

[K\_EUNKNOWN]

*actorcap* does not specify a reachable actor.

[K\_EFAULT]

Some of the data provided is out of the current actor's address space.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO***actorCreate(2K)* , *threadContext(2K)* , *threadStart(2K)* , *threadStop(2K)*

<b>NAME</b>	afexec, afexecl, afexecv, afexecle, afexecve, afexeclp, afexecvp – create a new <code>c_actor</code>
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int afexecve(const char * path, KnCap * cactorcap, const AcParam * param, char const * argv, char const * envp);  int afexecl(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecv(const char * path, KnCap * cactorcap, const AcParam * param, char *const * argv);  int afexecle(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */, char const * envp);  int afexeclp(const char * file, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecvp(const char * file, KnCap * cactorcap, const AcParam * param, char *const * argv);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p><i>afexec</i> family routines create a new <code>c_actor</code> with a single thread (the main thread). This main thread will execute with the priority defined by the <i>am.afexecschedprio</i> configuration parameter. It will have an automatically allocated stack, which size is defined by the <i>am.afexec.userstacksize</i> configuration parameter for user actors, and by <i>kern.exec.dflSysStackSize</i> for supervisor actors. (See <i>intro(2K)</i>).</p> <p>The main thread starts executing the runtime startup routine (usually called <i>_start</i>) specified as the program entry point in the binary header of the <i>new c_actor file</i>. After initialization of the threadsafe C library, the runtime startup routine calls the C program as follows:</p> <pre>int main (int argc, char* argv[], char* envp[])</pre> <p>where <i>argc</i> is the argument count, <i>argv</i> is an array of character pointers to the arguments themselves, and <i>envp</i> is an array of character pointers to the environment strings. As will be described later, <i>argc</i> is conventionally at least 1, and the first member of the array points to a string containing the name (or pathname) of the executable file.</p> <p><i>cactorcap</i> is the returned capability of the new <code>c_actor</code>, and is a pointer to a <i>KnCap</i> structure (see <i>actorCreate(2K)</i>).</p> <p>If <i>param</i> is <code>NULL</code>, default parameters are used, <i>acFlags</i> is set to <code>AFX_ANY_SPACE</code> and all others to un-specified. See <i>acreate(2K)</i> for a description of the <code>AcParam</code> structure.</p>

By convention, *afexec* routines search for a binary, either fully linked or relocatable, in *path*. If no such binary exists, *afexec* routines search for a file with a *.r* suffix and then with a *.O* suffix. If the *GZ\_FEATURE* is set to *true*, for each step, the *afexec* routines also search for a file with an additional *.gz* (for example: *path*, *path.gz*, *path.r*, *path.r.gz*).

*path* may also refer to a relocatable binary stored in a memory buffer in supervisor space, using the syntax "ram:0x<start>:0x<length>". <start> is the start address at which the binary is stored. <length> is the length of the binary. The buffer is not used any more after loading, except if the program uses the dynamic library API. Freeing the buffer after *afexec* may have unpredictable results if the program uses the dynamic library API.

Files to be loaded may be fully linked or relocatable files. They may also include dependencies on dynamic, but not shared, libraries. These libraries are looked for and loaded within the actor address space as part of the *afexec* call.

*arg0*, ... , *argn* are pointers to null-terminated character strings. These strings constitute the argument list available to the new *c\_actor*. By convention, at least *arg0* must be present and point to a string that is the same as *path* (or *file* or the last component of *path*).

*argv* is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new *c\_actor*. By convention, *argv* must have at least one member, and it must point to a string that is the same as *path* (or *file* or the last component of *path*). *argv* is terminated by a null pointer.

*envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new *c\_actor*. *envp* is terminated by a null pointer. For *afexecl*, *afexecv*, *afexeclp* and *afexecvp*, the C runtime start-off routine places a pointer to the environment of the calling *c\_actor* in the global object:

```
extern char** environ;
```

and it is used to pass the environment of the calling *c\_actor* to the new *c\_actor*.

*afexeclp* and *afexecvp* are called with the same arguments as *afexecl* and *afexecv*, but duplicate the shell's actions in searching for an executable file in a list of directories. The directory list is obtained from the *PATH* environment variable, and *file* must point to an executable file name for the new *c\_actor*.

## LIMITATIONS

Incompatible requests (such as *afexec* of a supervisor binary with the *AFX\_USER\_SPACE* flag set) are not normally detected. Applications can use the *AFX\_ANY\_SPACE* (see *acreate(2K)*) flag to force the system to deduce the type (*user* or *supervisor*) of the binary. In this case incompatible requests will be detected and refused.

Correct functioning of a supervisor actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) is not guaranteed. Even if the supervisor actor functions normally, only one supervisor actor can be executed in place at any one time from a given binary file. Attempting to execute more than one supervisor actor in place from the same binary file will be rejected by the system.

Correct functioning of any actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) when in flat memory mode is not guaranteed. Even if the actor functions normally, when in flat memory mode only one actor can be executed in place at any one time from a given binary file. Attempting to execute more than one actor in place in this situation will be rejected by the system.

**RETURN VALUES**

Upon successful completion, *afexec* returns a non-negative integer that is the *aid* of the new `c_actor`. Otherwise *afexec* returns `-1` and sets `errno` to indicate one of the following error conditions.

**ERRORS**

[E2BIG]	The number of bytes used by the new <code>c_actor</code> 's image argument list and environment list is greater than <code>ARG_MAX</code> bytes.
[EACCES]	Search permission is denied for a directory listed in the new <code>c_actor</code> 's file path prefix.
[EBADSYM]	The runtime link-editor cannot find a symbol needed by the new <code>c_actor</code> . Possible reasons are that ChorusOS 4.0 <code>imake</code> rules have not been used to compile the program, or the link editor has found unexpected dynamic libraries.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	Inconsistent attribute flags were specified.
[EINVAL]	The address of the text or data region of the new <code>c_actor</code> is not consistent with the <code>AFX_USER_SPACE</code> or <code>AFX_SUPERVISOR_SPACE</code> attribute flag.
[EIO]	An I/O error occurred while reading from the file system.
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds <code>NAME_MAX</code> characters or the length of <i>path</i> exceeds <code>PATH_MAX</code> characters.
[ENOENT]	The new <code>c_actor</code> is a dynamic actor and one of its dependencies cannot be found either in the directories listed in the <code>LD_LIBRARY_PATH</code>

	environment variable, or in its runpath, or in /usr/lib.
[ENOEXEC]	The new c_actor file has the appropriate access permission but is not in the correct format.
[ENOTDIR]	A component of the new c_actor path of the file prefix is not a directory.
[EPERM]	The calling c_actor is not Trusted and is trying either to create a Trusted c_actor or set the new c_actor's credentials.
[ELOOP]	Too many symbolic links have been encountered during analysis of <i>path</i> .

---

**Note** - Error codes generated by *read* and *open* may also be returned.

---

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`aconf(2K)`, `acreate(2K)`, `actorCreate(2K)`, `akill(2K)`, `aload(2K)`, `astart(2K)`, `astat(2K)`, `await(2K)`, `dlopen(2K)`, `fcntl(2POSIX)`, `intro(2K)`

<b>NAME</b>	afexec, afexecl, afexecv, afexecle, afexecve, afexeclp, afexecvp – create a new <code>c_actor</code>
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int afexecve(const char * path, KnCap * cactorcap, const AcParam * param, char const * argv, char const * envp);  int afexecl(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecv(const char * path, KnCap * cactorcap, const AcParam * param, char *const * argv);  int afexecle(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */, char const * envp);  int afexeclp(const char * file, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecvp(const char * file, KnCap * cactorcap, const AcParam * param, char *const * argv);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p><i>afexec</i> family routines create a new <code>c_actor</code> with a single thread (the main thread). This main thread will execute with the priority defined by the <i>am.afexecschedprio</i> configuration parameter. It will have an automatically allocated stack, which size is defined by the <i>am.afexec.userstacksize</i> configuration parameter for user actors, and by <i>kern.exec.dflSysStackSize</i> for supervisor actors. (See <i>intro(2K)</i>).</p> <p>The main thread starts executing the runtime startup routine (usually called <i>_start</i>) specified as the program entry point in the binary header of the <i>new c_actor file</i>. After initialization of the threadsafe C library, the runtime startup routine calls the C program as follows:</p> <pre>int main (int argc, char* argv[], char* envp[])</pre> <p>where <i>argc</i> is the argument count, <i>argv</i> is an array of character pointers to the arguments themselves, and <i>envp</i> is an array of character pointers to the environment strings. As will be described later, <i>argc</i> is conventionally at least 1, and the first member of the array points to a string containing the name (or pathname) of the executable file.</p> <p><i>cactorcap</i> is the returned capability of the new <code>c_actor</code>, and is a pointer to a <i>KnCap</i> structure (see <i>actorCreate(2K)</i>).</p> <p>If <i>param</i> is <code>NULL</code>, default parameters are used, <i>acFlags</i> is set to <code>AFX_ANY_SPACE</code> and all others to un-specified. See <i>acreate(2K)</i> for a description of the <code>AcParam</code> structure.</p>

By convention, *afexec* routines search for a binary, either fully linked or relocatable, in *path*. If no such binary exists, *afexec* routines search for a file with a *.r* suffix and then with a *.O* suffix. If the *GZ\_FEATURE* is set to *true*, for each step, the *afexec* routines also search for a file with an additional *.gz* (for example: *path*, *path.gz*, *path.r*, *path.r.gz*).

*path* may also refer to a relocatable binary stored in a memory buffer in supervisor space, using the syntax "ram:0x<start>:0x<length>". <start> is the start address at which the binary is stored. <length> is the length of the binary. The buffer is not used any more after loading, except if the program uses the dynamic library API. Freeing the buffer after *afexec* may have unpredictable results if the program uses the dynamic library API.

Files to be loaded may be fully linked or relocatable files. They may also include dependencies on dynamic, but not shared, libraries. These libraries are looked for and loaded within the actor address space as part of the *afexec* call.

*arg0*, ... , *argn* are pointers to null-terminated character strings. These strings constitute the argument list available to the new *c\_actor*. By convention, at least *arg0* must be present and point to a string that is the same as *path* (or *file* or the last component of *path*).

*argv* is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new *c\_actor*. By convention, *argv* must have at least one member, and it must point to a string that is the same as *path* (or *file* or the last component of *path*). *argv* is terminated by a null pointer.

*envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new *c\_actor*. *envp* is terminated by a null pointer. For *afexecl*, *afexecv*, *afexeclp* and *afexecvp*, the C runtime start-off routine places a pointer to the environment of the calling *c\_actor* in the global object:

```
extern char** environ;
```

and it is used to pass the environment of the calling *c\_actor* to the new *c\_actor*.

*afexeclp* and *afexecvp* are called with the same arguments as *afexecl* and *afexecv*, but duplicate the shell's actions in searching for an executable file in a list of directories. The directory list is obtained from the *PATH* environment variable, and *file* must point to an executable file name for the new *c\_actor*.

## LIMITATIONS

Incompatible requests (such as *afexec* of a supervisor binary with the *AFX\_USER\_SPACE* flag set) are not normally detected. Applications can use the *AFX\_ANY\_SPACE* (see *acreate(2K)*) flag to force the system to deduce the type (*user* or *supervisor*) of the binary. In this case incompatible requests will be detected and refused.

Correct functioning of a supervisor actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) is not guaranteed. Even if the supervisor actor functions normally, only one supervisor actor can be executed in place at any one time from a given binary file. Attempting to execute more than one supervisor actor in place from the same binary file will be rejected by the system.

Correct functioning of any actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) when in flat memory mode is not guaranteed. Even if the actor functions normally, when in flat memory mode only one actor can be executed in place at any one time from a given binary file. Attempting to execute more than one actor in place in this situation will be rejected by the system.

**RETURN VALUES**

Upon successful completion, *afexec* returns a non-negative integer that is the *aid* of the new `c_actor`. Otherwise *afexec* returns `-1` and sets `errno` to indicate one of the following error conditions.

**ERRORS**

[E2BIG]	The number of bytes used by the new <code>c_actor</code> 's image argument list and environment list is greater than <code>ARG_MAX</code> bytes.
[EACCES]	Search permission is denied for a directory listed in the new <code>c_actor</code> 's file path prefix.
[EBADSYM]	The runtime link-editor cannot find a symbol needed by the new <code>c_actor</code> . Possible reasons are that ChorusOS 4.0 <code>imake</code> rules have not been used to compile the program, or the link editor has found unexpected dynamic libraries.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	Inconsistent attribute flags were specified.
[EINVAL]	The address of the text or data region of the new <code>c_actor</code> is not consistent with the <code>AFX_USER_SPACE</code> or <code>AFX_SUPERVISOR_SPACE</code> attribute flag.
[EIO]	An I/O error occurred while reading from the file system.
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds <code>NAME_MAX</code> characters or the length of <i>path</i> exceeds <code>PATH_MAX</code> characters.
[ENOENT]	The new <code>c_actor</code> is a dynamic actor and one of its dependencies cannot be found either in the directories listed in the <code>LD_LIBRARY_PATH</code>

- environment variable, or in its runpath, or in /usr/lib .
- [ENOEXEC] The new c\_actor file has the appropriate access permission but is not in the correct format.
- [ENOTDIR] A component of the new c\_actor path of the file prefix is not a directory.
- [EPERM] The calling c\_actor is not Trusted and is trying either to create a Trusted c\_actor or set the new c\_actor's credentials.
- [ELOOP] Too many symbolic links have been encountered during analysis of *path* .

---

**Note** - Error codes generated by *read* and *open* may also be returned.

---

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

aconf(2K) , acreate(2K) , actorCreate(2K) , akill(2K) , aload(2K) , astart(2K) , astat(2K) , await(2K) , dlopen(2K) , fcntl(2POSIX) , intro(2K)

<b>NAME</b>	afexec, afexecl, afexecv, afexecle, afexecve, afexeclp, afexecvp – create a new <code>c_actor</code>
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int afexecve(const char * path, KnCap * cactorcap, const AcParam * param, char const * argv, char const * envp);  int afexecl(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecv(const char * path, KnCap * cactorcap, const AcParam * param, char *const * argv);  int afexecle(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */, char const * envp);  int afexeclp(const char * file, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecvp(const char * file, KnCap * cactorcap, const AcParam * param, char *const * argv);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p><i>afexec</i> family routines create a new <code>c_actor</code> with a single thread (the main thread). This main thread will execute with the priority defined by the <i>am.afexecschedprio</i> configuration parameter. It will have an automatically allocated stack, which size is defined by the <i>am.afexec.userstacksize</i> configuration parameter for user actors, and by <i>kern.exec.dflSysStackSize</i> for supervisor actors. (See <i>intro(2K)</i>).</p> <p>The main thread starts executing the runtime startup routine (usually called <i>_start</i>) specified as the program entry point in the binary header of the <i>new c_actor file</i>. After initialization of the threadsafe C library, the runtime startup routine calls the C program as follows:</p> <pre>int main (int argc, char* argv[], char* envp[])</pre> <p>where <i>argc</i> is the argument count, <i>argv</i> is an array of character pointers to the arguments themselves, and <i>envp</i> is an array of character pointers to the environment strings. As will be described later, <i>argc</i> is conventionally at least 1, and the first member of the array points to a string containing the name (or pathname) of the executable file.</p> <p><i>cactorcap</i> is the returned capability of the new <code>c_actor</code>, and is a pointer to a <i>KnCap</i> structure (see <i>actorCreate(2K)</i>).</p> <p>If <i>param</i> is <code>NULL</code>, default parameters are used, <i>acFlags</i> is set to <code>AFX_ANY_SPACE</code> and all others to un-specified. See <i>acreate(2K)</i> for a description of the <code>AcParam</code> structure.</p>

By convention, *afexec* routines search for a binary, either fully linked or relocatable, in *path*. If no such binary exists, *afexec* routines search for a file with a *.r* suffix and then with a *.O* suffix. If the *GZ\_FEATURE* is set to *true*, for each step, the *afexec* routines also search for a file with an additional *.gz* (for example: *path*, *path.gz*, *path.r*, *path.r.gz*).

*path* may also refer to a relocatable binary stored in a memory buffer in supervisor space, using the syntax "ram:0x<start>:0x<length>". <start> is the start address at which the binary is stored. <length> is the length of the binary. The buffer is not used any more after loading, except if the program uses the dynamic library API. Freeing the buffer after *afexec* may have unpredictable results if the program uses the dynamic library API.

Files to be loaded may be fully linked or relocatable files. They may also include dependencies on dynamic, but not shared, libraries. These libraries are looked for and loaded within the actor address space as part of the *afexec* call.

*arg0*, ..., *argn* are pointers to null-terminated character strings. These strings constitute the argument list available to the new *c\_actor*. By convention, at least *arg0* must be present and point to a string that is the same as *path* (or *file* or the last component of *path*).

*argv* is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new *c\_actor*. By convention, *argv* must have at least one member, and it must point to a string that is the same as *path* (or *file* or the last component of *path*). *argv* is terminated by a null pointer.

*envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new *c\_actor*. *envp* is terminated by a null pointer. For *afexecl*, *afexecv*, *afexeclp* and *afexecvp*, the C runtime start-off routine places a pointer to the environment of the calling *c\_actor* in the global object:

```
extern char** environ;
```

and it is used to pass the environment of the calling *c\_actor* to the new *c\_actor*.

*afexeclp* and *afexecvp* are called with the same arguments as *afexecl* and *afexecv*, but duplicate the shell's actions in searching for an executable file in a list of directories. The directory list is obtained from the *PATH* environment variable, and *file* must point to an executable file name for the new *c\_actor*.

## LIMITATIONS

Incompatible requests (such as *afexec* of a supervisor binary with the *AFX\_USER\_SPACE* flag set) are not normally detected. Applications can use the *AFX\_ANY\_SPACE* (see *acreate(2K)*) flag to force the system to deduce the type (*user* or *supervisor*) of the binary. In this case incompatible requests will be detected and refused.

Correct functioning of a supervisor actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) is not guaranteed. Even if the supervisor actor functions normally, only one supervisor actor can be executed in place at any one time from a given binary file. Attempting to execute more than one supervisor actor in place from the same binary file will be rejected by the system.

Correct functioning of any actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) when in flat memory mode is not guaranteed. Even if the actor functions normally, when in flat memory mode only one actor can be executed in place at any one time from a given binary file. Attempting to execute more than one actor in place in this situation will be rejected by the system.

**RETURN VALUES**

Upon successful completion, *afexec* returns a non-negative integer that is the *aid* of the new `c_actor`. Otherwise *afexec* returns `-1` and sets `errno` to indicate one of the following error conditions.

**ERRORS**

[E2BIG]	The number of bytes used by the new <code>c_actor</code> 's image argument list and environment list is greater than <code>ARG_MAX</code> bytes.
[EACCES]	Search permission is denied for a directory listed in the new <code>c_actor</code> 's file path prefix.
[EBADSYM]	The runtime link-editor cannot find a symbol needed by the new <code>c_actor</code> . Possible reasons are that ChorusOS 4.0 <code>imake</code> rules have not been used to compile the program, or the link editor has found unexpected dynamic libraries.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	Inconsistent attribute flags were specified.
[EINVAL]	The address of the text or data region of the new <code>c_actor</code> is not consistent with the <code>AFX_USER_SPACE</code> or <code>AFX_SUPERVISOR_SPACE</code> attribute flag.
[EIO]	An I/O error occurred while reading from the file system.
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds <code>NAME_MAX</code> characters or the length of <i>path</i> exceeds <code>PATH_MAX</code> characters.
[ENOENT]	The new <code>c_actor</code> is a dynamic actor and one of its dependencies cannot be found either in the directories listed in the <code>LD_LIBRARY_PATH</code>

	environment variable, or in its runpath, or in /usr/lib.
[ENOEXEC]	The new c_actor file has the appropriate access permission but is not in the correct format.
[ENOTDIR]	A component of the new c_actor path of the file prefix is not a directory.
[EPERM]	The calling c_actor is not Trusted and is trying either to create a Trusted c_actor or set the new c_actor's credentials.
[ELOOP]	Too many symbolic links have been encountered during analysis of <i>path</i> .

---

**Note** - Error codes generated by *read* and *open* may also be returned.

---

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*aconf(2K)*, *acreate(2K)*, *actorCreate(2K)*, *akill(2K)*, *aload(2K)*, *astart(2K)*, *astat(2K)*, *await(2K)*, *dlopen(2K)*, *fcntl(2POSIX)*, *intro(2K)*

<b>NAME</b>	afexec, afexecl, afexecv, afexecle, afexecve, afexeclp, afexecvp – create a new <code>c_actor</code>
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int afexecve(const char * path, KnCap * cactorcap, const AcParam * param, char const * argv, char const * envp);  int afexecl(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecv(const char * path, KnCap * cactorcap, const AcParam * param, char *const * argv);  int afexecle(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */, char const * envp);  int afexeclp(const char * file, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecvp(const char * file, KnCap * cactorcap, const AcParam * param, char *const * argv);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p><i>afexec</i> family routines create a new <code>c_actor</code> with a single thread (the main thread). This main thread will execute with the priority defined by the <i>am.afexecschedprio</i> configuration parameter. It will have an automatically allocated stack, which size is defined by the <i>am.afexec.userstacksize</i> configuration parameter for user actors, and by <i>kern.exec.dflSysStackSize</i> for supervisor actors. (See <i>intro(2K)</i>).</p> <p>The main thread starts executing the runtime startup routine (usually called <i>_start</i>) specified as the program entry point in the binary header of the <i>new c_actor file</i>. After initialization of the threadsafe C library, the runtime startup routine calls the C program as follows:</p> <pre>int main (int argc, char* argv[], char* envp[])</pre> <p>where <i>argc</i> is the argument count, <i>argv</i> is an array of character pointers to the arguments themselves, and <i>envp</i> is an array of character pointers to the environment strings. As will be described later, <i>argc</i> is conventionally at least 1, and the first member of the array points to a string containing the name (or pathname) of the executable file.</p> <p><i>cactorcap</i> is the returned capability of the new <code>c_actor</code>, and is a pointer to a <i>KnCap</i> structure (see <i>actorCreate(2K)</i>).</p> <p>If <i>param</i> is <code>NULL</code>, default parameters are used, <i>acFlags</i> is set to <code>AFX_ANY_SPACE</code> and all others to un-specified. See <i>acreate(2K)</i> for a description of the <code>AcParam</code> structure.</p>

By convention, *afexec* routines search for a binary, either fully linked or relocatable, in *path*. If no such binary exists, *afexec* routines search for a file with a *.r* suffix and then with a *.O* suffix. If the *GZ\_FEATURE* is set to *true*, for each step, the *afexec* routines also search for a file with an additional *.gz* (for example: *path*, *path.gz*, *path.r*, *path.r.gz*).

*path* may also refer to a relocatable binary stored in a memory buffer in supervisor space, using the syntax "ram:0x<start>:0x<length>". <start> is the start address at which the binary is stored. <length> is the length of the binary. The buffer is not used any more after loading, except if the program uses the dynamic library API. Freeing the buffer after *afexec* may have unpredictable results if the program uses the dynamic library API.

Files to be loaded may be fully linked or relocatable files. They may also include dependencies on dynamic, but not shared, libraries. These libraries are looked for and loaded within the actor address space as part of the *afexec* call.

*arg0*, ... , *argn* are pointers to null-terminated character strings. These strings constitute the argument list available to the new *c\_actor*. By convention, at least *arg0* must be present and point to a string that is the same as *path* (or *file* or the last component of *path*).

*argv* is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new *c\_actor*. By convention, *argv* must have at least one member, and it must point to a string that is the same as *path* (or *file* or the last component of *path*). *argv* is terminated by a null pointer.

*envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new *c\_actor*. *envp* is terminated by a null pointer. For *afexecl*, *afexecv*, *afexeclp* and *afexecvp*, the C runtime start-off routine places a pointer to the environment of the calling *c\_actor* in the global object:

```
extern char** environ;
```

and it is used to pass the environment of the calling *c\_actor* to the new *c\_actor*.

*afexeclp* and *afexecvp* are called with the same arguments as *afexecl* and *afexecv*, but duplicate the shell's actions in searching for an executable file in a list of directories. The directory list is obtained from the *PATH* environment variable, and *file* must point to an executable file name for the new *c\_actor*.

## LIMITATIONS

Incompatible requests (such as *afexec* of a supervisor binary with the *AFX\_USER\_SPACE* flag set) are not normally detected. Applications can use the *AFX\_ANY\_SPACE* (see *acreate(2K)*) flag to force the system to deduce the type (*user* or *supervisor*) of the binary. In this case incompatible requests will be detected and refused.

Correct functioning of a supervisor actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) is not guaranteed. Even if the supervisor actor functions normally, only one supervisor actor can be executed in place at any one time from a given binary file. Attempting to execute more than one supervisor actor in place from the same binary file will be rejected by the system.

Correct functioning of any actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) when in flat memory mode is not guaranteed. Even if the actor functions normally, when in flat memory mode only one actor can be executed in place at any one time from a given binary file. Attempting to execute more than one actor in place in this situation will be rejected by the system.

**RETURN VALUES**

Upon successful completion, *afexec* returns a non-negative integer that is the *aid* of the new `c_actor`. Otherwise *afexec* returns `-1` and sets `errno` to indicate one of the following error conditions.

**ERRORS**

[E2BIG]	The number of bytes used by the new <code>c_actor</code> 's image argument list and environment list is greater than <code>ARG_MAX</code> bytes.
[EACCES]	Search permission is denied for a directory listed in the new <code>c_actor</code> 's file path prefix.
[EBADSYM]	The runtime link-editor cannot find a symbol needed by the new <code>c_actor</code> . Possible reasons are that ChorusOS 4.0 <code>imake</code> rules have not been used to compile the program, or the link editor has found unexpected dynamic libraries.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	Inconsistent attribute flags were specified.
[EINVAL]	The address of the text or data region of the new <code>c_actor</code> is not consistent with the <code>AFX_USER_SPACE</code> or <code>AFX_SUPERVISOR_SPACE</code> attribute flag.
[EIO]	An I/O error occurred while reading from the file system.
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds <code>NAME_MAX</code> characters or the length of <i>path</i> exceeds <code>PATH_MAX</code> characters.
[ENOENT]	The new <code>c_actor</code> is a dynamic actor and one of its dependencies cannot be found either in the directories listed in the <code>LD_LIBRARY_PATH</code>

	environment variable, or in its runpath, or in /usr/lib.
[ENOEXEC]	The new c_actor file has the appropriate access permission but is not in the correct format.
[ENOTDIR]	A component of the new c_actor path of the file prefix is not a directory.
[EPERM]	The calling c_actor is not Trusted and is trying either to create a Trusted c_actor or set the new c_actor's credentials.
[ELOOP]	Too many symbolic links have been encountered during analysis of <i>path</i> .

---

**Note** - Error codes generated by *read* and *open* may also be returned.

---

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*aconf(2K)*, *acreate(2K)*, *actorCreate(2K)*, *akill(2K)*, *aload(2K)*, *astart(2K)*, *astat(2K)*, *await(2K)*, *dlopen(2K)*, *fcntl(2POSIX)*, *intro(2K)*

<b>NAME</b>	afexec, afexecl, afexecv, afexecle, afexecve, afexeclp, afexecvp – create a new <code>c_actor</code>
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int afexecve(const char * path, KnCap * cactorcap, const AcParam * param, char const * argv, char const * envp);  int afexecl(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecv(const char * path, KnCap * cactorcap, const AcParam * param, char *const * argv);  int afexecle(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */, char const * envp);  int afexeclp(const char * file, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecvp(const char * file, KnCap * cactorcap, const AcParam * param, char *const * argv);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p><i>afexec</i> family routines create a new <code>c_actor</code> with a single thread (the main thread). This main thread will execute with the priority defined by the <i>am.afexecschedprio</i> configuration parameter. It will have an automatically allocated stack, which size is defined by the <i>am.afexec.userstacksize</i> configuration parameter for user actors, and by <i>kern.exec.dflSysStackSize</i> for supervisor actors. (See <i>intro(2K)</i>).</p> <p>The main thread starts executing the runtime startup routine (usually called <i>_start</i>) specified as the program entry point in the binary header of the <i>new c_actor file</i>. After initialization of the threadsafe C library, the runtime startup routine calls the C program as follows:</p> <pre>int main (int argc, char* argv[], char* envp[])</pre> <p>where <i>argc</i> is the argument count, <i>argv</i> is an array of character pointers to the arguments themselves, and <i>envp</i> is an array of character pointers to the environment strings. As will be described later, <i>argc</i> is conventionally at least 1, and the first member of the array points to a string containing the name (or pathname) of the executable file.</p> <p><i>cactorcap</i> is the returned capability of the new <code>c_actor</code>, and is a pointer to a <i>KnCap</i> structure (see <i>actorCreate(2K)</i>).</p> <p>If <i>param</i> is <code>NULL</code>, default parameters are used, <i>acFlags</i> is set to <code>AFX_ANY_SPACE</code> and all others to un-specified. See <i>acreate(2K)</i> for a description of the <code>AcParam</code> structure.</p>

By convention, *afexec* routines search for a binary, either fully linked or relocatable, in *path*. If no such binary exists, *afexec* routines search for a file with a *.r* suffix and then with a *.O* suffix. If the *GZ\_FEATURE* is set to *true*, for each step, the *afexec* routines also search for a file with an additional *.gz* (for example: *path*, *path.gz*, *path.r*, *path.r.gz*).

*path* may also refer to a relocatable binary stored in a memory buffer in supervisor space, using the syntax "ram:0x<start>:0x<length>". <start> is the start address at which the binary is stored. <length> is the length of the binary. The buffer is not used any more after loading, except if the program uses the dynamic library API. Freeing the buffer after *afexec* may have unpredictable results if the program uses the dynamic library API.

Files to be loaded may be fully linked or relocatable files. They may also include dependencies on dynamic, but not shared, libraries. These libraries are looked for and loaded within the actor address space as part of the *afexec* call.

*arg0*, ..., *argn* are pointers to null-terminated character strings. These strings constitute the argument list available to the new *c\_actor*. By convention, at least *arg0* must be present and point to a string that is the same as *path* (or *file* or the last component of *path*).

*argv* is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new *c\_actor*. By convention, *argv* must have at least one member, and it must point to a string that is the same as *path* (or *file* or the last component of *path*). *argv* is terminated by a null pointer.

*envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new *c\_actor*. *envp* is terminated by a null pointer. For *afexecl*, *afexecv*, *afexeclp* and *afexecvp*, the C runtime start-off routine places a pointer to the environment of the calling *c\_actor* in the global object:

```
extern char** environ;
```

and it is used to pass the environment of the calling *c\_actor* to the new *c\_actor*.

*afexeclp* and *afexecvp* are called with the same arguments as *afexecl* and *afexecv*, but duplicate the shell's actions in searching for an executable file in a list of directories. The directory list is obtained from the *PATH* environment variable, and *file* must point to an executable file name for the new *c\_actor*.

## LIMITATIONS

Incompatible requests (such as *afexec* of a supervisor binary with the *AFX\_USER\_SPACE* flag set) are not normally detected. Applications can use the *AFX\_ANY\_SPACE* (see *acreate(2K)*) flag to force the system to deduce the type (*user* or *supervisor*) of the binary. In this case incompatible requests will be detected and refused.

Correct functioning of a supervisor actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) is not guaranteed. Even if the supervisor actor functions normally, only one supervisor actor can be executed in place at any one time from a given binary file. Attempting to execute more than one supervisor actor in place from the same binary file will be rejected by the system.

Correct functioning of any actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) when in flat memory mode is not guaranteed. Even if the actor functions normally, when in flat memory mode only one actor can be executed in place at any one time from a given binary file. Attempting to execute more than one actor in place in this situation will be rejected by the system.

**RETURN VALUES**

Upon successful completion, *afexec* returns a non-negative integer that is the *aid* of the new `c_actor`. Otherwise *afexec* returns `-1` and sets `errno` to indicate one of the following error conditions.

**ERRORS**

[E2BIG]	The number of bytes used by the new <code>c_actor</code> 's image argument list and environment list is greater than <code>ARG_MAX</code> bytes.
[EACCES]	Search permission is denied for a directory listed in the new <code>c_actor</code> 's file path prefix.
[EBADSYM]	The runtime link-editor cannot find a symbol needed by the new <code>c_actor</code> . Possible reasons are that ChorusOS 4.0 <code>imake</code> rules have not been used to compile the program, or the link editor has found unexpected dynamic libraries.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	Inconsistent attribute flags were specified.
[EINVAL]	The address of the text or data region of the new <code>c_actor</code> is not consistent with the <code>AFX_USER_SPACE</code> or <code>AFX_SUPERVISOR_SPACE</code> attribute flag.
[EIO]	An I/O error occurred while reading from the file system.
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds <code>NAME_MAX</code> characters or the length of <i>path</i> exceeds <code>PATH_MAX</code> characters.
[ENOENT]	The new <code>c_actor</code> is a dynamic actor and one of its dependencies cannot be found either in the directories listed in the <code>LD_LIBRARY_PATH</code>

	environment variable, or in its runpath, or in /usr/lib.
[ENOEXEC]	The new c_actor file has the appropriate access permission but is not in the correct format.
[ENOTDIR]	A component of the new c_actor path of the file prefix is not a directory.
[EPERM]	The calling c_actor is not Trusted and is trying either to create a Trusted c_actor or set the new c_actor's credentials.
[ELOOP]	Too many symbolic links have been encountered during analysis of <i>path</i> .

---

**Note** - Error codes generated by *read* and *open* may also be returned.

---

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*aconf(2K)*, *acreate(2K)*, *actorCreate(2K)*, *akill(2K)*, *aload(2K)*, *astart(2K)*, *astat(2K)*, *await(2K)*, *dlopen(2K)*, *fcntl(2POSIX)*, *intro(2K)*

<b>NAME</b>	afexec, afexecl, afexecv, afexecle, afexecve, afexecvp, afexecvp – create a new <code>c_actor</code>
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int afexecve(const char * path, KnCap * cactorcap, const AcParam * param, char const * argv, char const * envp);  int afexecl(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecv(const char * path, KnCap * cactorcap, const AcParam * param, char *const * argv);  int afexecle(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */, char const * envp);  int afexecvp(const char * file, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecvp(const char * file, KnCap * cactorcap, const AcParam * param, char *const * argv);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p><i>afexec</i> family routines create a new <code>c_actor</code> with a single thread (the main thread). This main thread will execute with the priority defined by the <i>am.afexecschedprio</i> configuration parameter. It will have an automatically allocated stack, which size is defined by the <i>am.afexec.userstacksize</i> configuration parameter for user actors, and by <i>kern.exec.dflSysStackSize</i> for supervisor actors. (See <i>intro(2K)</i>).</p> <p>The main thread starts executing the runtime startup routine (usually called <i>_start</i>) specified as the program entry point in the binary header of the <i>new c_actor file</i>. After initialization of the threadsafe C library, the runtime startup routine calls the C program as follows:</p> <pre>int main (int argc, char* argv[], char* envp[])</pre> <p>where <i>argc</i> is the argument count, <i>argv</i> is an array of character pointers to the arguments themselves, and <i>envp</i> is an array of character pointers to the environment strings. As will be described later, <i>argc</i> is conventionally at least 1, and the first member of the array points to a string containing the name (or pathname) of the executable file.</p> <p><i>cactorcap</i> is the returned capability of the new <code>c_actor</code>, and is a pointer to a <i>KnCap</i> structure (see <i>actorCreate(2K)</i>).</p> <p>If <i>param</i> is <code>NULL</code>, default parameters are used, <i>acFlags</i> is set to <code>AFX_ANY_SPACE</code> and all others to un-specified. See <i>acreate(2K)</i> for a description of the <code>AcParam</code> structure.</p>

By convention, *afexec* routines search for a binary, either fully linked or relocatable, in *path*. If no such binary exists, *afexec* routines search for a file with a *.r* suffix and then with a *.O* suffix. If the *GZ\_FEATURE* is set to *true*, for each step, the *afexec* routines also search for a file with an additional *.gz* (for example: *path*, *path.gz*, *path.r*, *path.r.gz*).

*path* may also refer to a relocatable binary stored in a memory buffer in supervisor space, using the syntax "ram:0x<start>:0x<length>". <start> is the start address at which the binary is stored. <length> is the length of the binary. The buffer is not used any more after loading, except if the program uses the dynamic library API. Freeing the buffer after *afexec* may have unpredictable results if the program uses the dynamic library API.

Files to be loaded may be fully linked or relocatable files. They may also include dependencies on dynamic, but not shared, libraries. These libraries are looked for and loaded within the actor address space as part of the *afexec* call.

*arg0*, ... , *argn* are pointers to null-terminated character strings. These strings constitute the argument list available to the new *c\_actor*. By convention, at least *arg0* must be present and point to a string that is the same as *path* (or *file* or the last component of *path*).

*argv* is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new *c\_actor*. By convention, *argv* must have at least one member, and it must point to a string that is the same as *path* (or *file* or the last component of *path*). *argv* is terminated by a null pointer.

*envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new *c\_actor*. *envp* is terminated by a null pointer. For *afexecl*, *afexecv*, *afexeclp* and *afexecvp*, the C runtime start-off routine places a pointer to the environment of the calling *c\_actor* in the global object:

```
extern char** environ;
```

and it is used to pass the environment of the calling *c\_actor* to the new *c\_actor*.

*afexeclp* and *afexecvp* are called with the same arguments as *afexecl* and *afexecv*, but duplicate the shell's actions in searching for an executable file in a list of directories. The directory list is obtained from the *PATH* environment variable, and *file* must point to an executable file name for the new *c\_actor*.

## LIMITATIONS

Incompatible requests (such as *afexec* of a supervisor binary with the *AFX\_USER\_SPACE* flag set) are not normally detected. Applications can use the *AFX\_ANY\_SPACE* (see *acreate(2K)*) flag to force the system to deduce the type (*user* or *supervisor*) of the binary. In this case incompatible requests will be detected and refused.

Correct functioning of a supervisor actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) is not guaranteed. Even if the supervisor actor functions normally, only one supervisor actor can be executed in place at any one time from a given binary file. Attempting to execute more than one supervisor actor in place from the same binary file will be rejected by the system.

Correct functioning of any actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) when in flat memory mode is not guaranteed. Even if the actor functions normally, when in flat memory mode only one actor can be executed in place at any one time from a given binary file. Attempting to execute more than one actor in place in this situation will be rejected by the system.

**RETURN VALUES**

Upon successful completion, *afexec* returns a non-negative integer that is the *aid* of the new `c_actor`. Otherwise *afexec* returns `-1` and sets `errno` to indicate one of the following error conditions.

**ERRORS**

[E2BIG]	The number of bytes used by the new <code>c_actor</code> 's image argument list and environment list is greater than <code>ARG_MAX</code> bytes.
[EACCES]	Search permission is denied for a directory listed in the new <code>c_actor</code> 's file path prefix.
[EBADSYM]	The runtime link-editor cannot find a symbol needed by the new <code>c_actor</code> . Possible reasons are that ChorusOS 4.0 <code>imake</code> rules have not been used to compile the program, or the link editor has found unexpected dynamic libraries.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	Inconsistent attribute flags were specified.
[EINVAL]	The address of the text or data region of the new <code>c_actor</code> is not consistent with the <code>AFX_USER_SPACE</code> or <code>AFX_SUPERVISOR_SPACE</code> attribute flag.
[EIO]	An I/O error occurred while reading from the file system.
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds <code>NAME_MAX</code> characters or the length of <i>path</i> exceeds <code>PATH_MAX</code> characters.
[ENOENT]	The new <code>c_actor</code> is a dynamic actor and one of its dependencies cannot be found either in the directories listed in the <code>LD_LIBRARY_PATH</code>

	environment variable, or in its runpath, or in /usr/lib.
[ENOEXEC]	The new c_actor file has the appropriate access permission but is not in the correct format.
[ENOTDIR]	A component of the new c_actor path of the file prefix is not a directory.
[EPERM]	The calling c_actor is not Trusted and is trying either to create a Trusted c_actor or set the new c_actor's credentials.
[ELOOP]	Too many symbolic links have been encountered during analysis of <i>path</i> .

---

**Note** - Error codes generated by *read* and *open* may also be returned.

---

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*aconf(2K)*, *acreate(2K)*, *actorCreate(2K)*, *akill(2K)*, *aload(2K)*, *astart(2K)*, *astat(2K)*, *await(2K)*, *dlopen(2K)*, *fcntl(2POSIX)*, *intro(2K)*

<b>NAME</b>	afexec, afexecl, afexecv, afexecle, afexecve, afexeclp, afexecvp – create a new <code>c_actor</code>
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int afexecve(const char * path, KnCap * cactorcap, const AcParam * param, char const * argv, char const * envp);  int afexecl(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecv(const char * path, KnCap * cactorcap, const AcParam * param, char *const * argv);  int afexecle(const char * path, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */, char const * envp);  int afexeclp(const char * file, KnCap * cactorcap, const AcParam * param, const char * arg0, ..., const char * argn, char * /*NULL */);  int afexecvp(const char * file, KnCap * cactorcap, const AcParam * param, char *const * argv);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p><i>afexec</i> family routines create a new <code>c_actor</code> with a single thread (the main thread). This main thread will execute with the priority defined by the <i>am.afexecschedprio</i> configuration parameter. It will have an automatically allocated stack, which size is defined by the <i>am.afexec.userstacksize</i> configuration parameter for user actors, and by <i>kern.exec.dflSysStackSize</i> for supervisor actors. (See <i>intro(2K)</i>).</p> <p>The main thread starts executing the runtime startup routine (usually called <i>_start</i>) specified as the program entry point in the binary header of the <i>new c_actor file</i>. After initialization of the threadsafe C library, the runtime startup routine calls the C program as follows:</p> <pre>int main (int argc, char* argv[], char* envp[])</pre> <p>where <i>argc</i> is the argument count, <i>argv</i> is an array of character pointers to the arguments themselves, and <i>envp</i> is an array of character pointers to the environment strings. As will be described later, <i>argc</i> is conventionally at least 1, and the first member of the array points to a string containing the name (or pathname) of the executable file.</p> <p><i>cactorcap</i> is the returned capability of the new <code>c_actor</code>, and is a pointer to a <i>KnCap</i> structure (see <i>actorCreate(2K)</i>).</p> <p>If <i>param</i> is <code>NULL</code>, default parameters are used, <i>acFlags</i> is set to <code>AFX_ANY_SPACE</code> and all others to un-specified. See <i>acreate(2K)</i> for a description of the <code>AcParam</code> structure.</p>

By convention, *afexec* routines search for a binary, either fully linked or relocatable, in *path*. If no such binary exists, *afexec* routines search for a file with a *.r* suffix and then with a *.O* suffix. If the *GZ\_FEATURE* is set to *true*, for each step, the *afexec* routines also search for a file with an additional *.gz* (for example: *path*, *path.gz*, *path.r*, *path.r.gz*).

*path* may also refer to a relocatable binary stored in a memory buffer in supervisor space, using the syntax "ram:0x<start>:0x<length>". <start> is the start address at which the binary is stored. <length> is the length of the binary. The buffer is not used any more after loading, except if the program uses the dynamic library API. Freeing the buffer after *afexec* may have unpredictable results if the program uses the dynamic library API.

Files to be loaded may be fully linked or relocatable files. They may also include dependencies on dynamic, but not shared, libraries. These libraries are looked for and loaded within the actor address space as part of the *afexec* call.

*arg0*, ... , *argn* are pointers to null-terminated character strings. These strings constitute the argument list available to the new *c\_actor*. By convention, at least *arg0* must be present and point to a string that is the same as *path* (or *file* or the last component of *path*).

*argv* is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new *c\_actor*. By convention, *argv* must have at least one member, and it must point to a string that is the same as *path* (or *file* or the last component of *path*). *argv* is terminated by a null pointer.

*envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new *c\_actor*. *envp* is terminated by a null pointer. For *afexecl*, *afexecv*, *afexeclp* and *afexecvp*, the C runtime start-off routine places a pointer to the environment of the calling *c\_actor* in the global object:

```
extern char** environ;
```

and it is used to pass the environment of the calling *c\_actor* to the new *c\_actor*.

*afexeclp* and *afexecvp* are called with the same arguments as *afexecl* and *afexecv*, but duplicate the shell's actions in searching for an executable file in a list of directories. The directory list is obtained from the *PATH* environment variable, and *file* must point to an executable file name for the new *c\_actor*.

## LIMITATIONS

Incompatible requests (such as *afexec* of a supervisor binary with the *AFX\_USER\_SPACE* flag set) are not normally detected. Applications can use the *AFX\_ANY\_SPACE* (see *acreate(2K)*) flag to force the system to deduce the type (*user* or *supervisor*) of the binary. In this case incompatible requests will be detected and refused.

Correct functioning of a supervisor actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) is not guaranteed. Even if the supervisor actor functions normally, only one supervisor actor can be executed in place at any one time from a given binary file. Attempting to execute more than one supervisor actor in place from the same binary file will be rejected by the system.

Correct functioning of any actor executed in place (using the `AFX_EXECUTE_IN_PLACE` flag) when in flat memory mode is not guaranteed. Even if the actor functions normally, when in flat memory mode only one actor can be executed in place at any one time from a given binary file. Attempting to execute more than one actor in place in this situation will be rejected by the system.

**RETURN VALUES**

Upon successful completion, *afexec* returns a non-negative integer that is the *aid* of the new `c_actor`. Otherwise *afexec* returns `-1` and sets `errno` to indicate one of the following error conditions.

**ERRORS**

[E2BIG]	The number of bytes used by the new <code>c_actor</code> 's image argument list and environment list is greater than <code>ARG_MAX</code> bytes.
[EACCES]	Search permission is denied for a directory listed in the new <code>c_actor</code> 's file path prefix.
[EBADSYM]	The runtime link-editor cannot find a symbol needed by the new <code>c_actor</code> . Possible reasons are that ChorusOS 4.0 <code>imake</code> rules have not been used to compile the program, or the link editor has found unexpected dynamic libraries.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	Inconsistent attribute flags were specified.
[EINVAL]	The address of the text or data region of the new <code>c_actor</code> is not consistent with the <code>AFX_USER_SPACE</code> or <code>AFX_SUPERVISOR_SPACE</code> attribute flag.
[EIO]	An I/O error occurred while reading from the file system.
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds <code>NAME_MAX</code> characters or the length of <i>path</i> exceeds <code>PATH_MAX</code> characters.
[ENOENT]	The new <code>c_actor</code> is a dynamic actor and one of its dependencies cannot be found either in the directories listed in the <code>LD_LIBRARY_PATH</code>

	environment variable, or in its runpath, or in /usr/lib.
[ENOEXEC]	The new c_actor file has the appropriate access permission but is not in the correct format.
[ENOTDIR]	A component of the new c_actor path of the file prefix is not a directory.
[EPERM]	The calling c_actor is not Trusted and is trying either to create a Trusted c_actor or set the new c_actor's credentials.
[ELOOP]	Too many symbolic links have been encountered during analysis of <i>path</i> .

---

**Note** - Error codes generated by *read* and *open* may also be returned.

---

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`aconf(2K)`, `acreate(2K)`, `actorCreate(2K)`, `akill(2K)`, `aload(2K)`, `astart(2K)`, `astat(2K)`, `await(2K)`, `dlopen(2K)`, `fcntl(2POSIX)`, `intro(2K)`

<b>NAME</b>	aload, alParamBuild, alParamUnpack, agetalparam – Manage the loading of an actor
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int <b>aload</b>(const KnCap * <i>cactorcap</i>, const char * <i>path</i>, const AlParam * <i>ldParam</i>, KnPc <i>entry</i>);  void <b>alParamBuild</b>(void * <i>buffer</i>, unsigned int <i>bufferSize</i>, const char *const * <i>argv</i>, const char *const * <i>envp</i>);  int <b>alParamUnpack</b>(const Alparam * <i>alParamp</i>, int * <i>argcp</i>, char *** <i>argvp</i>);  unsigned int <b>alParamSize</b>(const Alparam * <i>alParamp</i>);  int <b>agetalparam</b>(const AlParam ** <i>argp</i>);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p>The <code>aload( )</code> call initializes the text and data regions of the target actor specified by <code>cactorcap</code>. This call also initializes the environment and arguments of the target actor.</p> <p>The target actor is loaded from an executable object file whose pathname is specified by <code>path</code>. This file consists of a binary header, a text segment, and the initialized part of the data segment. The file may be in any of the executable file formats supported by the system.</p> <p>The <code>param</code> argument points to an <code>AlParam</code> data structure, which is a packed representation of the environment, and arguments which will be copied to a region of the target actor address space. The address of this region can subsequently be obtained by the target actor using the <code>agetalparam( )</code> call. This region is freed when the target actor starts executing (see <code>start(2K)</code>), after the <code>args</code> and <code>env</code> have been installed in the actor address space (heap or stack).</p> <p>An <code>AlParam</code> structure can be built using <code>alParamBuild( )</code>.</p> <p>If <code>path</code> is a <code>NULL</code> pointer, no executable file is loaded.</p> <p>If <code>ldParam</code> is a <code>NULL</code> pointer, the packed arguments and environment currently installed in the actor are not affected.</p> <p>If <code>entry</code> is not a <code>NULL</code> pointer, it must point to a valid location large enough to store a <code>KnPc</code> object. In that case, the entry-point as determined by the loader is returned into the location pointed to by <code>entry</code>.</p> <p>The <code>aload( )</code> call can be used on any actor as many times as necessary. It is up to the caller to ensure that the address ranges to be used in the target actor are actually free, and that the target actor is in a state where it is safe to modify its address space.</p> <p>The <code>alParamBuild( )</code> command builds an <code>AlParam</code> structure from the <code>argv</code> and <code>envp</code> arrays provided, and stores the result in the buffer <code>buffer</code>. The <code>buffer</code></p>

field must point to an area of memory large enough to hold the arguments and environment described by *argv* and *envp*.

The `alParamUnpack()` command retrieves the arguments and environment from the data structure pointed to by *alParam*, copies them to the newly allocated memory, builds the environment string pointer arrays and builds the argument's string pointers array whose addresses are stored in the areas pointed to by *argvp*. The number of arguments is returned into the location pointed to by *argcp*.

Do not build an `AlParam` data structure containing more than `ARG_MAX` bytes of data. The only limit is the amount of memory available when unpacking it, and the amount of memory available in the system to store a copy of it in the *restart parameters* of a new actor.

The *alparamsize* field shows the total size used of the `AlParam` object referred to by *alParam*. The size returned accounts for the `AlParam` object and the strings that follow.

The `agetalparam()` call returns, at the location pointed to by *argp*, the address where the packed arguments and environment are stored in the calling actor's address space.

Note that `agetalparam()` and `alParamUnpack()` are not normally used by the application; they are used by the C library startup code to build the arguments and environment in the traditional UNIX representation. The packed arguments and environment themselves are not affected, nor used by, the `getenv()` and `setenv()` library routines. If permitted by the nucleus configuration, the region containing the packed arguments and environment will be read-only.

#### LIMITATIONS

It should be noted that `aload()` is usually the first event that allocates regions to an actor created using `acreate()`.

#### BUGS

When installing the new packed arguments and environment, `aload()` does not free the packed arguments and environment region resulting from a previous call to `aload()`.

#### RETURN VALUE

Upon successful completion, `aload()` returns 0. Otherwise `aload()` returns -1 and sets `errno` to indicate one of the following error conditions:

[ENOENT]	One or more components of the new actor's pathname do not exist.
[ENOTDIR]	A component of the new actor's pathname prefix is not a directory.
[EACCES]	Search permission is denied for a directory listed in the new actor file's path prefix.

[ENOEXEC]	The new actor file has the appropriate access permission but is not in the proper format.
[EFAULT]	The new actor file is not as long as indicated by the size values in its header.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	The address of the text or data region of the new actor is not consistent with the new actor privilege (SUPERVISOR, USER, SYSTEM).
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds NAME_MAX characters or the length of <i>path</i> exceeds PATH_MAX characters.
[ELOOP]	Too many symbolic links have been encountered during analysis of the file.
[EIO]	An I/O error occurred while reading from or writing to the file system.
[ESRCH]	The actor referred to by the given capability does not exist.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`acreate(2K)` , `afexec(2K)` , `akill(2K)` , `astart(2K)` , `hrfexec(2RESTART)`

**NAME** agetId – get c\_actor’s ID

**SYNOPSIS** #include <am/afexec.h>  
int agetId(void);

**FEATURES** ACTOR\_EXTENDED\_MNGT

**RETURN VALUE** The agetId( ) call returns the current c\_actor’s ID (see intro(2K)).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** afexec(2K), acap(2K)

<b>NAME</b>	akill – kill or restart an actor				
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int <b>akill</b>(const KnCap *cactorcap);</pre>				
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT				
<b>DESCRIPTION</b>	The <code>akill()</code> call stops and deletes the actor identified by <i>cactorcap</i> (see <code>afexec(2K)</code> ), unless the actor is a restartable actor. If the actor is a restartable actor, calling <code>akill()</code> simply restarts the actor, as if it had terminated abnormally. To permanently kill (and not simply restart) a restartable actor, call <code>hrKillGroup(2RESTART)</code> .				
<b>RETURN VALUES</b>	The <code>akill()</code> call returns 0 if the actor was killed successfully (even if it restarts). Otherwise, <code>akill()</code> returns -1 and sets <code>errno</code> to indicate one of the following error conditions: <ul style="list-style-type: none"> <li>[EFAULT] <i>cactorcap</i> points to an illegal address.</li> <li>[ESRCH] <i>cactorcap</i> points to an illegal capability or the actor with the capability given doesn't exist.</li> <li>[EPERM] the calling actor is not a trusted actor or a supervisor actor, and the user-id of the calling actor does not match the user-id of the killed actor.</li> </ul>				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> 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;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>afexec(2K)</code> , <code>await(2K)</code> , <code>hrKillGroup(2RESTART)</code>				

<b>NAME</b>	aload, alParamBuild, alParamUnpack, agetalparam – Manage the loading of an actor
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int <b>aload</b>(const KnCap * <i>cactorcap</i>, const char * <i>path</i>, const AlParam * <i>ldParam</i>, KnPc <i>entry</i>);  void <b>alParamBuild</b>(void * <i>buffer</i>, unsigned int <i>bufferSize</i>, const char *const * <i>argv</i>, const char *const * <i>envp</i>);  int <b>alParamUnpack</b>(const Alparam * <i>alParamp</i>, int * <i>argcp</i>, char *** <i>argvp</i>);  unsigned int <b>alParamSize</b>(const Alparam * <i>alParamp</i>);  int <b>agetalparam</b>(const AlParam ** <i>argp</i>);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p>The <code>aload( )</code> call initializes the text and data regions of the target actor specified by <i>cactorcap</i> . This call also initializes the environment and arguments of the target actor.</p> <p>The target actor is loaded from an executable object file whose pathname is specified by <i>path</i> . This file consists of a binary header, a text segment, and the initialized part of the data segment. The file may be in any of the executable file formats supported by the system.</p> <p>The <i>param</i> argument points to an <i>AlParam</i> data structure, which is a packed representation of the environment, and arguments which will be copied to a region of the target actor address space. The address of this region can subsequently be obtained by the target actor using the <code>agetalparam( )</code> call. This region is freed when the target actor starts executing (see <code>start(2K)</code> ), after the <i>args</i> and <i>env</i> have been installed in the actor address space (heap or stack).</p> <p>An <i>AlParam</i> structure can be built using <code>alParamBuild( )</code> .</p> <p>If <i>path</i> is a NULL pointer, no executable file is loaded.</p> <p>If <i>ldParam</i> is a NULL pointer, the packed arguments and environment currently installed in the actor are not affected.</p> <p>If <i>entry</i> is not a NULL pointer, it must point to a valid location large enough to store a KnPc object. In that case, the entry-point as determined by the loader is returned into the location pointed to by <i>entry</i> .</p> <p>The <code>aload( )</code> call can be used on any actor as many times as necessary. It is up to the caller to ensure that the address ranges to be used in the target actor are actually free, and that the target actor is in a state where it is safe to modify its address space.</p> <p>The <code>alParamBuild( )</code> command builds an <i>AlParam</i> structure from the <i>argv</i> and <i>envp</i> arrays provided, and stores the result in the buffer <i>buffer</i> . The <i>buffer</i></p>

field must point to an area of memory large enough to hold the arguments and environment described by *argv* and *envp* .

The `alParamUnpack()` command retrieves the arguments and environment from the data structure pointed to by *alParam* , copies them to the newly allocated memory, builds the environment string pointer arrays and builds the argument's string pointers array whose addresses are stored in the areas pointed to by *argv* . The number of arguments is returned into the location pointed to by *argcp* .

Do not build an `AlParam` data structure containing more than `ARG_MAX` bytes of data. The only limit is the amount of memory available when unpacking it, and the amount of memory available in the system to store a copy of it in the *restart parameters* of a new actor.

The *alparamsize* field shows the total size used of the `AlParam` object referred to by *alParam* . The size returned accounts for the `AlParam` object and the strings that follow.

The `agetalparam()` call returns, at the location pointed to by *argp* , the address where the packed arguments and environment are stored in the calling actor's address space.

Note that `agetalparam()` and `alParamUnpack()` are not normally used by the application; they are used by the C library startup code to build the arguments and environment in the traditional UNIX representation. The packed arguments and environment themselves are not affected, nor used by, the `getenv()` and `setenv()` library routines. If permitted by the nucleus configuration, the region containing the packed arguments and environment will be read-only.

#### LIMITATIONS

It should be noted that `aload()` is usually the first event that allocates regions to an actor created using `acreate()` .

#### BUGS

When installing the new packed arguments and environment, `aload()` does not free the packed arguments and environment region resulting from a previous call to `aload()` .

#### RETURN VALUE

Upon successful completion, `aload()` returns 0 . Otherwise `aload()` returns -1 and sets `errno` to indicate one of the following error conditions:

[ENOENT]	One or more components of the new actor's pathname do not exist.
[ENOTDIR]	A component of the new actor's pathname prefix is not a directory.
[EACCES]	Search permission is denied for a directory listed in the new actor file's path prefix.

- [ENOEXEC]            The new actor file has the appropriate access permission but is not in the proper format.
- [EFAULT]            The new actor file is not as long as indicated by the size values in its header.
- [EFAULT]            *path* , *argv* , or *envp* point to an illegal address.
- [EINVAL]            The address of the text or data region of the new actor is not consistent with the new actor privilege (SUPERVISOR, USER, SYSTEM).
- [ENAMETOOLONG]    The length of a component of *path* exceeds NAME\_MAX characters or the length of *path* exceeds PATH\_MAX characters.
- [ELOOP]            Too many symbolic links have been encountered during analysis of the file.
- [EIO]                An I/O error occurred while reading from or writing to the file system.
- [ESRCH]            The actor referred to by the given capability does not exist.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`acreate(2K)` , `afexec(2K)` , `akill(2K)` , `astart(2K)` , `hrfexec(2RESTART)`

<b>NAME</b>	aload, alParamBuild, alParamUnpack, agetalparam – Manage the loading of an actor
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int <b>aload</b>(const KnCap * <i>cactorcap</i>, const char * <i>path</i>, const AlParam * <i>ldParam</i>, KnPc <i>entry</i>);  void <b>alParamBuild</b>(void * <i>buffer</i>, unsigned int <i>bufferSize</i>, const char *const * <i>argv</i>, const char *const * <i>envp</i>);  int <b>alParamUnpack</b>(const Alparam * <i>alParamp</i>, int * <i>argcp</i>, char *** <i>argvp</i>);  unsigned int <b>alParamSize</b>(const Alparam * <i>alParamp</i>);  int <b>agetalparam</b>(const AlParam ** <i>argp</i>);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p>The <code>aload()</code> call initializes the text and data regions of the target actor specified by <code>cactorcap</code>. This call also initializes the environment and arguments of the target actor.</p> <p>The target actor is loaded from an executable object file whose pathname is specified by <code>path</code>. This file consists of a binary header, a text segment, and the initialized part of the data segment. The file may be in any of the executable file formats supported by the system.</p> <p>The <code>param</code> argument points to an <code>AlParam</code> data structure, which is a packed representation of the environment, and arguments which will be copied to a region of the target actor address space. The address of this region can subsequently be obtained by the target actor using the <code>agetalparam()</code> call. This region is freed when the target actor starts executing (see <code>start(2K)</code>), after the <code>args</code> and <code>env</code> have been installed in the actor address space (heap or stack).</p> <p>An <code>AlParam</code> structure can be built using <code>alParamBuild()</code>.</p> <p>If <code>path</code> is a <code>NULL</code> pointer, no executable file is loaded.</p> <p>If <code>ldParam</code> is a <code>NULL</code> pointer, the packed arguments and environment currently installed in the actor are not affected.</p> <p>If <code>entry</code> is not a <code>NULL</code> pointer, it must point to a valid location large enough to store a <code>KnPc</code> object. In that case, the entry-point as determined by the loader is returned into the location pointed to by <code>entry</code>.</p> <p>The <code>aload()</code> call can be used on any actor as many times as necessary. It is up to the caller to ensure that the address ranges to be used in the target actor are actually free, and that the target actor is in a state where it is safe to modify its address space.</p> <p>The <code>alParamBuild()</code> command builds an <code>AlParam</code> structure from the <code>argv</code> and <code>envp</code> arrays provided, and stores the result in the buffer <code>buffer</code>. The <code>buffer</code></p>

field must point to an area of memory large enough to hold the arguments and environment described by *argv* and *envp* .

The `alParamUnpack()` command retrieves the arguments and environment from the data structure pointed to by *alParam* , copies them to the newly allocated memory, builds the environment string pointer arrays and builds the argument's string pointers array whose addresses are stored in the areas pointed to by *argvp* . The number of arguments is returned into the location pointed to by *argcp* .

Do not build an `AlParam` data structure containing more than `ARG_MAX` bytes of data. The only limit is the amount of memory available when unpacking it, and the amount of memory available in the system to store a copy of it in the *restart parameters* of a new actor.

The *alparamsize* field shows the total size used of the `AlParam` object referred to by *alParam* . The size returned accounts for the `AlParam` object and the strings that follow.

The `agetalparam()` call returns, at the location pointed to by *argp* , the address where the packed arguments and environment are stored in the calling actor's address space.

Note that `agetalparam()` and `alParamUnpack()` are not normally used by the application; they are used by the C library startup code to build the arguments and environment in the traditional UNIX representation. The packed arguments and environment themselves are not affected, nor used by, the `getenv()` and `setenv()` library routines. If permitted by the nucleus configuration, the region containing the packed arguments and environment will be read-only.

#### LIMITATIONS

It should be noted that `aload()` is usually the first event that allocates regions to an actor created using `acreate()` .

#### BUGS

When installing the new packed arguments and environment, `aload()` does not free the packed arguments and environment region resulting from a previous call to `aload()` .

#### RETURN VALUE

Upon successful completion, `aload()` returns 0 . Otherwise `aload()` returns -1 and sets `errno` to indicate one of the following error conditions:

[ENOENT]	One or more components of the new actor's pathname do not exist.
[ENOTDIR]	A component of the new actor's pathname prefix is not a directory.
[EACCES]	Search permission is denied for a directory listed in the new actor file's path prefix.

[ENOEXEC]	The new actor file has the appropriate access permission but is not in the proper format.
[EFAULT]	The new actor file is not as long as indicated by the size values in its header.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	The address of the text or data region of the new actor is not consistent with the new actor privilege (SUPERVISOR, USER, SYSTEM).
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds NAME_MAX characters or the length of <i>path</i> exceeds PATH_MAX characters.
[ELOOP]	Too many symbolic links have been encountered during analysis of the file.
[EIO]	An I/O error occurred while reading from or writing to the file system.
[ESRCH]	The actor referred to by the given capability does not exist.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`acreate(2K)` , `afexec(2K)` , `akill(2K)` , `astart(2K)` , `hrfexec(2RESTART)`

<b>NAME</b>	aload, alParamBuild, alParamUnpack, agetalparam – Manage the loading of an actor
<b>SYNOPSIS</b>	<pre>#include &lt;am/afexec.h&gt; int <b>aload</b>(const KnCap * <i>cactorcap</i>, const char * <i>path</i>, const AlParam * <i>ldParam</i>, KnPc <i>entry</i>);  void <b>alParamBuild</b>(void * <i>buffer</i>, unsigned int <i>bufferSize</i>, const char *const * <i>argv</i>, const char *const * <i>envp</i>);  int <b>alParamUnpack</b>(const Alparam * <i>alParamp</i>, int * <i>argcp</i>, char *** <i>argvp</i>);  unsigned int <b>alParamSize</b>(const Alparam * <i>alParamp</i>);  int <b>agetalparam</b>(const AlParam ** <i>argp</i>);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p>The <code>aload( )</code> call initializes the text and data regions of the target actor specified by <code>cactorcap</code>. This call also initializes the environment and arguments of the target actor.</p> <p>The target actor is loaded from an executable object file whose pathname is specified by <code>path</code>. This file consists of a binary header, a text segment, and the initialized part of the data segment. The file may be in any of the executable file formats supported by the system.</p> <p>The <code>param</code> argument points to an <code>AlParam</code> data structure, which is a packed representation of the environment, and arguments which will be copied to a region of the target actor address space. The address of this region can subsequently be obtained by the target actor using the <code>agetalparam( )</code> call. This region is freed when the target actor starts executing (see <code>start(2K)</code>), after the <code>args</code> and <code>env</code> have been installed in the actor address space (heap or stack).</p> <p>An <code>AlParam</code> structure can be built using <code>alParamBuild( )</code>.</p> <p>If <code>path</code> is a <code>NULL</code> pointer, no executable file is loaded.</p> <p>If <code>ldParam</code> is a <code>NULL</code> pointer, the packed arguments and environment currently installed in the actor are not affected.</p> <p>If <code>entry</code> is not a <code>NULL</code> pointer, it must point to a valid location large enough to store a <code>KnPc</code> object. In that case, the entry-point as determined by the loader is returned into the location pointed to by <code>entry</code>.</p> <p>The <code>aload( )</code> call can be used on any actor as many times as necessary. It is up to the caller to ensure that the address ranges to be used in the target actor are actually free, and that the target actor is in a state where it is safe to modify its address space.</p> <p>The <code>alParamBuild( )</code> command builds an <code>AlParam</code> structure from the <code>argv</code> and <code>envp</code> arrays provided, and stores the result in the buffer <code>buffer</code>. The <code>buffer</code></p>

field must point to an area of memory large enough to hold the arguments and environment described by *argv* and *envp* .

The `alParamUnpack()` command retrieves the arguments and environment from the data structure pointed to by *alParam* , copies them to the newly allocated memory, builds the environment string pointer arrays and builds the argument's string pointers array whose addresses are stored in the areas pointed to by *argvp* . The number of arguments is returned into the location pointed to by *argcp* .

Do not build an `AlParam` data structure containing more than `ARG_MAX` bytes of data. The only limit is the amount of memory available when unpacking it, and the amount of memory available in the system to store a copy of it in the *restart parameters* of a new actor.

The *alparamsize* field shows the total size used of the `AlParam` object referred to by *alParam* . The size returned accounts for the `AlParam` object and the strings that follow.

The `agetalparam()` call returns, at the location pointed to by *argp* , the address where the packed arguments and environment are stored in the calling actor's address space.

Note that `agetalparam()` and `alParamUnpack()` are not normally used by the application; they are used by the C library startup code to build the arguments and environment in the traditional UNIX representation. The packed arguments and environment themselves are not affected, nor used by, the `getenv()` and `setenv()` library routines. If permitted by the nucleus configuration, the region containing the packed arguments and environment will be read-only.

#### LIMITATIONS

It should be noted that `aload()` is usually the first event that allocates regions to an actor created using `acreate()` .

#### BUGS

When installing the new packed arguments and environment, `aload()` does not free the packed arguments and environment region resulting from a previous call to `aload()` .

#### RETURN VALUE

Upon successful completion, `aload()` returns 0 . Otherwise `aload()` returns -1 and sets `errno` to indicate one of the following error conditions:

[ENOENT]	One or more components of the new actor's pathname do not exist.
[ENOTDIR]	A component of the new actor's pathname prefix is not a directory.
[EACCES]	Search permission is denied for a directory listed in the new actor file's path prefix.

[ENOEXEC]	The new actor file has the appropriate access permission but is not in the proper format.
[EFAULT]	The new actor file is not as long as indicated by the size values in its header.
[EFAULT]	<i>path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
[EINVAL]	The address of the text or data region of the new actor is not consistent with the new actor privilege (SUPERVISOR, USER, SYSTEM).
[ENAMETOOLONG]	The length of a component of <i>path</i> exceeds NAME_MAX characters or the length of <i>path</i> exceeds PATH_MAX characters.
[ELOOP]	Too many symbolic links have been encountered during analysis of the file.
[EIO]	An I/O error occurred while reading from or writing to the file system.
[ESRCH]	The actor referred to by the given capability does not exist.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`acreate(2K)` , `afexec(2K)` , `akill(2K)` , `astart(2K)` , `hrfexec(2RESTART)`

<b>NAME</b>	astart – activates a <code>c_actor</code>				
<b>SYNOPSIS</b>	<code>#include &lt;am/afexec.h&gt;</code> <code>int astart(const KnCap *cactorcap);</code>				
<b>FEATURES</b>	RESTART_EXTENDED				
<b>DESCRIPTION</b>	<p>The <code>astart()</code> call creates one thread in the <code>c_actor</code> specified by <code>cactorcap</code>. That thread executes the <code>c_actor</code>'s text at the entry-point defined by the executable file header determined during the last use of the <code>aload(2K)</code> operation applied to that <code>c_actor</code>.</p> <p>The <code>cactorcap</code> pointer indicates the capability of the <code>c_actor</code> to be activated.</p> <p>Although unlikely, it is possible that you may want to start a number of threads at the actor's entry point by calling <code>astart()</code> repeatedly.</p>				
<b>LIMITATIONS</b>	The <code>astart()</code> cannot act on a remote <code>c_actor</code> . If the capability specified refers to a remote actor, this call fails and returns the <code>ENOTIMPLEMENTED</code> error code.				
<b>RETURN VALUE</b>	<p>Upon successful completion, these routines return 0. Otherwise they return -1 and set <code>errno</code> to indicate the following error condition:</p> <p>[ESRCH]                      The <code>c_actor</code> referred to by the given capability does not exist.</p>				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>acreate(2K)</code> , <code>afexec(2K)</code> , <code>akill(2K)</code> , <code>aload(2K)</code>				

**NAME**            astat – list all active c\_actors

**SYNOPSIS**        #include <am/astat.h>  
int **astat**(astatEntry \*entries, int maxEntries);

**FEATURES**        ACTOR\_EXTENDED\_MNGT

**DESCRIPTION**    The `astat()` call fills an entry defined by the `entries` array for each `c_actor` present. The maximum number of entries is defined by `maxEntries`.

The `entries` array is a structure containing the following members:

```
int     astatFlags;      /* c_actor's status */
int     astatAid;       /* c_actor's ID */
int     astatUid;       /* c_actor's user ID */
char    astatName[];    /* c_actor's name */
```

The status information `astatFlags` is bit-encoded using the following bits:

```
ASTAT_DEBUGGED_CACTOR  the c_actor is being debugged
```

**RETURN VALUE**    Upon successful completion, `astat` returns the number of active `c_actors`; otherwise it returns `-1` and sets `errno` to indicate the following error condition:

[EFAULT]                    `entries` points outside the allocated address space of the `c_actor`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**        `afexec(2K)`

<b>NAME</b>	atrace – c_actor trace
<b>SYNOPSIS</b>	<pre>#include &lt;chorus.h&gt; #include &lt;am/reg.h&gt; #include &lt;am/atrace.h&gt; int <b>atrace</b>(int <i>request</i>, const KnCap *<i>cactorcap</i>, char *<i>addr</i>, int <i>data</i>, char *<i>addr2</i>);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT
<b>DESCRIPTION</b>	<p>The <code>atrace( )</code> call allows a <code>c_actor</code> to control the execution of another <code>c_actor</code>. This system call is similar to the <code>ptrace( )</code> system call in UNIX. Its primary use is for the implementation of breakpoint debugging. The traced <code>c_actor</code> behaves normally until it encounters an exception, at which point it enters a stopped state and the debugger is notified using <code>awaits(2K)</code>.</p> <p>The system maintains the concept of a <code>target thread</code>, which is the identifier of the thread that encountered the exception. Some <code>atrace( )</code> requests apply to the thread identified by the value of the <code>target thread</code>. (<code>ATRACE_SINGLESTEP</code>, <code>ATRACE_SETREGS</code> and <code>ATRACE_GETREGS</code>.) The <code>target thread</code> is (re)set by the system as soon as a thread causes the <code>c_actor</code> to stop. It can also be reset by the debugger using a specific <code>atrace( )</code> request.</p> <p>When the <code>c_actor</code> is stopped, the debugger can examine and modify its <i>core image</i> using <code>atrace( )</code>. The debugger can also make the <code>c_actor</code> terminate or continue. By using the <code>ATRACE_THREADSUSPEND</code> and <code>ATRACE_THREADRESUME</code> requests before making the <code>c_actor</code> continue execution using <code>ATRACE_CONT</code>, the debugger can restart a subset of the threads within the <code>c_actor</code>.</p> <p>The <i>request</i> argument determines the precise action to be taken by <code>atrace( )</code>. For each request, <i>cactorcap</i> must designate a debugged <code>c_actor</code>. If the semantics of the <i>addr</i>, <i>data</i> and/or <i>addr2</i> arguments are not specified for a request, those arguments are ignored. The debugged <code>c_actor</code> must be stopped before the requests can be made.</p> <p><code>ATRACE_PEEKTEXT</code>, <code>ATRACE_PEEKDATA</code></p> <p>These requests return the word at the location <i>addr</i> in the address space of the <code>c_actor</code> to the debugger. These two requests will fail if <i>addr</i> is not the start address of a word or is outside the <code>c_actor</code>'s address space; in this case, a value of <code>-1</code> is returned to the debugger and <code>errno</code> is set to <code>EIO</code>.</p> <p><code>ATRACE_POKETEXT</code>, <code>ATRACE_POKEDATA</code></p> <p>These requests put the value given by the <i>data</i> argument into the address space of the <code>c_actor</code> at the location <i>addr</i>. Upon successful completion, the value written into the address space of the <code>c_actor</code> is returned to the debugger. These two requests will fail if <i>addr</i> is not the start address of a word or is outside the <code>c_actor</code>'s address space; in this case, a value of <code>-1</code> is returned to the debugger and <code>errno</code> is set to <code>EFAULT</code>.</p>

**ATRACE\_CONT**

This request causes the `c_actor` to resume execution. The target thread resumes execution at the address `addr`. All other threads will also resume execution, unless suspended by `ATRACE_THREADSUSPEND`. For this request, the `data` argument must be equal to 0. Upon successful completion, 0 is returned to the debugger; otherwise -1 is returned and `errno` is set to `EFAULT`.

**ATRACE\_KILL**

This request causes the `c_actor` to terminate with the same consequences as `_exit(2K)`.

**ATRACE\_SINGLESTEP**

This request sets the trace bit in the Processor Status Word of the target thread, it then executes the same steps listed above for the `ATRACE_CONT` request, with the exception that only the target thread will be restarted. All other threads remain in a stopped state. The trace bit causes an exception upon completion of one machine instruction, which allows single-stepping of the target thread.

**ATRACE\_ATTACH**

Start tracing a running `c_actor`, designated by `cactorcap`.

**ATRACE\_DETACH**

Stop tracing the `c_actor` designated by `cactorcap`. The `c_actor` continues its execution from the address `addr`. If `addr` is defined as `(char *) 1` the execution continues from where it stopped. The `data` argument must be 0.

**ATRACE\_GETREGS, ATRACE\_SETREGS**

`ATRACE_GETREGS` returns the registers of the target thread in the `c_actor` designated by `cactorcap` to a table pointed to by `addr`. The registers table type is `regs` as described in `<am/reg.h>`. `ATRACE_SETREGS` writes the registers of the target thread in the `c_actor` designated by `cactorcap` from a table pointed to by `addr`. The registers table type is `regs` as described in `<am/reg.h>`.

**ATRACE\_GETTHREADREGS, ATRACE\_SETTHREADREGS**

`ATRACE_GETTHREADREGS` returns the registers of the thread identified by `data` in the `c_actor` designated by `cactorcap` to a structure pointed to by `addr`. The registers table type is `regs` as described in `<am/reg.h>`. `ATRACE_SETTHREADREGS` writes the registers of the thread identified by `data` in the `c_actor` designated by `cactorcap` from a structure pointed to by `addr`. The registers table type is `regs` as described in `<am/reg.h>`.

**ATRACE\_GETTARGETTHREAD**

The identifier of the current target thread is returned.

**ATTRACE\_SETTARGETTHREAD**

Set the target thread to the thread identified by the *data* argument.

**ATTRACE\_THREADSUSPEND**

Suspend the thread identified by *data* (see `threadSuspend(2K)`). If that thread is the thread which caused the `c_actor` to stop, the thread will only be suspended at the time the debugger restarts the `c_actor` using `ATTRACE_CONT` or `ATTRACE_DETACH`.

**ATTRACE\_THREADRESUME**

Resume the thread identified by *data* (see `threadResume(2K)`).

**ATTRACE\_THREADLIST**

Return the number of threads in the `c_actor` designated by *cactorcap*. For each thread within the `c_actor`, return the `KnThreadStat` structure in an array pointed to by *addr*. The number of bytes copied to *addr* is limited to the amount specified in *data* (see `threadStat(2K)`).

**ATTRACE\_GETTHREADNAME, ATTRACE\_SETTHREADNAME**

Get or set the symbolic name of the thread identified by *data*.

Note: `-1` can be a legitimate return value of `atrace()` in the case of a `PEEK` request. Therefore, the only way to check whether `atrace()` has failed is to set `errno` to `0` before calling `atrace()` and then check the value of the `errno` variable.

**RETURN VALUES**

If `atrace()` fails, it will set `errno` to one of the following values:

[EIO]	<i>request</i> is an illegal number.
[EFAULT]	One or more of the arguments provided is outside the caller's memory space.
[ESRCH]	The thread designated by <i>data</i> does not correspond to a valid thread identifier within the <code>c_actor</code> designated by <i>cactorcap</i> .
[ESRCH]	<i>cactorcap</i> designated a <code>c_actor</code> that does not exist or is not debugged by the calling <code>c_actor</code> .

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`awaits(2K)`, `threadSuspend(2K)`, `threadResume(2K)`, `threadStat(2K)`, `_exit(2K)`

<b>NAME</b>	await, awaits – wait for <i>c_actor</i> to terminate or stop												
<b>SYNOPSIS</b>	<pre>#include &lt;am/await.h&gt; int awaits(const KnCap * <i>cactorcap</i>, int * <i>statusp</i>);  IS_EXIT_EVT(int <i>status</i>); GET_EXIT_EVT(int <i>status</i>);  IS_KILL_EVT(int <i>status</i>); GET_KILL_EVT(int <i>status</i>);  IS_STP_EVT(int <i>status</i>); GET_STP_EVT(int <i>status</i>);  int await(const KnCap * <i>cactorcap</i>);</pre>												
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT												
<b>DESCRIPTION</b>	<p>The <i>awaits()</i> call blocks the caller until the <i>c_actor</i> identified by <i>cactorcap</i> (see <i>afexec(2K)</i>) terminates or stops due to tracing.</p> <p>If <i>statusp</i> is not NULL, the status of the <i>c_actor</i> designated by <i>cactorcap</i> is stored in the location pointed to by <i>statusp</i> upon successful return from <i>awaits()</i>. It indicates why the <i>c_actor</i> terminated or stopped.</p> <p>The IS_EXIT_EVT, GET_EXIT_EVT, IS_KILL_EVT, GET_KILL_EVT, IS_STP_EVT and GET_STP_EVT macros take an argument of the int type as returned by <i>awaits()</i> in the location pointed to by <i>statusp</i>. If the <i>c_actor</i> terminated due to a call to <i>_exit(2K)</i>, IS_EXIT_EVT will be true (1) and GET_EXIT_EVT will be the low-order byte of the argument the <i>c_actor</i> passed to <i>_exit()</i>. If the <i>c_actor</i> terminated due to an asynchronous event (exception, <i>actorDelete(2K)</i>, and so forth) IS_KILL_EVT will be true (1). If the <i>c_actor</i> was being debugged and caught an exception, IS_STP_EVT will be true (1). In both of these cases GET_KILL_EVT or GET_STP_EVT will be the number of the event which caused the <i>c_actor</i> to terminate or stop.</p> <p>The following events are valid:</p> <table border="0"> <tr> <td>EVT_ILL</td> <td>(illegal instruction exception)</td> </tr> <tr> <td>EVT_BKPT</td> <td>(breakpoint exception)</td> </tr> <tr> <td>EVT_FPE</td> <td>(floating point exception)</td> </tr> <tr> <td>EVT_SEGV</td> <td>(memory protection exception)</td> </tr> <tr> <td>EVT_AKILL</td> <td>(<i>c_actor</i> killed by <i>akill(2K)</i>)</td> </tr> <tr> <td>EVT_DELETE</td> <td>(<i>c_actor</i> killed by <i>actorDelete(2K)</i>)</td> </tr> </table>	EVT_ILL	(illegal instruction exception)	EVT_BKPT	(breakpoint exception)	EVT_FPE	(floating point exception)	EVT_SEGV	(memory protection exception)	EVT_AKILL	( <i>c_actor</i> killed by <i>akill(2K)</i> )	EVT_DELETE	( <i>c_actor</i> killed by <i>actorDelete(2K)</i> )
EVT_ILL	(illegal instruction exception)												
EVT_BKPT	(breakpoint exception)												
EVT_FPE	(floating point exception)												
EVT_SEGV	(memory protection exception)												
EVT_AKILL	( <i>c_actor</i> killed by <i>akill(2K)</i> )												
EVT_DELETE	( <i>c_actor</i> killed by <i>actorDelete(2K)</i> )												

The `await()` call performs the same function as `awaits()`, and should be used when `statusp` is a `NULL` pointer.

**RETURN VALUES**

Both `awaits()` and `await()` return the *aid* of the `c_actor` which terminated or stopped due to tracing; otherwise they return `-1` and set `errno` to indicate one of the following error conditions:

- [EINVAL] *cactorcap* refers to the calling `c_actor`.
- [ESRCH] *cactorcap* doesn't designate a valid `c_actor`.
- [EFAULT] *cactorcap* or *statusp* points to an illegal address.
- [EINTR] The calling thread was aborted before the execution of the `c_actor` was terminated.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`actorDelete(2K)`, `afexec(2K)`, `akill(2K)`, `_exit(2K)`

**RESTRICTIONS**

It is not currently possible to distinguish between terminated `c_actors` and invalid actor capabilities.

<b>NAME</b>	await, awaits – wait for <i>c_actor</i> to terminate or stop												
<b>SYNOPSIS</b>	<pre>#include &lt;am/await.h&gt; int awaits(const KnCap * <i>cactorcap</i>, int * <i>statusp</i>);  IS_EXIT_EVT(int <i>status</i>); GET_EXIT_EVT(int <i>status</i>);  IS_KILL_EVT(int <i>status</i>); GET_KILL_EVT(int <i>status</i>);  IS_STP_EVT(int <i>status</i>); GET_STP_EVT(int <i>status</i>);  int await(const KnCap * <i>cactorcap</i>);</pre>												
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT												
<b>DESCRIPTION</b>	<p>The <code>awaits()</code> call blocks the caller until the <i>c_actor</i> identified by <i>cactorcap</i> (see <code>afexec(2K)</code>) terminates or stops due to tracing.</p> <p>If <i>statusp</i> is not NULL, the status of the <i>c_actor</i> designated by <i>cactorcap</i> is stored in the location pointed to by <i>statusp</i> upon successful return from <code>awaits()</code>. It indicates why the <i>c_actor</i> terminated or stopped.</p> <p>The <code>IS_EXIT_EVT</code>, <code>GET_EXIT_EVT</code>, <code>IS_KILL_EVT</code>, <code>GET_KILL_EVT</code>, <code>IS_STP_EVT</code> and <code>GET_STP_EVT</code> macros take an argument of the <code>int</code> type as returned by <code>awaits()</code> in the location pointed to by <i>statusp</i>. If the <i>c_actor</i> terminated due to a call to <code>_exit(2K)</code>, <code>IS_EXIT_EVT</code> will be true (1) and <code>GET_EXIT_EVT</code> will be the low-order byte of the argument the <i>c_actor</i> passed to <code>_exit()</code>. If the <i>c_actor</i> terminated due to an asynchronous event (exception, <code>actorDelete(2K)</code>, and so forth) <code>IS_KILL_EVT</code> will be true (1). If the <i>c_actor</i> was being debugged and caught an exception, <code>IS_STP_EVT</code> will be true (1). In both of these cases <code>GET_KILL_EVT</code> or <code>GET_STP_EVT</code> will be the number of the event which caused the <i>c_actor</i> to terminate or stop.</p> <p>The following events are valid:</p> <table border="0"> <tr> <td><code>EVT_ILL</code></td> <td>(illegal instruction exception)</td> </tr> <tr> <td><code>EVT_BKPT</code></td> <td>(breakpoint exception)</td> </tr> <tr> <td><code>EVT_FPE</code></td> <td>(floating point exception)</td> </tr> <tr> <td><code>EVT_SEGV</code></td> <td>(memory protection exception)</td> </tr> <tr> <td><code>EVT_AKILL</code></td> <td>(<i>c_actor</i> killed by <code>akill(2K)</code>)</td> </tr> <tr> <td><code>EVT_DELETE</code></td> <td>(<i>c_actor</i> killed by <code>actorDelete(2K)</code>)</td> </tr> </table>	<code>EVT_ILL</code>	(illegal instruction exception)	<code>EVT_BKPT</code>	(breakpoint exception)	<code>EVT_FPE</code>	(floating point exception)	<code>EVT_SEGV</code>	(memory protection exception)	<code>EVT_AKILL</code>	( <i>c_actor</i> killed by <code>akill(2K)</code> )	<code>EVT_DELETE</code>	( <i>c_actor</i> killed by <code>actorDelete(2K)</code> )
<code>EVT_ILL</code>	(illegal instruction exception)												
<code>EVT_BKPT</code>	(breakpoint exception)												
<code>EVT_FPE</code>	(floating point exception)												
<code>EVT_SEGV</code>	(memory protection exception)												
<code>EVT_AKILL</code>	( <i>c_actor</i> killed by <code>akill(2K)</code> )												
<code>EVT_DELETE</code>	( <i>c_actor</i> killed by <code>actorDelete(2K)</code> )												

The `await()` call performs the same function as `awaits()`, and should be used when `statusp` is a `NULL` pointer.

**RETURN VALUES**

Both `awaits()` and `await()` return the *aid* of the `c_actor` which terminated or stopped due to tracing; otherwise they return `-1` and set `errno` to indicate one of the following error conditions:

- [EINVAL] *cactorcap* refers to the calling `c_actor`.
- [ESRCH] *cactorcap* doesn't designate a valid `c_actor`.
- [EFAULT] *cactorcap* or *statusp* points to an illegal address.
- [EINTR] The calling thread was aborted before the execution of the `c_actor` was terminated.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`actorDelete(2K)`, `afexec(2K)`, `akill(2K)`, `_exit(2K)`

**RESTRICTIONS**

It is not currently possible to distinguish between terminated `c_actors` and invalid actor capabilities.

<b>NAME</b>	dladdr – translate address to symbolic information								
<b>SYNOPSIS</b>	<pre>#include &lt;cx/dlfcn.h&gt; int dladdr(void *address, Dl_info *dli);</pre>								
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT, DYNAMIC_LIB								
<b>DESCRIPTION</b>	<p>dladdr( ) is a member of the dynamic linking API, a family of routines that give the user direct access to dynamic linking functionality.</p> <hr/> <p>dladdr( ) is <i>only</i> available to programs and libraries compiled with imake rules of the form <code>DynamicTypeTarget( )</code>, where <i>Type</i> is one of <code>User</code>, <code>Sup</code>, <code>CXXUser</code>, <code>CXXSup</code> or <code>Library</code>.</p> <p>See <i>ChorusOS 4.0 Introduction</i> for more information on building executables with dynamic libraries.</p> <hr/> <p>dladdr( ) determines whether the specified address is located within one of the mapped objects that make up the address space of the current application. An address is considered to fall within a mapped object when it is between the base address and the <code>_end</code> address of that object. If a mapped object fits this criterion, the symbol table made available to the run-time linker is searched to locate the symbol nearest the address specified. The nearest symbol is the one that has a value less than or equal to the required address.</p>								
<b>PARAMETERS</b>	<p>The <code>Dl_info</code> structure must be preallocated by the user. The structure members are filled by <code>dladdr( )</code> based on the specified <i>address</i>. The <code>Dl_info</code> structure includes the following members:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>const char *dli_fname</code></td> <td>Contains a pointer to the filename of the containing object.</td> </tr> <tr> <td style="padding-right: 20px;"><code>void *dli_fbase</code></td> <td>Contains the base address of the containing object.</td> </tr> <tr> <td style="padding-right: 20px;"><code>const char *dli_sname</code></td> <td>Contains a pointer to the symbol name nearest to the address specified. This symbol either has the same address or is the nearest symbol with a lower address.</td> </tr> <tr> <td style="padding-right: 20px;"><code>void *dli_saddr</code></td> <td>Contains the actual address of the above symbol.</td> </tr> </table>	<code>const char *dli_fname</code>	Contains a pointer to the filename of the containing object.	<code>void *dli_fbase</code>	Contains the base address of the containing object.	<code>const char *dli_sname</code>	Contains a pointer to the symbol name nearest to the address specified. This symbol either has the same address or is the nearest symbol with a lower address.	<code>void *dli_saddr</code>	Contains the actual address of the above symbol.
<code>const char *dli_fname</code>	Contains a pointer to the filename of the containing object.								
<code>void *dli_fbase</code>	Contains the base address of the containing object.								
<code>const char *dli_sname</code>	Contains a pointer to the symbol name nearest to the address specified. This symbol either has the same address or is the nearest symbol with a lower address.								
<code>void *dli_saddr</code>	Contains the actual address of the above symbol.								
<b>RETURN VALUES</b>	<p>If the specified <i>address</i> cannot be matched to a mapped object, 0 is returned. Otherwise, a non-zero value is returned and the associated <code>Dl_info</code> elements are filled.</p>								
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:								

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`dlclose(2K)`, `dlerror(2K)`, `dlopen(2K)`, `dlsym(2K)`

*ChorusOS 4.0 Introduction***NOTES**

The *DL\_info* pointer elements point to addresses within the mapped objects. These addresses may become invalid if objects are removed before the pointer elements are used (see `dlclose(2K)`). If no symbol is found to describe the specified address, both `dli_sname` and `dli_saddr` members are set to 0.

<b>NAME</b>	dlclose – close a dynamic object				
<b>SYNOPSIS</b>	<pre>#include &lt;cx/dlfcn.h&gt; int dlclose(void *handle);</pre>				
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT, DYNAMIC_LIB				
<b>DESCRIPTION</b>	<p>dlclose( ) is a member of the dynamic linking API, a family of routines that give the user direct access to dynamic linking functionality.</p> <hr/> <p>dlclose( ) is <i>only</i> available to programs and libraries compiled with imake rules of the form <code>DynamicTypeTarget( )</code>, where <i>Type</i> is one of <code>User</code>, <code>Sup</code>, <code>CXXUser</code>, <code>CXXSup</code> or <code>Library</code>.</p> <p>See the <i>ChorusOS 4.0 Introduction</i> for more information on building executables with dynamic libraries.</p> <hr/> <p>dlclose( ) dissociates a dynamic object previously opened by <code>dlopen( )</code> from the current actor. Once an object has been closed using <code>dlclose( )</code>, its symbols are no longer available to <code>dlsym( )</code>. All objects loaded automatically as a result of invoking <code>dlopen( )</code> on the referenced object are also closed. The <code>handle</code> argument is the value returned by <code>dlopen( )</code>.</p>				
<b>RETURN VALUES</b>	<p>Upon successful completion, 0 is returned. If the object could not be closed, or if <i>handle</i> does not refer to an open object, <code>dlclose( )</code> returns a non-zero value. Detailed diagnostic information is available through <code>dlerror( )</code>.</p>				
<b>ATTRIBUTES</b>	<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;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>dladdr(2K)</code> , <code>dlerror(2K)</code> , <code>dlopen(2K)</code> , <code>dlsym(2K)</code> <i>ChorusOS 4.0 Introduction</i>				
<b>NOTES</b>	<p>Successful invocation of <code>dlclose( )</code> does not guarantee that the objects associated with <i>handle</i> will actually be removed from the address space of the actor. The same object may be opened multiple times. Objects loaded automatically when the program was started, or as a result of invoking <code>dlopen( )</code> on another object, may also be loaded by another invocation of <code>dlopen( )</code>. An object is not removed from the address space until all references to that object through explicit <code>dlopen( )</code> calls have been closed and all other objects implicitly referencing that object have also been closed.</p> <p>Once an object has been closed by <code>dlclose( )</code>, referencing symbols contained in that object can cause undefined behavior.</p>				

<b>NAME</b>	dlerror – get diagnostic information				
<b>SYNOPSIS</b>	<pre>#include &lt;cx/dlfcn.h&gt; char *dlerror(void);</pre>				
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT, DYNAMIC_LIB				
<b>DESCRIPTION</b>	<p>dlerror( ) is a member of the dynamic linking API, a family of routines that give the user direct access to dynamic linking functionality.</p> <hr/> <p>dlerror( ) is <i>only</i> available to programs and libraries compiled with imake rules of the form <code>DynamicTypeTarget( )</code>, where <i>Type</i> is one of User, Sup, CXXUser, CXXSup or Library.</p> <p>See <i>ChorusOS 4.0 Introduction</i> for more information on building executables with dynamic libraries.</p> <hr/> <p>dlerror( ) returns a null-terminated character string with no trailing newline that describes the last error that occurred during dynamic linking processing. If no dynamic linking errors have occurred since the last invocation of dlerror( ), dlerror( ) returns NULL. Therefore, invoking dlerror( ) twice in a row returns NULL.</p>				
<b>ATTRIBUTES</b>	<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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	dladdr(2K), dlclose(2K), dlopen(2K), dlsym(2K)				
<b>NOTES</b>	<p><i>Introducint ChorusOS 4.0</i></p> <p>Messages returned by dlerror( ) may reside in a static buffer that is overwritten on each call to dlerror( ). Application code should not write to this buffer. Programs that need to preserve an error message should make their own copies of the message.</p>				

<b>NAME</b>	dlopen – gain access to a dynamic object file
<b>SYNOPSIS</b>	<pre>#include &lt;cx/dlfcn.h&gt; void *dlopen(const char *pathname, int mode);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT, DYNAMIC_LIB
<b>DESCRIPTION</b>	<p>dlopen( ) is a member of the dynamic linking API, a family of routines that give the user direct access to dynamic linking functionality.</p> <hr/> <p>dlopen( ) is <i>only</i> available to programs and libraries compiled with imake rules of the form <code>DynamicTypeTarget( )</code>, where <i>Type</i> is one of <code>User</code>, <code>Sup</code>, <code>CXXUser</code>, <code>CXXSup</code> or <code>Library</code>.</p> <p>See the <i>ChorusOS 4.0 Introduction</i> for more information on building executables with dynamic libraries.</p> <hr/> <p>dlopen( ) makes a dynamic object file available to a running actor. dlopen( ) returns a <i>handle</i> that an actor may use when calling <code>dlsym( )</code> and <code>dlclose( )</code>.</p>
<b>PARAMETERS</b>	<p>The value of the <i>handle</i> should not be interpreted in any way by the actor. The <i>pathname</i> is the path name of the object to be opened. Any <i>pathname</i> containing / is interpreted as an absolute path or as a path relative to the current directory. Otherwise, the set of search paths currently in effect for the runtime linker is used to locate the specific file. See NOTES.</p> <p>Any dependencies recorded within the <i>pathname</i> object are also loaded as part of the dlopen( ) call. The dependencies are searched in the order they are loaded to locate any additional dependencies. The search continues until all dependencies of <i>pathname</i> are loaded.</p> <p>If the value of <i>pathname</i> is 0, then dlopen( ) provides a <i>handle</i> for the global symbol object. The global symbol object provides access to the symbols from an ordered set of objects consisting of the original program image file and any dependencies loaded at program startup, and any objects loaded using dlopen( ). Because the set of objects loaded using dlopen( ) can change during actor execution, the set identified by <i>handle</i> can also change dynamically.</p> <p>An object may be opened multiple times by dlopen( ). As long as the object has not been removed from memory using <code>dlclose( )</code>, the same <i>handle</i> is returned each time dlopen( ) is called. A counter is maintained internally to record the number of consecutive openings.</p> <p>The <i>mode</i> parameter describes how dlopen( ) operates on <i>pathname</i> with regard to relocation processing and scope of symbols from <i>pathname</i> and its dependencies. When an object is brought into the address space of an actor, it may contain references to symbols for which addresses are not known until the object is loaded. These types of references must be relocated before the</p>

symbols can be accessed. The *mode* parameter governs how these relocations are processed. The following *modes* are supported:

RTLD_NOW	Performs all necessary relocations when the object is first loaded.
RTLD_GLOBAL	Makes the object's global symbols available for relocation processing of any other object.
RTLD_LOCAL	Makes the object's global symbols available only for relocation processing of other objects in the same group.

The program image file and any objects loaded at program startup use the RTLD\_GLOBAL *mode*. By default, all objects accessed using `dlopen()` use RTLD\_LOCAL *mode*. Any local object may depend on more than one group. Any object using RTLD\_LOCAL *mode* referenced as a dependency of an object using RTLD\_GLOBAL *mode* is promoted to RTLD\_GLOBAL, and RTLD\_LOCAL is ignored. Any object loaded by `dlopen()` that requires relocation of global symbols can reference the symbols in any RTLD\_GLOBAL object, including, at least, the program image file, any objects loaded at program startup, the object itself and any dependencies that the object references. However, the *mode* parameter may also be bitwise-ORed with the following values that affect the scope of symbol availability:

RTLD_GROUP	Makes only symbols from the associated group available for relocation. The group is the object and all dependencies of the object. A group must be completely self-contained. All dependencies between members of the group must satisfy the relocation requirements of each object in the group.
RTLD_WORLD	Makes only symbols from RTLD_GLOBAL objects available for relocation.

The default *modes* for `dlopen()` are RTLD\_WORLD and RTLD\_GROUP. Both *modes* are bitwise-ORed together if the object is required by different dependencies specifying different *modes*.

---

RTLD\_LAZY (lazy-binding) and RTLD\_PARENT are not supported.

---

The following *modes* provide additional functionality beyond relocation processing:

RTLD_NODELETE	Does not delete the specified object from the address space as part of <code>dlclose()</code> . The opening counter is still decremented, however.
---------------	--

**RTLD\_NOLOAD** Does not load the specified object as part of the `dlopen( )` call, but a valid *handle* is returned if the object already exists in the actor address space. Additional *modes* can be specified and are bitwise-ORed with the current *mode* of the object and its dependencies. `RTLD_NOLOAD` provides a means of querying the presence, or promoting the *modes*, of an existing dependency. Note that `RTLD_NOLOAD` does not prevent the opening counter of the object from being incremented. Therefore, using it on an existing object requires an additional `dlclose( )` to effectively remove the object from memory.

**RETURN VALUES**

If the specified *pathname* cannot be found, cannot be opened for reading, is not a relocatable object, or if an error occurs while loading *pathname* or while relocating its symbolic references, `dlopen( )` returns `NULL`.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`dladdr(2K)`, `dlclose(2K)`, `dlerror(2K)`, `dlsym(2K)`

*ChorusOS 4.0 Introduction*

**NOTES**

If other objects are were link-edited with *pathname* when *pathname* was built — that is, if *pathname* has dependencies on other objects — those objects are automatically loaded by `dlopen( )`. Unless *pathname* contains `/`, the directory search path used to find both *pathname* and the objects it depends upon may be set using either the environment variable `LD_LIBRARY_PATH`, which is analyzed at actor startup, or the runpath setting within the object calling `dlopen( )`. Objects whose names resolve to the same absolute or relative *pathname* may be opened any number of times using `dlopen( )`. However, the object referenced is loaded only once into the address space of the current actor.

<b>NAME</b>	dlsym – get the address of a symbol in a dynamic object
<b>SYNOPSIS</b>	<pre>#include &lt;cx/dlfcn.h&gt; void *dlsym(void *handle, const char *name);</pre>
<b>FEATURES</b>	ACTOR_EXTENDED_MNGT, DYNAMIC_LIB
<b>DESCRIPTION</b>	<p>dlsym( ) is a member of the dynamic linking API, a family of routines that give the user direct access to dynamic linking functionality.</p> <hr/> <p>dlsym( ) is <i>only</i> available to programs and libraries compiled with imake rules of the form <code>DynamicTypeTarget( )</code>, where <i>Type</i> is one of <code>User</code>, <code>Sup</code>, <code>CXXUser</code>, <code>CXXSup</code> or <code>Library</code>.</p> <p>See the <i>ChorusOS 4.0 Introduction</i> for more information on building executables with dynamic libraries.</p> <hr/> <p>dlsym( ) allows an actor to obtain the address of a symbol defined within a dynamic object. The <i>handle</i> is either the value returned from a call to <code>dlopen( )</code> or one of the special flags <code>RTLD_NEXT</code> or <code>RTLD_DEFAULT</code>. The <i>name</i> is the symbol's name as a character string.</p> <p>If the <i>handle</i> used has been returned by <code>dlopen( )</code>, the corresponding dynamic object must not have been closed using <code>dlclose( )</code>. <code>dlsym( )</code> searches for the named symbol in all dynamic objects loaded automatically as a result of loading the object referenced by <i>handle</i>. For details, see <code>dlopen(2K)</code>.</p> <p>If the <i>handle</i> used is <code>RTLD_NEXT</code>, <code>dlsym( )</code> searches for the named symbol in the objects that were loaded following the object from which the <code>dlsym( )</code> call is being made.</p> <p>If the <i>handle</i> used is <code>RTLD_DEFAULT</code>, <code>dlsym( )</code> searches for the named symbol, starting with the first object loaded and proceeding through the list of loaded objects until a match is found. This search follows the default model employed to relocate all objects within the actor.</p> <p>If the <i>handle</i> used is either <code>RTLD_NEXT</code> or <code>RTLD_DEFAULT</code> and if the objects being searched have been loaded using <code>dlopen( )</code>, <code>dlsym( )</code> searches the object only if the caller is part of the same <code>dlopen( )</code> dependency hierarchy, or if the object was given global search access. See <code>dlopen(2K)</code> for a discussion of the <code>RTLD_GLOBAL</code> mode.</p>
<b>RETURN VALUES</b>	<p>If the specified <i>handle</i> does not refer to a valid object opened using <code>dlopen( )</code>, is not the special flag <code>RTLD_NEXT</code>, or if the named symbol cannot be found within any of the objects associated with <i>handle</i>, then <code>dlsym( )</code> returns <code>NULL</code>. More detailed diagnostic information is available through <code>dlerror( )</code>.</p>

**EXAMPLES**

The following example shows how to use `dlopen( )` and `dlsym( )` to access either function or data objects. For simplicity, error checking has been omitted.

```
void *handle;
int *iptr, (*fptr)(int);
/* open the object needed */
handle = dlopen("/usr/home/me/libfoo.so.1", RTLD_NOW);

/* find the address of function and data objects */
fptr = (int (*)(int))dlsym(handle, "my_function");
iptr = (int *)dlsym(handle, "my_object");

/* invoke function, passing value of integer as a parameter */
(*fptr)(*iptr);
```

The following example shows how to use `dlsym( )` to check that a particular function is defined and to call it only if it is defined.

```
int (*fptr)();
if ((fptr = (int (*)())dlsym(RTLD_DEFAULT,
    "my_function")) != NULL) {
    (*fptr)();
}
```

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`dladdr(2K)`, `dlclose(2K)`, `dLError(2K)`, `dlopen(2K)`  
*ChorusOS 4.0 Introduction*

<b>NAME</b>	ethIpcStackAttach – attach an IPC stack to an Ethernet device				
<b>SYNOPSIS</b>	<pre>#include &lt;ddi/net/netFrame.h&gt; #include &lt;iom/ipcStackAttach.h&gt; int ethIpcStackAttach(const *device);</pre>				
<b>FEATURES</b>	IOM_IPC, IPC_REMOTE				
<b>DESCRIPTION</b>	<p>If the IOM_IPC feature is set to true, the IOM actor includes an IPC stack. ethIpcStackAttach( ) attaches the IPC stack to an Ethernet <i>device</i> handled by the IOM.</p> <p>ethIpcStackAttach( ) should be executed only after the corresponding Ethernet <i>device</i> is up and running and has been assigned an IP address.</p> <p>ethIpcStackAttach( ) works by causing the IOM to scan its network its ifnet interfaces until it finds the interface attached to the corresponding Ethernet hardware <i>device</i>. Once that interface is found, it processes IPC traffic. If <i>device</i> is NULL, the first IOM Ethernet network interface is used.</p> <p>Once attached, the IPC stack cannot be detached. Furthermore, the IPC stack may only be attached to a single Ethernet hardware <i>device</i>.</p>				
<b>RETURN VALUES</b>	Upon successful completion, ethIpcStackAttach( ) returns 0; otherwise it returns -1 and sets errno to indicate the error.				
<b>ERRORS</b>	<p>[EACCES] ethIpcStackAttach( ) was called by a user actor. Only supervisor actors can invoke ethIpcStackAttach( ).</p> <p>[ENODEV] No network interface attached to <i>device</i> was found.</p> <p>[EINVAL] The IPC stack could not be correctly initialized.</p>				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	C_INIT(1M)				

<b>NAME</b>	ethOsiStackAttach – attach an OSI stack to an Ethernet device
<b>SYNOPSIS</b>	<pre>#include &lt;ddi/net/netFrame.h&gt; #include &lt;iom/osiStackAttach.h&gt; int ethOsiStackAttach(const *device, OsiToEthConf* in, EthToOsiConf* out);</pre>
<b>FEATURES</b>	IOM_OSI
<b>DESCRIPTION</b>	<p>ethOsiStackAttach( ) attaches an OSI stack that you provide to an Ethernet <i>device</i> handled by the IOM.</p> <p>ethOsiStackAttach( ) works by causing the IOM to scan its network its ifnet interfaces until it finds the interface attached to the corresponding Ethernet hardware <i>device</i>. Once that interface is found, it processes OSI traffic. If <i>device</i> is NULL, the first IOM Ethernet network interface is used.</p> <p>Once attached, the OSI stack cannot be detached. OSI stacks may be attached to multiple Ethernet hardware <i>devices</i> and each <i>device</i> may be attached to multiple OSI stacks.</p>
<b>EXTENDED DESCRIPTION</b>	<p>The <i>in</i> and <i>out</i> parameters are defined as follows:</p> <pre>typedef void      (*EthFrameSend)      (void* ethCookie, struct NetFrame* outFrame); typedef void      (*OsiFrameReceive)   (void* osiCookie, struct NetFrame* inFrame);  typedef struct {     void*          osiCookie;     OsiFrameReceive osiFrameReceive; } OsiToEthConf;  typedef struct {     void*          ethCookie;     char           etherAddr[6];     EthFrameSend  ethFrameSend; } EthToOsiConf;</pre> <p>The OSI stack provider must provide both (i) a cookie and (ii) an upcall for processing OSI network frames received in the <i>in</i> parameter.</p> <p>It must also return (i) an Ethernet cookie, (ii) the Ethernet address of the Ethernet <i>device</i>, and (iii) a function pointer for sending OSI frames in the <i>out</i> parameter.</p> <p>After ethOsiStackAttach( ) succeeds, the IOM will invoke the osiFrameReceive( ) function you provide each time it receives an OSI network frame. The first argument is the cookie given before in the <i>in</i> parameter passed to ethOsiStackAttach( ), the second parameter is the frame whose type is defined as follows (&lt;ddi/net/netFrame.h&gt;):</p> <pre>typedef struct NetBuf {     struct NetBuf* next;    /* next buffer in the list or NULL */</pre>

```

        uint32_f      bufSize; /* size of memory buffer */
        char*        bufAddr; /* starting address of memory buffer */
    } NetBuf;

typedef struct NetFrame* NetFramePtr;
typedef void (*NetFrameFree)(NetFramePtr);

typedef struct NetFrame {
    struct NetFrame* next; /* to build list of NetFrames */
    uint32_f      frameSize; /* total length of frame */
    NetBuf*       bufList; /* list of memory buffers holding the data */
    NetFrameFree  freeFrame; /* the free function of the frame */
} NetFrame;

```

The OSI stack is given ownership of the `NetFrame`. Once done with it, it must invoke the `freeFrame()` function to free it.

The `osiFrameReceive()` function is called at interrupt level. It is highly recommended that you optimize it in order to process the frame as fast as possible.

#### RETURN VALUES

Upon successful completion, `ethOsiStackAttach()` returns 0; otherwise it returns -1 and sets `errno` to indicate the error.

#### ERRORS

[EACCES] `ethOsiStackAttach()` was called by a user actor.  
 Only supervisor actors can invoke `ethOsiStackAttach()`.

[ENODEV] No network interface attached to *device* was found.

#### ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	eventInit, eventClear, eventPost, eventWait – initialize an event flag set; Clear events in an event flag set; Post events in an event flag set; Wait for events in an event flag set
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chEvent.h&gt; int eventInit(KnEventSet * eventSet);  int eventClear(KnEventSet * eventSet, unsigned int mask);  int eventPost(KnEventSet * eventSet, unsigned int mask);  int eventWait(KnEventSet * eventSet, unsigned int inMask, unsigned int option, unsigned int * outMask, KnTimeVal * waitLimit);</pre>
<b>FEATURES</b>	EVENT
<b>DESCRIPTION</b>	<p>Event flag sets are <i>KnEventSet</i> structures allocated in user memory. An event flag set is a set of bits in memory associated with a thread wait queue. Each bit is associated with one event. The set is implemented as an unsigned integer; the maximum number of flags in a set is <math>8 * \text{sizeof}(int)</math>. Within a set, each event flag is designated by an integer number (<math>[0 .. 8 * \text{sizeof}(int) - 1]</math>).</p> <p>When a flag is set, it is said to be POSTED, and the associated event is considered to have occurred. Otherwise, the flag is said to be UNPOSTED, and the associated event has not yet occurred. Both threads and interrupt handlers can use event flag sets for signaling purposes.</p> <p>The <i>eventInit</i> call initializes the event flag set whose address is <i>eventSet</i>. All events are initialized in the UNPOSTED state.</p> <p>The <i>eventPost</i> call posts (sets to POSTED) one or more event flags within <i>eventSet</i>. Any thread waiting for posted events will be awakened, and subsequently <i>eventWait</i> calls referring to the events posted will return immediately. The event flags remain in the posted state until explicitly cleared using <i>eventClear</i>. The <i>mask</i> field specifies the mask of posted events (event number <i>i</i> is posted if bit <i>i</i> of <i>mask</i> is set). You can call <i>eventPost</i> within an interrupt handler, or with preemption disabled.</p> <p>The <i>eventWait</i> call makes the current thread wait conditionally for one or more events on the event flag set <i>eventSet</i>.</p> <p>The <i>inMask</i> field specifies the mask of awaited events (event number { <i>i</i> } is awaited if the bit <i>i</i> of <i>inMask</i> is set). If <i>inMask</i> is zero (no awaited flag), <i>eventWait</i> returns immediately, regardless the value of <i>waitLimit</i>. This can be used to obtain the current state of the event flag set.</p> <p>The <i>option</i> field can be:</p>

**K\_EVENT\_OR** The thread waits for an OR condition to be satisfied by the events indicated in *inMask*; it is sufficient that one of the events be posted for the thread to be awakened.

**K\_EVENT\_AND** The thread waits for an AND condition to be satisfied by the events indicated in *inMask*; all the events in the mask must be posted for the thread to be awakened.

Posted events remain posted after *eventWait* returns, until cleared explicitly using *eventClear*. If a subsequent *eventWait* with the same *inMask* is launched before clearing, it will return immediately with the same status as the previous call.

The *outMask* field shows the state of the event flag set when the thread was awakened.

If the thread event-set state does not satisfy the conditions described in *inMask* and *option*, the thread is blocked according to the options described in *waitLimit* in *intro(K)*. The *waitLimit* pointer indicates a *KnTimeVal* structure which contains a timeout interval as described in *sysTime(K)*.

The *eventClear* call clears one or more event flags within the event flag set *eventSet*. The *mask* field shows the mask of cleared events (event number *i* is cleared if the bit *i* of *mask* is set).

#### RESTRICTIONS

A user application and a supervisor application may not share an event.

However, two user applications may share an event by mapping it in both user address spaces.

#### RETURN VALUE

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

#### ERRORS

[K\_EABORT] *eventWait* has been aborted.

[K\_EFAULT] *eventSet* or *outMask* points outside the current actor's address space.

[K\_EINVAL] The *eventSet* structure has not been correctly initialized, or the *waitLimit* is an invalid *KnTimeVal*.

[K\_ETIMEOUT] The time-out occurred.

#### ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	eventInit, eventClear, eventPost, eventWait – initialize an event flag set; Clear events in an event flag set; Post events in an event flag set; Wait for events in an event flag set
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chEvent.h&gt; int eventInit(KnEventSet * eventSet);  int eventClear(KnEventSet * eventSet, unsigned int mask);  int eventPost(KnEventSet * eventSet, unsigned int mask);  int eventWait(KnEventSet * eventSet, unsigned int inMask, unsigned int option, unsigned int * outMask, KnTimeVal * waitLimit);</pre>
<b>FEATURES</b>	EVENT
<b>DESCRIPTION</b>	<p>Event flag sets are <i>KnEventSet</i> structures allocated in user memory. An event flag set is a set of bits in memory associated with a thread wait queue. Each bit is associated with one event. The set is implemented as an unsigned integer; the maximum number of flags in a set is <math>8 * \text{sizeof}(int)</math>. Within a set, each event flag is designated by an integer number (<math>[0 .. 8 * \text{sizeof}(int) - 1]</math>).</p> <p>When a flag is set, it is said to be POSTED, and the associated event is considered to have occurred. Otherwise, the flag is said to be UNPOSTED, and the associated event has not yet occurred. Both threads and interrupt handlers can use event flag sets for signaling purposes.</p> <p>The <i>eventInit</i> call initializes the event flag set whose address is <i>eventSet</i>. All events are initialized in the UNPOSTED state.</p> <p>The <i>eventPost</i> call posts (sets to POSTED) one or more event flags within <i>eventSet</i>. Any thread waiting for posted events will be awakened, and subsequently <i>eventWait</i> calls referring to the events posted will return immediately. The event flags remain in the posted state until explicitly cleared using <i>eventClear</i>. The <i>mask</i> field specifies the mask of posted events (event number <i>i</i> is posted if bit <i>i</i> of <i>mask</i> is set). You can call <i>eventPost</i> within an interrupt handler, or with preemption disabled.</p> <p>The <i>eventWait</i> call makes the current thread wait conditionally for one or more events on the event flag set <i>eventSet</i>.</p> <p>The <i>inMask</i> field specifies the mask of awaited events (event number { <i>i</i> } is awaited if the bit <i>i</i> of <i>inMask</i> is set). If <i>inMask</i> is zero (no awaited flag), <i>eventWait</i> returns immediately, regardless the value of <i>waitLimit</i>. This can be used to obtain the current state of the event flag set.</p> <p>The <i>option</i> field can be:</p>

**K\_EVENT\_OR** The thread waits for an OR condition to be satisfied by the events indicated in *inMask*; it is sufficient that one of the events be posted for the thread to be awakened.

**K\_EVENT\_AND** The thread waits for an AND condition to be satisfied by the events indicated in *inMask*; all the events in the mask must be posted for the thread to be awakened.

Posted events remain posted after *eventWait* returns, until cleared explicitly using *eventClear*. If a subsequent *eventWait* with the same *inMask* is launched before clearing, it will return immediately with the same status as the previous call.

The *outMask* field shows the state of the event flag set when the thread was awakened.

If the thread event-set state does not satisfy the conditions described in *inMask* and *option*, the thread is blocked according to the options described in *waitLimit* in *intro(K)*. The *waitLimit* pointer indicates a *KnTimeVal* structure which contains a timeout interval as described in *sysTime(K)*.

The *eventClear* call clears one or more event flags within the event flag set *eventSet*. The *mask* field shows the mask of cleared events (event number *i* is cleared if the bit *i* of *mask* is set).

#### RESTRICTIONS

A user application and a supervisor application may not share an event.

However, two user applications may share an event by mapping it in both user address spaces.

#### RETURN VALUE

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

#### ERRORS

[K\_EABORT] *eventWait* has been aborted.

[K\_EFAULT] *eventSet* or *outMask* points outside the current actor's address space.

[K\_EINVAL] The *eventSet* structure has not been correctly initialized, or the *waitLimit* is an invalid *KnTimeVal*.

[K\_ETIMEOUT] The time-out occurred.

#### ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	eventInit, eventClear, eventPost, eventWait – initialize an event flag set; Clear events in an event flag set; Post events in an event flag set; Wait for events in an event flag set
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chEvent.h&gt; int eventInit(KnEventSet * eventSet);  int eventClear(KnEventSet * eventSet, unsigned int mask);  int eventPost(KnEventSet * eventSet, unsigned int mask);  int eventWait(KnEventSet * eventSet, unsigned int inMask, unsigned int option, unsigned int * outMask, KnTimeVal * waitLimit);</pre>
<b>FEATURES</b>	EVENT
<b>DESCRIPTION</b>	<p>Event flag sets are <i>KnEventSet</i> structures allocated in user memory. An event flag set is a set of bits in memory associated with a thread wait queue. Each bit is associated with one event. The set is implemented as an unsigned integer; the maximum number of flags in a set is <math>8 * \text{sizeof}(int)</math>. Within a set, each event flag is designated by an integer number (<math>[0 .. 8 * \text{sizeof}(int) - 1]</math>).</p> <p>When a flag is set, it is said to be POSTED, and the associated event is considered to have occurred. Otherwise, the flag is said to be UNPOSTED, and the associated event has not yet occurred. Both threads and interrupt handlers can use event flag sets for signaling purposes.</p> <p>The <i>eventInit</i> call initializes the event flag set whose address is <i>eventSet</i>. All events are initialized in the UNPOSTED state.</p> <p>The <i>eventPost</i> call posts (sets to POSTED) one or more event flags within <i>eventSet</i>. Any thread waiting for posted events will be awakened, and subsequently <i>eventWait</i> calls referring to the events posted will return immediately. The event flags remain in the posted state until explicitly cleared using <i>eventClear</i>. The <i>mask</i> field specifies the mask of posted events (event number <i>i</i> is posted if bit <i>i</i> of <i>mask</i> is set). You can call <i>eventPost</i> within an interrupt handler, or with preemption disabled.</p> <p>The <i>eventWait</i> call makes the current thread wait conditionally for one or more events on the event flag set <i>eventSet</i>.</p> <p>The <i>inMask</i> field specifies the mask of awaited events (event number { <i>i</i> } is awaited if the bit <i>i</i> of <i>inMask</i> is set). If <i>inMask</i> is zero (no awaited flag), <i>eventWait</i> returns immediately, regardless the value of <i>waitLimit</i>. This can be used to obtain the current state of the event flag set.</p> <p>The <i>option</i> field can be:</p>

**K\_EVENT\_OR** The thread waits for an OR condition to be satisfied by the events indicated in *inMask*; it is sufficient that one of the events be posted for the thread to be awakened.

**K\_EVENT\_AND** The thread waits for an AND condition to be satisfied by the events indicated in *inMask*; all the events in the mask must be posted for the thread to be awakened.

Posted events remain posted after *eventWait* returns, until cleared explicitly using *eventClear*. If a subsequent *eventWait* with the same *inMask* is launched before clearing, it will return immediately with the same status as the previous call.

The *outMask* field shows the state of the event flag set when the thread was awakened.

If the thread event-set state does not satisfy the conditions described in *inMask* and *option*, the thread is blocked according to the options described in *waitLimit* in *intro(K)*. The *waitLimit* pointer indicates a *KnTimeVal* structure which contains a timeout interval as described in *sysTime(K)*.

The *eventClear* call clears one or more event flags within the event flag set *eventSet*. The *mask* field shows the mask of cleared events (event number *i* is cleared if the bit *i* of *mask* is set).

#### RESTRICTIONS

A user application and a supervisor application may not share an event.

However, two user applications may share an event by mapping it in both user address spaces.

#### RETURN VALUE

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

#### ERRORS

[K\_EABORT] *eventWait* has been aborted.

[K\_EFAULT] *eventSet* or *outMask* points outside the current actor's address space.

[K\_EINVAL] The *eventSet* structure has not been correctly initialized, or the *waitLimit* is an invalid *KnTimeVal*.

[K\_ETIMEOUT] The time- out occurred.

#### ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	eventInit, eventClear, eventPost, eventWait – initialize an event flag set; Clear events in an event flag set; Post events in an event flag set; Wait for events in an event flag set
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chEvent.h&gt; int eventInit(KnEventSet * eventSet);  int eventClear(KnEventSet * eventSet, unsigned int mask);  int eventPost(KnEventSet * eventSet, unsigned int mask);  int eventWait(KnEventSet * eventSet, unsigned int inMask, unsigned int option, unsigned int * outMask, KnTimeVal * waitLimit);</pre>
<b>FEATURES</b>	EVENT
<b>DESCRIPTION</b>	<p>Event flag sets are <i>KnEventSet</i> structures allocated in user memory. An event flag set is a set of bits in memory associated with a thread wait queue. Each bit is associated with one event. The set is implemented as an unsigned integer; the maximum number of flags in a set is <math>8 * \text{sizeof}(int)</math>. Within a set, each event flag is designated by an integer number (<math>[0 .. 8 * \text{sizeof}(int) - 1]</math>).</p> <p>When a flag is set, it is said to be POSTED, and the associated event is considered to have occurred. Otherwise, the flag is said to be UNPOSTED, and the associated event has not yet occurred. Both threads and interrupt handlers can use event flag sets for signaling purposes.</p> <p>The <i>eventInit</i> call initializes the event flag set whose address is <i>eventSet</i>. All events are initialized in the UNPOSTED state.</p> <p>The <i>eventPost</i> call posts (sets to POSTED) one or more event flags within <i>eventSet</i>. Any thread waiting for posted events will be awakened, and subsequently <i>eventWait</i> calls referring to the events posted will return immediately. The event flags remain in the posted state until explicitly cleared using <i>eventClear</i>. The <i>mask</i> field specifies the mask of posted events (event number <i>i</i> is posted if bit <i>i</i> of <i>mask</i> is set). You can call <i>eventPost</i> within an interrupt handler, or with preemption disabled.</p> <p>The <i>eventWait</i> call makes the current thread wait conditionally for one or more events on the event flag set <i>eventSet</i>.</p> <p>The <i>inMask</i> field specifies the mask of awaited events (event number { <i>i</i> } is awaited if the bit <i>i</i> of <i>inMask</i> is set). If <i>inMask</i> is zero (no awaited flag), <i>eventWait</i> returns immediately, regardless the value of <i>waitLimit</i>. This can be used to obtain the current state of the event flag set.</p> <p>The <i>option</i> field can be:</p>

**K\_EVENT\_OR** The thread waits for an OR condition to be satisfied by the events indicated in *inMask*; it is sufficient that one of the events be posted for the thread to be awakened.

**K\_EVENT\_AND** The thread waits for an AND condition to be satisfied by the events indicated in *inMask*; all the events in the mask must be posted for the thread to be awakened.

Posted events remain posted after *eventWait* returns, until cleared explicitly using *eventClear*. If a subsequent *eventWait* with the same *inMask* is launched before clearing, it will return immediately with the same status as the previous call.

The *outMask* field shows the state of the event flag set when the thread was awakened.

If the thread event-set state does not satisfy the conditions described in *inMask* and *option*, the thread is blocked according to the options described in *waitLimit* in *intro(K)*. The *waitLimit* pointer indicates a *KnTimeVal* structure which contains a timeout interval as described in *sysTime(K)*.

The *eventClear* call clears one or more event flags within the event flag set *eventSet*. The *mask* field shows the mask of cleared events (event number *i* is cleared if the bit *i* of *mask* is set).

#### RESTRICTIONS

A user application and a supervisor application may not share an event.

However, two user applications may share an event by mapping it in both user address spaces.

#### RETURN VALUE

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

#### ERRORS

[K\_EABORT] *eventWait* has been aborted.

[K\_EFAULT] *eventSet* or *outMask* points outside the current actor's address space.

[K\_EINVAL] The *eventSet* structure has not been correctly initialized, or the *waitLimit* is an invalid *KnTimeVal*.

[K\_ETIMEOUT] The time- out occurred.

#### ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**NAME** | \_exit – terminate a c\_actor

**SYNOPSIS** | #include <unistd.h>  
void \_exit(int status);

**FEATURES** | ACTOR\_EXTENDED\_MNGT

**DESCRIPTION** | The *\_exit* call terminates the calling *c\_actor* with the following consequences:

- All open file descriptors in the calling *c\_actor* are closed.
- If other *c\_actors* are executing an *await(2K)* or an *awaits(2K)* system call, they are notified of the calling *c\_actor*'s termination.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | *exit(3STDC)*, *awaits(2K)*

<b>NAME</b>	grpAllocate – allocate a port group capability
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int grpAllocate(int options, KnCap *groupcap, int stamp);</pre>
<b>FEATURES</b>	IPC
<b>DESCRIPTION</b>	<p><i>grpAllocate</i> allocates a port group capability.</p> <p><i>groupcap</i> is the returned capability for the port group, and is a pointer to a <i>KnCap</i> structure the members of which are the following:</p> <pre>KnUniqueId ui ; /* entity name */ KnKey      key ; /* modification key */</pre> <p><i>ui</i> is the unique identifier of the port group, used to send messages to the group (see <i>ipcSend(2K)</i> and <i>ipcTarget(2K)</i>). <i>key</i> is the key needed for modification of the group (i.e. inserting or removing ports by <i>grpPortInsert(2K)</i> and <i>grpPortRemove(2K)</i>).</p> <p><i>options</i> may have one of the following values, which describe the way in which the port group capability is to be allocated:</p> <p><b>K_DYNAMIC</b>      The call will return a new unique capability for a port group. <i>stamp</i> is ignored.</p> <p><b>K_STATUSER</b>     The group is a static user port group. The call will not return a new group capability, but a capability of one of the pre-allocated user port group capabilities. The right group is identified by <i>stamp</i> which may vary between 0 and 0xffffffff.</p> <p><b>K_STATSYS</b>      The group is a static system port group. The call will not return a new group name but rather the name of one of the pre-allocated system port group capabilities. The right group is identified by <i>stamp</i> which may vary between 0 and 0xffffffff. This option is only allowed for threads of <i>SYSTEM</i> actors or <i>SUPERVISOR</i> threads.</p> <p>Static and dynamic groups only differ by the way their name is allocated. They are used (by <i>ipcSend(2K)</i> or <i>ipcCall(2K)</i>) and managed (by <i>grpPortInsert(2K)</i> or <i>grpPortRemove(2K)</i>) in exactly the same way. Static groups are used in order to ease applications binding. Their main property resides in the fact that <i>grpAllocate</i> will always return the same capability for a given value of <i>stamp</i>. The different actors of an application independently invoke <i>grpAllocate</i>. After this operation, some of them insert ports in the group. The others simply send messages to the group. Dynamic groups are used when applications need new groups with</p>

unique names. Only one actor of the application invokes *grpAllocate*. It needs to transmit the group UI to other application actors to allow them to use the group.

**RETURN VALUE**

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EPRIV] The K\_STATSYS option is not allowed to the current thread.

[K\_EFAULT] Some of the provided data are outside the current actor's address space.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*grpPortInsert(2K)*, *ipcSend(2K)*, *ipcTarget(2K)*, *grpPortRemove(2K)*

<b>NAME</b>	grpPortInsert, grpPortRemove – insert a port into a port group; remove a port from a port group				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int grpPortInsert(KnCap * groupcap, KnUniqueId * portui); int grpPortRemove(KnCap * groupcap, KnUniqueId * portui);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p><i>grpPortInsert</i> inserts the port, the name of which is given by <i>portui</i> into the port group, the capability of which is pointed to by <i>groupcap</i>.</p> <p><i>grpPortRemove</i> removes the port, the name of which is given by <i>portui</i> from the port group, the capability of which is pointed to by <i>groupcap</i>.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned. <i>grpPortInsert</i> returns 0 if the port was already in the port group.				
<b>ERRORS</b>	<p>[K_EINVAL] <i>groupcap</i> is an invalid port group capability, or there was an attempt to remove a port from a port group to which it did not belong.</p> <p>[K_EUNKNOWN] <i>portui</i> is an invalid port name.</p> <p>[K_EFAULT] Some of the provided data are outside the current actor's address space.</p> <p>[K_ENOMEM] The system is out of resources.</p> <p>[K_EPRIV] The current thread is neither a supervisor thread nor a thread of a system actor, and attempts to insert a port or remove a port in a system static group.</p>				
<b>RESTRICTIONS</b>	The port and the current actor must be located on the same site.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>grpAllocate(2K)</code>				

<b>NAME</b>	grpPortInsert, grpPortRemove – insert a port into a port group; remove a port from a port group				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int grpPortInsert(KnCap * groupcap, KnUniqueId * portui); int grpPortRemove(KnCap * groupcap, KnUniqueId * portui);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p><i>grpPortInsert</i> inserts the port, the name of which is given by <i>portui</i> into the port group, the capability of which is pointed to by <i>groupcap</i>.</p> <p><i>grpPortRemove</i> removes the port, the name of which is given by <i>portui</i> from the port group, the capability of which is pointed to by <i>groupcap</i>.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned. <i>grpPortInsert</i> returns 0 if the port was already in the port group.				
<b>ERRORS</b>	<p>[K_EINVAL] <i>groupcap</i> is an invalid port group capability, or there was an attempt to remove a port from a port group to which it did not belong.</p> <p>[K_EUNKNOWN] <i>portui</i> is an invalid port name.</p> <p>[K_EFAULT] Some of the provided data are outside the current actor's address space.</p> <p>[K_ENOMEM] The system is out of resources.</p> <p>[K_EPRIV] The current thread is neither a supervisor thread nor a thread of a system actor, and attempts to insert a port or remove a port in a system static group.</p>				
<b>RESTRICTIONS</b>	The port and the current actor must be located on the same site.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>grpAllocate(2K)</code>				

<b>NAME</b>	ipcCall – send an RPC request and wait for the reply
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int ipcCall(KnMsgDesc *reqmsg, int reqsrc, KnIpcDest *reqdest, KnMsgDesc *repmsg, int delay);</pre>
<b>FEATURES</b>	IPC
<b>DESCRIPTION</b>	<p>The <i>ipcCall</i> function sends a message in transactional (<i>RPC</i>) mode, and blocks the caller until a reply to the request is received or the optional <i>delay</i> expires.</p> <p>The <i>delay</i> field is the maximum waiting time, expressed in milliseconds. If <i>delay</i> has a negative value, the blocking time is unlimited.</p> <p>The <i>reqdest</i> field defines the destination of the message, and is a pointer to a <i>KnIpcDest</i> structure. Its meaning is described in <i>ipcSend(2K)</i>. Use of the <i>K_BROADMODE</i> addressing mode on a port group is forbidden.</p> <p>The <i>reqsrc</i> field is a local identifier for the source port of the message. If <i>reqsrc</i> is <i>K_DEFAULTPORT</i>, the default port of the current actor is used.</p> <p>The <i>reqmsg</i> field points to a descriptor for the request message to be sent, on condition that <i>reqmsg</i> points to a <i>KnMsgDesc</i> structure, as described in <i>ipcSend(2K)</i>.</p> <p>The <i>repmsg</i> field is a descriptor for the reply message to be received, as described in <i>ipcReceive(2K)</i>. As in <i>ipcReceive(2K)</i>, the <i>K_ABORTABLE</i> flag may be present in the <i>flags</i> field of this structure. In this case, the <i>ipcCall</i> is <i>ABORTABLE</i>. Otherwise, it is <i>NONABORTABLE</i>. These two different states determine the behavior of the thread when a <i>threadAbort(2K)</i> is applied to it. If <i>threadAbort(2K)</i> is applied to a <i>NONABORTABLE ipcCall</i>, the thread replies and stays in the <i>ABORTED</i> state (see <i>threadAbort(2K)</i>). If <i>threadAbort(2K)</i> is applied to an <i>ABORTABLE ipcCall</i>, the kernel attempts to pass the abortion on to the thread which is processing the request messages:</p> <ul style="list-style-type: none"> <li>■ If the request message is queued behind the destination port, the message is deleted and the <i>ipcCall</i> returns <i>K_EABORT</i>.</li> <li>■ If a thread (called the <i>server</i>) currently is processing the request message, the <i>server</i> is aborted.</li> </ul> <p>This propagation mechanism is recursive: in the case of nested <i>ipcCall (2K)</i>, (the <i>server</i> is itself blocked in an <i>ABORTABLE ipcCall</i>), the abortion of the <i>server</i> is also passed on to its <i>server</i>.</p> <p>Unless explicitly shown by the return of the <i>K_EABORT</i> error code, a <i>client</i> thread aborted during its <i>ipcCall</i> always enters the <i>ABORTED</i> state upon return (the abortion still has to be processed; it appears as if the abortion occurred after the <i>ipcCall</i> returned).</p>

<b>RETURN VALUE</b>	Upon return, the reply received does not become the current message for the current thread. As a consequence, <i>ipcReply(2K)</i> , <i>ipcSave(2K)</i> , <i>ipcGetData(2K)</i> and <i>ipcSysInfo(2K)</i> may never be applied to the reply of an <i>ipcCall</i> .																		
<b>ERRORS</b>	<p>Upon successful completion, the size of the message body received is returned. Otherwise, a negative error code is returned. The error codes are those listed below, except when the destination port is connected to a message handler. In this case the error code is the return value of the handler function.</p> <table border="0"> <tr> <td style="vertical-align: top;">[K_EABORT]</td> <td>The thread was aborted while waiting; the request message was never delivered to its destination.</td> </tr> <tr> <td style="vertical-align: top;">[K_EBADMODE]</td> <td>Attempt to broadcast an <i>RPC</i> request.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAIL]</td> <td>The <i>ipcCall</i> transaction has failed.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the data provided are outside the current actor's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFULL]</td> <td>The destination port is local and its queue is full, or the local remote-communication subsystem is saturated.</td> </tr> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td><i>reqsrc</i> is not a valid local port identifier.</td> </tr> <tr> <td style="vertical-align: top;">[K_ETIMEOUT]</td> <td>The timeout has occurred.</td> </tr> <tr> <td style="vertical-align: top;">[K_ETOOMUCH]</td> <td>The request message body is too big.</td> </tr> <tr> <td style="vertical-align: top;">[K_EUNKNOWN]</td> <td>Unreachable destination. When the destination port is connected to a message handler, the error code is the return value of the handler function.</td> </tr> </table>	[K_EABORT]	The thread was aborted while waiting; the request message was never delivered to its destination.	[K_EBADMODE]	Attempt to broadcast an <i>RPC</i> request.	[K_EFAIL]	The <i>ipcCall</i> transaction has failed.	[K_EFAULT]	Some of the data provided are outside the current actor's address space.	[K_EFULL]	The destination port is local and its queue is full, or the local remote-communication subsystem is saturated.	[K_EINVAL]	<i>reqsrc</i> is not a valid local port identifier.	[K_ETIMEOUT]	The timeout has occurred.	[K_ETOOMUCH]	The request message body is too big.	[K_EUNKNOWN]	Unreachable destination. When the destination port is connected to a message handler, the error code is the return value of the handler function.
[K_EABORT]	The thread was aborted while waiting; the request message was never delivered to its destination.																		
[K_EBADMODE]	Attempt to broadcast an <i>RPC</i> request.																		
[K_EFAIL]	The <i>ipcCall</i> transaction has failed.																		
[K_EFAULT]	Some of the data provided are outside the current actor's address space.																		
[K_EFULL]	The destination port is local and its queue is full, or the local remote-communication subsystem is saturated.																		
[K_EINVAL]	<i>reqsrc</i> is not a valid local port identifier.																		
[K_ETIMEOUT]	The timeout has occurred.																		
[K_ETOOMUCH]	The request message body is too big.																		
[K_EUNKNOWN]	Unreachable destination. When the destination port is connected to a message handler, the error code is the return value of the handler function.																		
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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																		
<b>SEE ALSO</b>	<i>ipcSend(2K)</i> , <i>ipcReply(2K)</i> , <i>ipcReceive(2K)</i> , <i>ipcTarget(2K)</i> , <i>threadAbort(2K)</i> , <i>svMsgHandler(2K)</i>																		

<b>NAME</b>	ipcGetData – get the current message body				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int ipcGetData(KnMsgDesc *msg);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p>The <i>ipcGetData</i> call delivers the body size of the current thread's current message to the thread's address space.</p> <p>The message body may only be obtained if the message was received with a NULL <i>bodyAddr</i> argument (see <i>ipcReceive(2K)</i>). The message body may only be obtained once; system buffers are freed after the operation.</p> <p>The <i>msg</i> field points to a <i>KnMsgDesc</i> structure, as described in <i>ipcReceive(2K)</i>.</p> <p>This function returns the real body size delivered.</p>				
<b>RETURN VALUE</b>	Upon successful completion, the body size delivered is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EINVAL]                    There is no current message (no message has been received by the current thread since the last <i>ipcReply(2K)</i> or <i>ipcSave(2K)</i>).</p> <p>[K_EFAULT]                    Some of the data provided are outside the current actor's address space.</p>				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>ipcReceive(2K)</i>				

<b>NAME</b>	ipcReceive – receive a message
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int ipcReceive(KnMsgDesc *msg, int *portli, int delay);</pre>
<b>FEATURES</b>	IPC
<b>DESCRIPTION</b>	<p><i>ipcReceive</i> blocks the caller until a message (sent by <i>ipcSend(2K)</i>, or <i>ipcReply(2K)</i>, or <i>ipcCall(2K)</i>) is received on the port(s) specified by <i>portli</i> or the optional <i>delay</i> expires.</p> <p><i>delay</i> is the blocking waiting time, expressed in milliseconds. A delay of 0 means a non blocking attempt. If <i>delay</i> has a negative value, the blocking time is unbound.</p> <p>At call time, <i>portli</i> is a pointer to the local identifier of the port on which a message is expected. If <i>portli</i> points to the value <code>K_DEFAULTPORT</code>, the default port of the current actor is used.</p> <p>If <i>portli</i> points to the value <code>K_ANYENABLED</code>, a message is expected on any of the enabled ports owned by the current actor (see <i>portEnable(2K)</i>). In that case, the value pointed by <i>portli</i> is set, at return-time, to the local identifier of the port on which a message has been received. If more than one enabled port are holding messages at the time of the call, a message is received on the enabled port with the highest priority.</p> <p><i>msg</i> points to a message descriptor for the expected message. <i>msg</i> is a pointer to a <i>KnMsgDesc</i> structure whose members are the following:</p> <pre>unsigned int  flags ;      /* message structure definition */ unsigned int  bodySize ;  /* body size */ VmAddr       bodyAddr ;  /* body address */ VmAddr       annexAddr ; /* fixed size annex address */ KnEvtNum     seqNum ;    /* sequence number */</pre> <p>The <i>annexAddr</i> member of the message descriptor gives the starting address at which the system must copy the message annex in the receiver's address space. If this member is set to <code>NULL</code>, the annex is not delivered. Otherwise, <code>K_CMSGANNEXSIZE</code> bytes are copied by the system in the area starting at <i>annexAddr</i> in the receiver's address space. If the message was not sent with an annex (see <i>ipcSend(2K)</i>), this area is not affected.</p> <p>The <i>bodySize</i> and <i>bodyAddr</i> members of the message descriptor respectively give the size of the expected message body and the starting address at which the system must deliver it in the receiver's address space. If the expected size is smaller than the real message size, only the expected amount is received and the remainder is discarded. If <i>bodyAddr</i> is aligned on a page boundary, and if <i>bodySize</i> is a multiple of the page size (see <i>vmPageSize(2K)</i>), the body will be delivered without copy.</p>

If *bodyAddr* is NULL, and if the message contains a body, the message body is not delivered to the receiver address space but rather is kept in system buffers. If the body was not delivered at receive time (NULL *bodyAddr*), it can be delivered by an invocation of *ipcGetData(2K)*. This allows the receiver to fix the message body location in its address space only when the message body size is known.

The *seqNum* member of the message descriptor is an output value indicating the count of messages that has been received on the port, before and including the delivered message. This member is only updated when the *ipcReceive* return value is positive or null.

This call invalidates the current thread's previous current message. If the *ipcReceive* is successful, the received message becomes the current thread's current message, on which *ipcReply(2K)*, *ipcSysInfo(2K)*, *ipcGetData(2K)* or *ipcSave(2K)* may be applied.

The *flags* field of the message descriptor is a logical combination of the following options:

- K\_ABORTABLE If this flag is set, *ipcReceive* is *ABORTABLE*. Otherwise, it is *NONABORTABLE*. These two different states determine the behavior of the thread when a *threadAbort(2K)* is applied to it. See *threadAbort(2K)* for a description of this behavior in the two different cases.
  
- K\_USERBODY If the caller thread is a *SUPERVISOR* thread, this flag indicates that the message body destination address is part of the current user address space. If this flag is not set while the caller is a *SUPERVISOR* thread, the message body destination is assumed to be part of the kernel address space. This flag is only to be used by trap handling routines, when a received message body has to be directly received in the user address space, without copying it into the kernel address space. If the caller thread is not a *SUPERVISOR* thread, this flag is ignored (the message body destination is always assumed to be part of the user address space).
  
- K\_USERANNEX This flag has exactly the same meaning as the K\_USERBODY, but concerns the message annex instead of the message body.

If *flag* is equal to 0, none of the previous options is selected: the blocking is *NONABORTABLE*, and message and annex are assumed to be part of the user (kernel) address space when the caller is a *USER (SUPERVISOR)* thread.

#### RETURN VALUE

Upon successful completion:

- If *bodyAddr* was not NULL, the size in bytes of the delivered message body. If the system encountered a memory access fault while copying the message body to the receiver address space, the returned value is set to the successfully delivered body size.
- If *bodyAddr* was NULL, the size in bytes of the message body.

Otherwise, a negative error code is returned.

**ERRORS**

- [K\_ENOPORT] No port corresponds to *portli*.
- [K\_EINVAL] A handler is attached to the port.
- [K\_EFAULT] Some of the provided data are outside the current actor's address space.
- [K\_ETIMEOUT] The time out occurred.
- [K\_EABORT] The thread has been aborted while waiting.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`ipcSend(2K)`, `ipcReply(2K)`, `ipcCall(2K)`, `ipcSave(2K)`, `ipcSysInfo(2K)`, `ipcGetData(2K)`, `threadAbort(2K)`, `svMsgHandler(2K)`

<b>NAME</b>	ipcReply – reply to the current message				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int ipcReply(KnMsgDesc *msg);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p>The <i>ipcReply</i> call replies to the sender of the current thread's current message.</p> <p>The <i>msg</i> field points to a descriptor for the destination of the message (a <i>KnMsgDesc</i> structure, as described in <i>ipcSend(2K)</i>).</p> <p>The message is sent asynchronously; the sender is only blocked during the time needed by the system to process the request.</p> <p>After the reply, if the call is successful, the current thread will not have a current message.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EFULL]                   The destination port is local and its queue is full, or the local remote-communication subsystem is saturated.</p> <p>[K_EINVAL]                 There is no current message (no message has been received by the current thread since the last <i>ipcReply(2K)</i>, or <i>ipcSave(2K)</i>).</p> <p>[K_EFAULT]                 Some of the data provided are outside the current actor's address space.</p> <p>[K_ETOOMUCH]              The message body is too big.</p> <p>[K_EUNKNOWN]              Unreachable destination.</p>				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>ipcSend(2K)</i> , <i>ipcReceive(2K)</i> , <i>ipcCall(2K)</i> , <i>ipcGetData(2K)</i>				

<b>NAME</b>	ipcSave, ipcRestore – Save the current message; Restore a saved message as the current message				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int ipcSave(int portli);  int ipcRestore(int msgid);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p>The <i>ipcSave</i> call saves the current thread's current message to a system buffer, and returns an identifier for the saved message. The <i>portli</i> argument reserved for future use. The caller should use the K_NONEPORT value for this parameter.</p> <p>Messages used to be saved when a thread receives several messages and has to reply in a different order (<i>ipcReply</i> (2K) is always applied to the current message; receiving a message supplants the current message). After <i>ipcSave</i>, there is no current message, and <i>ipcReply</i> (2K), <i>ipcGetData</i> (2K) or <i>ipcSysInfo</i> (2K) cannot be used unless a message is received (<i>ipcReceive</i> (2K)) or restored (<i>ipcRestore</i> (2K)).</p> <p>The <i>ipcRestore</i> call restores a saved message as the current message. If the call is successful, the previous current message is replaced by the current message .</p> <p>The number of messages that an actor may save is limited to K_CMSGSAVEDMAX.</p>				
<b>RETURN VALUE</b>	Upon successful completion, <i>ipcRestore</i> returns 0 and <i>ipcSave</i> returns a positive identifier for the saved message. Otherwise, a negative error code is returned. In case of failure, the current message is not erased.				
<b>ERRORS</b>	<p>[K_ENOMEM]                    <i>ipcSave</i> : K_CMSGSAVEDMAX messages have already been saved by the calling actor.</p> <p>[K_EINVAL]                    <i>ipcSave</i> : no current message; <i>ipcRestore</i> : invalid identifier.</p>				
<b>ATTRIBUTES</b>	See <a href="#">attributes(5)</a> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>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				
<b>SEE ALSO</b>	<a href="#">ipcReceive(2K)</a>				

<b>NAME</b>	ipcSave, ipcRestore – Save the current message; Restore a saved message as the current message				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int ipcSave(int portId);  int ipcRestore(int msgid);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p>The <i>ipcSave</i> call saves the current thread's current message to a system buffer, and returns an identifier for the saved message. The <i>portId</i> argument reserved for future use. The caller should use the K_NONEPORT value for this parameter.</p> <p>Messages used to be saved when a thread receives several messages and has to reply in a different order ( <i>ipcReply</i> (2K) is always applied to the current message; receiving a message supplants the current message). After <i>ipcSave</i>, there is no current message, and <i>ipcReply</i> (2K), <i>ipcGetData</i> (2K) or <i>ipcSysInfo</i> (2K) cannot be used unless a message is received (<i>ipcReceive</i> (2K)) or restored (<i>ipcRestore</i> (2K)).</p> <p>The <i>ipcRestore</i> call restores a saved message as the current message. If the call is successful, the previous current message is replaced by the current message .</p> <p>The number of messages that an actor may save is limited to K_CMSGSAVEDMAX.</p>				
<b>RETURN VALUE</b>	Upon successful completion, <i>ipcRestore</i> returns 0 and <i>ipcSave</i> returns a positive identifier for the saved message. Otherwise, a negative error code is returned. In case of failure, the current message is not erased.				
<b>ERRORS</b>	<p>[K_ENOMEM]                    <i>ipcSave</i> : K_CMSGSAVEDMAX messages have already been saved by the calling actor.</p> <p>[K_EINVAL]                    <i>ipcSave</i> : no current message; <i>ipcRestore</i> : invalid identifier.</p>				
<b>ATTRIBUTES</b>	See <i>attributes</i> (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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>ipcReceive</i> (2K)				

<b>NAME</b>	ipcSend – send a message
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int ipcSend(KnMsgDesc *msg, int msgsrc, KnIpcDest *msgdest);</pre>
<b>FEATURES</b>	IPC
<b>DESCRIPTION</b>	<p><i>ipcSend</i> sends an asynchronous message. <i>msgdest</i> defines the destination of the message. It is a pointer to a <i>KnIpcDest</i> structure the members of which are the following:</p> <pre>KnUniqueId    target ;           /* port or group */ KnUniqueId    coTarget ;        /* site qualifier */</pre> <p>If <i>target</i> is a port name, the message will be sent to that port. <i>target</i> may also be a port group name in which an addressing mode has been set using <i>ipcTarget(2K)</i>. If the addressing mode is <i>K_BROADMODE</i>, the message will be sent to each reachable member of the group (i.e. ports which have been inserted in the group - see <i>grpPortInsert(2K)</i>). If the addressing mode is <i>K_FUNCMODE</i>, the message will be sent to one of the reachable members of the group. If the addressing mode is <i>K_FUNCUMODE</i>, the message will be sent to one of the reachable members of the group which reside on the site denoted by <i>coTarget</i>. <i>coTarget</i> is a unique identifier. This addressing mode is called <i>associative functional mode</i>. A site is denoted by such a unique identifier, if this unique identifier has been declared on the site. For example, the names of the ports owned by the actors located on the site and the pre-defined site unique identifiers (see <i>uiSite(2K)</i>) are identifiers which are automatically registered on a site. If the addressing mode is <i>K_FUNCXMODE</i>, the message will be sent to one of the reachable members of the group, assuming that the UI of the selected member is different from the UI given by <i>coTarget</i>. This addressing mode is called <i>exclusive functional mode</i>.</p> <p><i>msgsrc</i> is the local identifier of a port owned by the sender, and specifies the source port of the message. If one of the receivers of the message performs a <i>ipcReply(2K)</i> on it, the reply will be sent to that port. The unique identifier of this port is given to a receiver which performs a <i>ipcSysInfo(2K)</i>, as the name of the source port of the message. If <i>msgsrc</i> is <i>K_DEFAULTPORT</i>, the current actor default port is used.</p> <p><i>msg</i> points to a descriptor for the message to be sent, and is a pointer to a <i>KnMsgDesc</i> structure, the members of which are the following:</p> <pre>unsigned int   flags ;           /* message structure definition */ unsigned int   bodySize ;        /* body size */ VmAddr        bodyAddr ;        /* body address */ VmAddr        annexAddr ;      /* fixed size annex address */ KnEvtNum      seqNum ;          /* not used */</pre>

The message data is composed of a `message body`, a byte string of variable size (limited to `K_CMSGsizEMAX`), to which a `message annex`, a small fixed sized (`K_CMSGANNEXSIZE`) byte string, might be associated. The body transmission may benefit from the (virtual) memory management to avoid physical copies (see below). When present, the annex is always physically copied from the sender address space to the receiver address space.

The `bodySize` and `bodyAddr` members of the message descriptor respectively give the size of the message body and its starting address in the sender address space. `bodySize` is limited to `K_CMSGsizEMAX`. If `bodySize` is set to 0, the message contains no body data.

The `annexAddr` member of the message descriptor gives the starting address of the message annex in the sender address space. If this field is `NULL`, no annex is sent with the message body.

The `flags` field of the message descriptor may be a combination of the following options, which describe the way in which the message body is to be transmitted:

**K\_MOVE** (for virtual memory management only) This option indicates that the sender will not retain the contents of the message body in its address space after the send operation. In that case, the system will try to transfer the message body without local copy, by transferring only the page descriptors. After the send, the virtual addresses corresponding to the message body are still valid, but their contents are undefined.

If the `K_MOVE` option is not set, the system will copy the message body.

**K\_MOVEAL** (for virtual memory management only) This option has the same effect as the `K_MOVE` option, except that even if the message body end is not aligned on a page boundary (see `vmPageSize(2K)`), the last page partially covered by the message may be moved by the system. That means that the user must be ready to lose the contents of its address space from the beginning of the message body to the next page boundary following the end of the message body.

**K\_USERBODY** If the caller thread is a `SUPERVISOR` thread, this flag indicates that the message body is part of the current user address space. If this flag is not set while the caller is a `SUPERVISOR` thread, the message body is assumed to be part of the kernel address space. This flag is only to be used by trap handling routines, when data located in the user address space have to be directly used as a message

body, without copying them into the kernel address space. If the caller thread is not a SUPERVISOR thread, this flag is ignored (the message body is always assumed to be part of the user address space).

K\_USERANNEX This flag has exactly the same meaning as the K\_USERBODY, but concerns the message annex instead of the message body.

If flag is equal to 0, none of the previous options is selected: the message body will be copied (never "moved"), and message and annex are assumed to be part of the user (kernel) address space when the caller is a USER (SUPERVISOR) thread.

The message is sent asynchronously: the sender is only blocked during the time needed by the system to queue the send request.

**RETURN VALUE**

Upon successful completion, a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

- [K\_EFULL] The destination port is local and its queue is full, or the local communication subsystem is saturated.
- [K\_EINVAL] msgsrc is not a valid local port identifier.
- [K\_EFAULT] Some of the provided data are outside the current actor's address space.
- [K\_ETOOMUCH] bodysize is too big.
- [K\_EUNKNOWN] Unreachable destination.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

ipcReply(2K), ipcReceive(2K), ipcTarget(2K), ipcSysInfo(2K)

<b>NAME</b>	ipcSysInfo – get system information about the current message
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int ipcSysInfo(KnMsgHead *msghead);</pre>
<b>FEATURES</b>	IPC
<b>DESCRIPTION</b>	<p>The <i>ipcSysInfo</i> call returns system information about the current message received. The <i>msghead</i> field points to a <i>KnMsgHead</i> structure whose members are the following:</p> <pre>int      bodySize ;      /* Body size */ char     flags [4] ; /* Flags */ KnUniqueId srcPort ;    /* source port */ KnUniqueId target ;     /* destination port or group */ KnUniqueId coTarget ;   /* site qualifier */ KnTransId transId ;     /* message transaction id. */ KnProtId  actorId ;     /* sender actor protection id. */ KnProtId  portId ;      /* source port protection id. */</pre> <p>The <i>bodySize</i> field is the size of the message body.</p> <p>The <i>srcPort</i> field is the unique identifier of the source port of the message.</p> <p>The <i>target</i> and <i>coTarget</i> parameters specified the destination when the message was sent (see <i>ipcSend(2K)</i> and <i>ipcCall(2K)</i>).</p> <p>The <i>transId</i> field is the transaction identifier of the message.</p> <p>The protection identifier of the actor originating the message is specified using <i>actorId</i>.</p> <p>The protection identifier of the source port of the message is specified using <i>portId</i>.</p> <p>Note that <i>ipcReply(2K)</i> does not affect the source port, the protection identifiers and the transaction identifier of a message.</p>
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	<p>[K_EINVAL] There is no current message (no message has been received by the current thread since the last <i>ipcReply(2K)</i> or <i>ipcSave(2K)</i>).</p> <p>[K_EFAULT] Some of the data provided are outside the current actor's address space.</p>
<b>RESTRICTIONS</b>	The fields <i>flags</i> and <i>transId</i> are obsolete and are never set by the kernel.
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

ipcReply(2K), ipcReceive(2K), ipcSend(2K), ipcCall(2K)

<b>NAME</b>	ipcTarget – build a message target
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int ipcTarget(KnUniqueId *groupui, int mode);</pre>
<b>FEATURES</b>	IPC
<b>DESCRIPTION</b>	<p><i>ipcTarget</i> builds a message destination using the port group unique identifier given by <i>groupui</i> and the addressing mode <i>mode</i>. It affects the port group identifier given by <i>groupui</i>. This identifier may then be used as a target for sending messages by <i>ipcSend(2K)</i> or <i>ipcCall(2K)</i>.</p> <p>The following addressing modes are possible:</p> <p><b>K_BROADMODE</b>            The message will be sent to each reachable member of the group (<i>ie.</i> ports which have been inserted in the group - see <i>grpPortInsert(2K)</i>).</p> <p><b>K_FUNCMODE</b>              The message will be sent to one of the reachable members of the group.</p> <p><b>K_FUNCUMODE</b>            The message will be sent to one of the reachable members of the group which resides on the site qualified by the <i>coTarget</i> member specified at send time (see <i>ipcSend(2K)</i>).</p> <p><b>K_FUNCXMODE</b>            The message will be sent to one of the reachable members of the group, assuming that the UI of this member must be different from the UI given by the <i>coTarget</i> member specified at send time (see <i>ipcSend(2K)</i>).</p> <p>A port group name may be used directly as a message target, without applying <i>ipcTarget(2K)</i> on it. In that case, the default addressing mode <b>K_BROADMODE</b> is used.</p>
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	<p>[<b>K_EINVAL</b>]              <i>groupui</i> is not a port group name, or <i>mode</i> is not a valid addressing mode.</p> <p>[<b>K_EFAULT</b>]              Some of the provided data are outside the current actor's address space.</p>
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

grpAllocate(2K), ipcSend(2K), ipcCall(2K), grpPortInsert(2K)

<b>NAME</b>	svLapCreate, lapDescZero, lapDescIsZero, lapDescDup – create a lap; reset a lap descriptor; test if a lap descriptor has been initialized; duplicate a lap descriptor
<b>SYNOPSIS</b>	<pre>#include &lt;lap/chLap.h&gt; int svLapCreate(KnCap * actorcap, KnLapHdl laphdl, void * cookie, unsigned int options, KnLapDesc * lapdesc);  void lapDescZero(KnLapDesc * lapdesc);  int lapDescIsZero(KnLapDesc * lapdesc);  void lapDescDup(KnLapDesc * olddesc, KnLapDesc * newdesc);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svLapCreate</i> (2K) system call creates a new local access point (lap) for the actor designated by the <i>actorcap</i> capability. If <i>actorcap</i> is equal to <code>K_MYACTOR</code>, the home actor of the calling thread is considered. If the call succeeds, the lap descriptor pointed to by <i>lapdesc</i> represents this new lap and can be passed to client threads that will use it to invoke the lap handler.</p> <p>On lap invocation (see <i>lapInvoke</i> (2K)), the <i>laphdl</i> handler is called. This handler is a function which takes two arguments:</p> <pre>void handler (arg, cookie) void          *arg ; void          *cookie ;</pre> <p>Where <i>arg</i> is the argument specified by <i>lapInvoke</i> (2K), and <i>cookie</i> is the value of the cookie specified by <i>svLapCreate</i> (2K).</p> <p>By default as a consequence of a lap invocation, the execution actor of the calling thread are set to the lap owning actor for the duration of the lap handler execution.</p> <p>Additional actions may take place during the lap invocation, depending on the combination of flags specified in the <i>options</i> parameter:</p> <p><b>K_LAP_SAFE</b>                    A "lap frame" descriptor is allocated to register the calling thread as a temporary resource of the invoked actor. Each "lap frame" has an associated level (lap frame level) which represents the number of lap frames for the considered thread.</p> <p>                                  This option is only valid if the LAPSAFE feature has been specified into the system. It will enforce a stronger checking during lap invocation. This guarantees in particular that the kernel will synchronize the <i>svLapDelete</i> (2K) operation with concurrent lap invocations.</p>

Futhermore, full context of the calling thread is saved prior to the lap invocation. This guarantees that the calling thread can return from its invocation even if a failure (exception, deletion of the lap owning actor, etc.) occurs during the execution of the lap handler.

This option is mandatory for laps to be called from user mode.

The two utility routines *lapDescZero* (2K) and *lapDescIsZero* (2K) are available to manipulate the state of a lap descriptor with the following semantics: if a lap descriptor has been initialized with *lapDescZero* (2K) or is implemented in a zero-filled memory region, *lapDescIsZero* (2K) returns a non-zero value until the lap descriptor has been successfully initialized by *svLapCreate* (2K) or *lapResolve* (2K).

*lapDescDup* (2K) must be used to duplicate the contents of a lap descriptor.

**RETURN VALUE**

On success, *svLapCreate* (2K) returns K\_OK. Otherwise, a negative error code is returned.

*lapDescIsZero* (2K) returns a non-zero value if the lap descriptor is not initialized.

*lapDescZero* (2K) and *lapDescDup* (2K) have no return values.

**ERRORS**

[K\_EINVAL] *actorcap* is an inconsistent actor capability.

[K\_ENOMEM] The system is out of resources.

**RESTRICTIONS**

The current implementation does not support the K\_LAP\_SETJMP option, which is silently ignored.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke*(2K) , *lapResolve*(2K) , *svLapBind*(2K) , *svLapDelete*(2K) , *svLapUnbind*(2K) , *threadStat*(2K)

<b>NAME</b>	svLapCreate, lapDescZero, lapDescIsZero, lapDescDup – create a lap; reset a lap descriptor; test if a lap descriptor has been initialized; duplicate a lap descriptor
<b>SYNOPSIS</b>	<pre>#include &lt;lap/chLap.h&gt; int svLapCreate(KnCap * actorcap, KnLapHdl laphdl, void * cookie, unsigned int options, KnLapDesc * lapdesc);  void lapDescZero(KnLapDesc * lapdesc);  int lapDescIsZero(KnLapDesc * lapdesc);  void lapDescDup(KnLapDesc * olddesc, KnLapDesc * newdesc);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svLapCreate</i> (2K) system call creates a new local access point (lap) for the actor designated by the <i>actorcap</i> capability. If <i>actorcap</i> is equal to <code>K_MYACTOR</code>, the home actor of the calling thread is considered. If the call succeeds, the lap descriptor pointed to by <i>lapdesc</i> represents this new lap and can be passed to client threads that will use it to invoke the lap handler.</p> <p>On lap invocation (see <i>lapInvoke</i> (2K)), the <i>laphdl</i> handler is called. This handler is a function which takes two arguments:</p> <pre>void handler (arg, cookie) void          *arg ; void          *cookie ;</pre> <p>Where <i>arg</i> is the argument specified by <i>lapInvoke</i> (2K), and <i>cookie</i> is the value of the cookie specified by <i>svLapCreate</i> (2K).</p> <p>By default as a consequence of a lap invocation, the execution actor of the calling thread are set to the lap owning actor for the duration of the lap handler execution.</p> <p>Additional actions may take place during the lap invocation, depending on the combination of flags specified in the <i>options</i> parameter:</p> <p><b>K_LAP_SAFE</b>                    A "lap frame" descriptor is allocated to register the calling thread as a temporary resource of the invoked actor. Each "lap frame" has an associated level (lap frame level) which represents the number of lap frames for the considered thread.</p> <p>                                  This option is only valid if the LAPSAFE feature has been specified into the system. It will enforce a stronger checking during lap invocation. This guarantees in particular that the kernel will synchronize the <i>svLapDelete</i> (2K) operation with concurrent lap invocations.</p>

Futhermore, full context of the calling thread is saved prior to the lap invocation. This guarantees that the calling thread can return from its invocation even if a failure (exception, deletion of the lap owning actor, etc.) occurs during the execution of the lap handler.

This option is mandatory for laps to be called from user mode.

The two utility routines *lapDescZero* (2K) and *lapDescIsZero* (2K) are available to manipulate the state of a lap descriptor with the following semantics: if a lap descriptor has been initialized with *lapDescZero* (2K) or is implemented in a zero-filled memory region, *lapDescIsZero* (2K) returns a non-zero value until the lap descriptor has been successfully initialized by *svLapCreate* (2K) or *lapResolve* (2K).

*lapDescDup* (2K) must be used to duplicate the contents of a lap descriptor.

**RETURN VALUE**

On success, *svLapCreate* (2K) returns K\_OK. Otherwise, a negative error code is returned.

*lapDescIsZero* (2K) returns a non-zero value if the lap descriptor is not initialized.

*lapDescZero* (2K) and *lapDescDup* (2K) have no return values.

**ERRORS**

[K\_EINVAL] *actorcap* is an inconsistent actor capability.

[K\_ENOMEM] The system is out of resources.

**RESTRICTIONS**

The current implementation does not support the K\_LAP\_SETJMP option, which is silently ignored.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke*(2K) , *lapResolve*(2K) , *svLapBind*(2K) , *svLapDelete*(2K) , *svLapUnbind*(2K) , *threadStat*(2K)

<b>NAME</b>	svLapCreate, lapDescZero, lapDescIsZero, lapDescDup – create a lap; reset a lap descriptor; test if a lap descriptor has been initialized; duplicate a lap descriptor
<b>SYNOPSIS</b>	<pre>#include &lt;lap/chLap.h&gt; int svLapCreate(KnCap * actorcap, KnLapHdl laphdl, void * cookie, unsigned int options, KnLapDesc * lapdesc);  void lapDescZero(KnLapDesc * lapdesc);  int lapDescIsZero(KnLapDesc * lapdesc);  void lapDescDup(KnLapDesc * olddesc, KnLapDesc * newdesc);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svLapCreate</i> (2K) system call creates a new local access point (lap) for the actor designated by the <i>actorcap</i> capability. If <i>actorcap</i> is equal to <code>K_MYACTOR</code>, the home actor of the calling thread is considered. If the call succeeds, the lap descriptor pointed to by <i>lapdesc</i> represents this new lap and can be passed to client threads that will use it to invoke the lap handler.</p> <p>On lap invocation (see <i>lapInvoke</i> (2K)), the <i>laphdl</i> handler is called. This handler is a function which takes two arguments:</p> <pre>void handler (arg, cookie) void          *arg ; void          *cookie ;</pre> <p>Where <i>arg</i> is the argument specified by <i>lapInvoke</i> (2K), and <i>cookie</i> is the value of the cookie specified by <i>svLapCreate</i> (2K).</p> <p>By default as a consequence of a lap invocation, the execution actor of the calling thread are set to the lap owning actor for the duration of the lap handler execution.</p> <p>Additional actions may take place during the lap invocation, depending on the combination of flags specified in the <i>options</i> parameter:</p> <p><b>K_LAP_SAFE</b>                    A "lap frame" descriptor is allocated to register the calling thread as a temporary resource of the invoked actor. Each "lap frame" has an associated level (lap frame level) which represents the number of lap frames for the considered thread.</p> <p>                                  This option is only valid if the LAPSAFE feature has been specified into the system. It will enforce a stronger checking during lap invocation. This guarantees in particular that the kernel will synchronize the <i>svLapDelete</i> (2K) operation with concurrent lap invocations.</p>

Futhermore, full context of the calling thread is saved prior to the lap invocation. This guarantees that the calling thread can return from its invocation even if a failure (exception, deletion of the lap owning actor, etc.) occurs during the execution of the lap handler.

This option is mandatory for laps to be called from user mode.

The two utility routines *lapDescZero* (2K) and *lapDescIsZero* (2K) are available to manipulate the state of a lap descriptor with the following semantics: if a lap descriptor has been initialized with *lapDescZero* (2K) or is implemented in a zero-filled memory region, *lapDescIsZero* (2K) returns a non-zero value until the lap descriptor has been successfully initialized by *svLapCreate* (2K) or *lapResolve* (2K).

*lapDescDup* (2K) must be used to duplicate the contents of a lap descriptor.

**RETURN VALUE**

On success, *svLapCreate* (2K) returns K\_OK. Otherwise, a negative error code is returned.

*lapDescIsZero* (2K) returns a non-zero value if the lap descriptor is not initialized.

*lapDescZero* (2K) and *lapDescDup* (2K) have no return values.

**ERRORS**

[K\_EINVAL] *actorcap* is an inconsistent actor capability.

[K\_ENOMEM] The system is out of resources.

**RESTRICTIONS**

The current implementation does not support the K\_LAP\_SETJMP option, which is silently ignored.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke*(2K) , *lapResolve*(2K) , *svLapBind*(2K) , *svLapDelete*(2K) , *svLapUnbind*(2K) , *threadStat*(2K)

<b>NAME</b>	lapInvoke – invoke a lap handler				
<b>SYNOPSIS</b>	<pre>#include &lt;lap/chLap.h&gt; int lapInvoke(KnLapDesc *lapdesc, void *arg);</pre>				
<b>DESCRIPTION</b>	<p>The <i>lapInvoke(2K)</i> system call invokes a lap handler in the context of the calling thread (see <i>svLapCreate(2K)</i> for a description of the exact invocation semantics). The lap considered here is represented by its lap descriptor pointed to by <i>lapdesc</i>.</p> <p>The lap handler is called with two parameters: the <i>arg</i> argument and the cookie specified at lap creation time (see <i>svLapCreate(2K)</i>).</p>				
<b>RETURN VALUES</b>	On success, <i>lapInvoke(2K)</i> returns K_OK, otherwise a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EFAIL]                   The lap was created with the K_LAP_SAFE option and the lap has been deleted when the invocation is in progress.</p> <p>[K_EFAIL]                   The lap was created with the K_LAP_SETJMP option and a failure occurred during the lap invocation.</p> <p>[K_EFAULT]                  Some of the provided data are outside the current actor's address space.</p> <p>[K_EINVAL]                  The lap was created with the K_LAP_SAFE option and it has been deleted. <i>lapInvoke(2K)</i> is called from user mode and the invoked lap was not created with the K_LAP_SAFE option.</p>				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>svLapCreate(2K)</i> , <i>svLapDelete(2K)</i>				

<b>NAME</b>	svLapBind, svLapUnbind, lapResolve – bind a symbolic name to a lap descriptor; unbind the symbolic name bound to a lap descriptor; get a lap descriptor from a lap symbolic name										
<b>SYNOPSIS</b>	<pre>#include &lt;lap/chLap.h&gt; int svLapBind(KnLapDesc * lapdesc, char * name, unsigned int options);  int svLapUnbind(char * name);  int lapResolve(KnLapDesc * lapdesc, char * name, unsigned int options);</pre>										
<b>FEATURES</b>	LAPBIND										
<b>DESCRIPTION</b>	<p>The <i>svLapBind</i> (2K) system call binds the lap descriptor pointed to by <i>lapdesc</i> with the symbolic name pointed to by <i>name</i> .</p> <p><i>name</i> points to a null-terminated string of K_LAPNAMEMAX characters at most (not including the null-terminating character).</p> <p>If the K_LAP_PROTECTED option is set in the <i>options</i> parameter, the binding will only be visible to trusted threads (e.g. threads executing in user actors will get an error from <i>lapResolve</i> (2K) on this name).</p> <p>The <i>svLapUnbind</i> (2K) system call removes the lap binding associated to <i>name</i> .</p> <p>The <i>lapResolve</i> (2K) system call initializes the lap descriptor pointed to by <i>lapdesc</i> with the lap descriptor bound to <i>name</i> . <i>lapResolve</i> (2K) will block until a lap descriptor is bound to <i>name</i> , except if the K_LAP_NOBLOCK option is set in <i>options</i> , in which case the call returns immediatly with an error.</p> <p><i>svLapBind</i> and <i>svLapUnbind</i> are restricted to SUPERVISOR threads.</p>										
<b>RETURN VALUE</b>	On success, these calls return K_OK. Otherwise, a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EABORT]</td> <td>The calling thread has been aborted in <i>lapResolve</i> (2K).</td> </tr> <tr> <td style="vertical-align: top;">[K_EBUSY]</td> <td>The name given to <i>svLapBind</i> (2K) is already in use.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>The <i>lapdesc</i> or <i>name</i> arguments of <i>lapResolve</i> points to the outside of the caller's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td>The <i>name</i> given to <i>svLapUnbind</i> (2K) is not bound to any lap, or the <i>name</i> given to <i>lapResolve</i> (2K) is not bound to any lap, and the K_LAP_NOBLOCK option is set.</td> </tr> <tr> <td style="vertical-align: top;">[K_ENOMEM]</td> <td>The <i>name</i> given to <i>svLapBind</i> (2K), <i>svLapUnbind</i> (2K) or <i>lapResolve</i> (2K) is too long. [K_ENOMEM] The system is out of resources.</td> </tr> </table>	[K_EABORT]	The calling thread has been aborted in <i>lapResolve</i> (2K).	[K_EBUSY]	The name given to <i>svLapBind</i> (2K) is already in use.	[K_EFAULT]	The <i>lapdesc</i> or <i>name</i> arguments of <i>lapResolve</i> points to the outside of the caller's address space.	[K_EINVAL]	The <i>name</i> given to <i>svLapUnbind</i> (2K) is not bound to any lap, or the <i>name</i> given to <i>lapResolve</i> (2K) is not bound to any lap, and the K_LAP_NOBLOCK option is set.	[K_ENOMEM]	The <i>name</i> given to <i>svLapBind</i> (2K), <i>svLapUnbind</i> (2K) or <i>lapResolve</i> (2K) is too long. [K_ENOMEM] The system is out of resources.
[K_EABORT]	The calling thread has been aborted in <i>lapResolve</i> (2K).										
[K_EBUSY]	The name given to <i>svLapBind</i> (2K) is already in use.										
[K_EFAULT]	The <i>lapdesc</i> or <i>name</i> arguments of <i>lapResolve</i> points to the outside of the caller's address space.										
[K_EINVAL]	The <i>name</i> given to <i>svLapUnbind</i> (2K) is not bound to any lap, or the <i>name</i> given to <i>lapResolve</i> (2K) is not bound to any lap, and the K_LAP_NOBLOCK option is set.										
[K_ENOMEM]	The <i>name</i> given to <i>svLapBind</i> (2K), <i>svLapUnbind</i> (2K) or <i>lapResolve</i> (2K) is too long. [K_ENOMEM] The system is out of resources.										

[K\_EPRIV]

*lapResolve* (2K) is called by a non trusted thread to resolve a name which has been bound with the K\_LAP\_PROTECTED option.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svLapCreate(2K)`

<b>NAME</b>	msgAllocate – allocates a message from a message space	
<b>SYNOPSIS</b>	<pre>#include &lt;mipc/chMipc.h&gt; int msgAllocate(int spaceLid, unsigned int poolIndex, unsigned int msgSize, KnTimeVal *waitLimit, char **msgAddr);</pre>	
<b>FEATURES</b>	MIPC	
<b>DESCRIPTION</b>	<p><i>msgAllocate</i> attempts to allocate a message from the most appropriate message pool of the message space designated by <i>spaceLid</i>.</p> <p>With MIPC_S, the only valid value for <i>spaceLid</i> is 0.</p> <p><i>poolIndex</i> is either the index of the message pool from which the message must be allocated or the specific K_ANY_MSGPOOL constant.</p> <p><i>msgSize</i> is the size of the application message.</p> <p><i>waitLimit</i> optionally specifies a timeout, and may also specify a non abortable wait.</p> <p><i>msgAddr</i> is the address of a pointer where to store the address of the allocated message.</p> <p><i>msgAllocate</i> considers the message pool which is designated by <i>poolIndex</i>, if it is not equal to K_ANY_MSGPOOL, or parses the list of message pools of the Message Space to find the pool which fits the requested message size. If the message pool is not empty, <i>msgAllocate</i> takes the first free message, computes its address within the addressing space of the caller, and stores it into the <i>msgAddr</i> pointer.</p> <p>If the pool is empty, <i>msgAllocate</i> does not try to allocate a message from a pool of bigger messages, if any. If <i>waitLimit</i> is equal to K_NOBLOCK(0), <i>msgAllocate</i> immediately returns a K_ETIMEOUT error code. Otherwise, it blocks the invoking thread until a message of the pool is freed by <i>msgFree</i>(2K) or until the timeout, if any, expires.</p>	
<b>RESTRICTIONS</b>	Interrupt handlers must always invoke <i>msgAllocate</i> with a delay of 0.	
<b>RETURN VALUE</b>	Upon successful completion, <i>msgAllocate</i> returns zero. Otherwise, a negative error code is returned.	
<b>ERRORS</b>	[K_EINVAL]	<i>spaceLid</i> is not valid, or <i>poolIndex</i> is not valid, or <i>waitLimit</i> is not a valid <i>KnTimeVal</i> .
	[K_ESIZE]	<i>msgSize</i> is greater than the maximum message size.
	[K_ETIMEOUT]	The maximum waiting time specified in <i>waitLimit</i> expired.

[K\_EFAULT] *msgAddr* is not a valid address within the addressing space of the caller.

[K\_EABORT] The invoking thread was aborted while waiting for a message.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`msgSpaceCreate(2K)`, `msgSpaceOpen(2K)`, `msgFree(2K)`, `msgPut(2K)`, `msgGet(2K)`, `msgRemove(2K)`

**NAME** | msgFree – free a message of a message space

**SYNOPSIS** | #include <mipc/chMipc.h>  
 int msgFree(int spaceLid, char \*msg);

**FEATURES** | MIPC

**DESCRIPTION** | *msgFree* returns to the message space identified by *spaceLid* a message, which was previously allocated by *msgAllocate(2K)* from the same message space, or which was previously received by *msgGet(2K)* from a message queue of the same message space.

With MIPC\_S, the only valid value for *spaceLid* is 0.

*msg* is the address of the message within the addressing space of the caller.

*msgFree* converts *msg* into the corresponding internal message resource, from which it finds its associated message pool. If a thread is blocked waiting for a message of the pool to become available, it is awakened, and the freed message is given to it. Otherwise, the freed message is appended to the pool.

**RETURN VALUE** | Upon successful completion, *msgFree* returns zero. Otherwise, a negative error code is returned.

**ERRORS** | [K\_EINVAL]                      *spaceLid* is not valid.  
 [K\_EBADMSG]                      *msg* is not a valid message address, or the corresponding internal message resource is not in the allocated state.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | msgSpaceCreate(2K), msgSpaceOpen(2K), msgAllocate(2K), msgPut(2K), msgGet(2K), msgRemove(2K)

<b>NAME</b>	<code>msgGet</code> – retrieves the first message of a message queue								
<b>SYNOPSIS</b>	<pre>#include &lt;mipc/chMipc.h&gt; int msgGet(int spaceLid, unsigned int msgQueueId, KnTimeVal *waitLimit, char **msgAddr, KnUniqueId *srcActor);</pre>								
<b>FEATURES</b>	MIPC								
<b>DESCRIPTION</b>	<p><code>msgGet</code> returns to the caller the first message with the highest priority pending behind the message queue of index <code>msgQueueId</code> within the set of message queues of the message space designated by <code>spaceLid</code>.</p> <p><code>waitLimit</code> optionally specifies a timeout, and may also specify a non abortable wait.</p> <p><code>msgAddr</code> is the address of a pointer where to store the address of the retrieved message.</p> <p><code>srcActor</code> is either equal to the specific NULL constant or points to a structure of type <code>KnUniqueId</code> where to store the unique identifier of the source actor of the retrieved message, if any.</p> <p>The first message with the highest priority pending behind the message queue is dequeued.</p> <p>If no message is pending behind the message queue, the invoking thread is blocked until a message is sent to the message queue or until the timeout, if any, expires. Multiple threads can be blocked behind the same message queue. In such a case, the policy for ordering incoming messages is a simple FIFO. No data copy is performed. The receiving thread is returned a pointer to the message allocated by the sender. This thread can immediately access the content of the message, and update it before sending the message (without any extra data copy) to another message queue of the same message space.</p>								
<b>RETURN VALUE</b>	Upon successful completion, <code>msgGet</code> returns 0. Otherwise, a negative error code is returned.								
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td><code>spaceLid</code> is not valid, or <code>waitLimit</code> is not a valid <code>KnTimeVal</code>.</td> </tr> <tr> <td style="vertical-align: top;">[K_EARGS]</td> <td><code>msgQueueId</code> is not a valid message queue.</td> </tr> <tr> <td style="vertical-align: top;">[K_ETIMEOUT]</td> <td>The maximum waiting time specified in <code>waitLimit</code> expired.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>One of the following pointers is invalid: <code>waitLimit</code>, <code>msgAddr</code> or <code>srcActor</code>. If <code>waitLimit</code> is an invalid pointer, <code>msgGet</code> has no effect on the message queue designated by <code>msgQueueId</code> and returns immediately K_EFAULT. If <code>srcActor</code> or</td> </tr> </table>	[K_EINVAL]	<code>spaceLid</code> is not valid, or <code>waitLimit</code> is not a valid <code>KnTimeVal</code> .	[K_EARGS]	<code>msgQueueId</code> is not a valid message queue.	[K_ETIMEOUT]	The maximum waiting time specified in <code>waitLimit</code> expired.	[K_EFAULT]	One of the following pointers is invalid: <code>waitLimit</code> , <code>msgAddr</code> or <code>srcActor</code> . If <code>waitLimit</code> is an invalid pointer, <code>msgGet</code> has no effect on the message queue designated by <code>msgQueueId</code> and returns immediately K_EFAULT. If <code>srcActor</code> or
[K_EINVAL]	<code>spaceLid</code> is not valid, or <code>waitLimit</code> is not a valid <code>KnTimeVal</code> .								
[K_EARGS]	<code>msgQueueId</code> is not a valid message queue.								
[K_ETIMEOUT]	The maximum waiting time specified in <code>waitLimit</code> expired.								
[K_EFAULT]	One of the following pointers is invalid: <code>waitLimit</code> , <code>msgAddr</code> or <code>srcActor</code> . If <code>waitLimit</code> is an invalid pointer, <code>msgGet</code> has no effect on the message queue designated by <code>msgQueueId</code> and returns immediately K_EFAULT. If <code>srcActor</code> or								

*msgAddr* is or becomes an invalid pointer (while *msgGet*, the message that has been dequeued is freed. Therefore the message is lost from the faulty application point of view.

[K\_EABORT]

The invoking thread was aborted while waiting for a message.

#### ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

#### SEE ALSO

`msgSpaceCreate(2K)`, `msgSpaceOpen(2K)`, `msgAllocate(2K)`, `msgFree(2K)`, `msgPut(2K)`, `msgRemove(2K)`

<b>NAME</b>	msgPut – post a message to a message queue						
<b>SYNOPSIS</b>	<pre>#include &lt;mipc/chMipc.h&gt; int msgPut(int spaceLid, unsigned int msgQueueId, char *msg, unsigned int prio);</pre>						
<b>FEATURES</b>	MIPC						
<b>DESCRIPTION</b>	<p><i>msgPut</i> sends a message to the message queue of index <i>msgQueueId</i> in the message space identified by <i>spaceLid</i>.</p> <p><i>msg</i> is the address within the addressing space of the caller of a message which was previously allocated (by <i>msgAllocate(2K)</i>) from the same message space, or which has been retrieved (by <i>msgGet(2K)</i>) from a message queue of the same message space.</p> <p><i>prio</i> is the priority of the message. Its value must fall between 0 (the lowest priority) and the specific K_MSG_PRIOMAX constant which is the highest priority. If <i>prio</i> is greater than K_MSG_PRIOMAX, it is silently dropped to K_MSG_PRIOMAX.</p> <p>If a thread is blocked waiting for a message to arrive, it is awakened, and the message is immediately given to it. Otherwise, the message is inserted in the message queue according to its priority.</p>						
<b>RETURN VALUE</b>	Upon successful completion, <i>msgPut</i> returns 0. Otherwise, a negative error code is returned.						
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EINVAL]</td> <td><i>spaceLid</i> is not valid.</td> </tr> <tr> <td>[K_EARGS]</td> <td><i>msgQueueId</i> is not a valid message queue.</td> </tr> <tr> <td>[K_EBADMSG]</td> <td><i>msg</i> is not a valid message address, or the corresponding message resource is not in the allocated state.</td> </tr> </table>	[K_EINVAL]	<i>spaceLid</i> is not valid.	[K_EARGS]	<i>msgQueueId</i> is not a valid message queue.	[K_EBADMSG]	<i>msg</i> is not a valid message address, or the corresponding message resource is not in the allocated state.
[K_EINVAL]	<i>spaceLid</i> is not valid.						
[K_EARGS]	<i>msgQueueId</i> is not a valid message queue.						
[K_EBADMSG]	<i>msg</i> is not a valid message address, or the corresponding message resource is not in the allocated state.						
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
<b>SEE ALSO</b>	<i>msgSpaceCreate(2K)</i> , <i>msgSpaceOpen(2K)</i> , <i>msgAllocate(2K)</i> , <i>msgFree(2K)</i> , <i>msgGet(2K)</i> , <i>msgRemove(2K)</i>						

<b>NAME</b>	msgRemove – remove a message from a message queue						
<b>SYNOPSIS</b>	<pre>#include &lt;mipc/chMipc.h&gt; int msgRemove(int spaceLid, unsigned int msgQueueId, char *msg, unsigned int prio);</pre>						
<b>FEATURES</b>	MIPC						
<b>DESCRIPTION</b>	<p>The <i>msgRemove</i> call removes the message whose address is <i>msg</i> from the message queue of the index <i>msgQueueId</i> in the message space <i>spaceLid</i>.</p> <p>The <i>prio</i> argument must be identical to the priority provided to the <i>msgPut</i> system call previously invoked to post the message.</p>						
<b>RETURN VALUE</b>	Upon successful completion, <i>msgRemove</i> returns 0. Otherwise, a negative error code is returned.						
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td><i>spaceLid</i> is not valid.</td> </tr> <tr> <td style="vertical-align: top;">[K_EARGS]</td> <td><i>msgQueueId</i> is not a valid message queue.</td> </tr> <tr> <td style="vertical-align: top;">[K_EBADMSG]</td> <td><i>msg</i> is not a valid message address, or the corresponding message has already been consumed, or the priority is not the one affected when the message was put in the queue</td> </tr> </table>	[K_EINVAL]	<i>spaceLid</i> is not valid.	[K_EARGS]	<i>msgQueueId</i> is not a valid message queue.	[K_EBADMSG]	<i>msg</i> is not a valid message address, or the corresponding message has already been consumed, or the priority is not the one affected when the message was put in the queue
[K_EINVAL]	<i>spaceLid</i> is not valid.						
[K_EARGS]	<i>msgQueueId</i> is not a valid message queue.						
[K_EBADMSG]	<i>msg</i> is not a valid message address, or the corresponding message has already been consumed, or the priority is not the one affected when the message was put in the queue						
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
<b>SEE ALSO</b>	<i>msgSpaceCreate(2K)</i> , <i>msgSpaceOpen(2K)</i> , <i>msgAllocate(2K)</i> , <i>msgFree(2K)</i> , <i>msgGet(2K)</i> , <i>msgPut(2K)</i>						

<b>NAME</b>	msgSpaceCreate – create a message space
<b>SYNOPSIS</b>	<pre>#include &lt;mipc/chMipc.h&gt; int msgSpaceCreate(KnMsgSpaceId spaceGid, unsigned int msgQueueNb, unsigned int msgPoolNb, const KnMsgPool *msgPools);</pre>
<b>FEATURES</b>	MIPC
<b>DESCRIPTION</b>	<p>msgSpaceCreate() creates a message space.</p> <p><i>spaceGid</i> is the global identifier of the new message space. It is under the responsibility of applications to assign unique global identifiers to their message spaces. If <i>spaceGid</i> is equal to the special <code>K_PRIVATEID</code> constant, the message space is private to the invoking actor: its message queues and message pools will only be accessible to threads executing within this actor.</p> <p><i>msgQueueNb</i> is the number of message queues of the message space. Each message queue is designated by its index within the set of message queues (by an unsigned integer within the range <math>[0 .. \text{msgQueueNb} - 1]</math>).</p> <p><i>msgPoolNb</i> is the number of message pools of the communication space. Its value must be within the range <math>[1 .. \text{K\_MSG\_POOLMAX}]</math> where <code>K_MSG_POOLMAX</code> (16) is a constant defined at kernel compile time.</p> <p><i>msgPool</i> is a pointer to an array of <code>poolNb</code> structures of type <code>KnMsgPool</code>. The <code>KnMsgPool</code> structure includes the following information:</p> <pre>unsigned int msgSize ; unsigned int msgNumber ;</pre>
<b>RETURN VALUES</b>	Upon successful completion, the memory allocated for the message pools is mapped within the addressing space of the actor, and <code>msgSpaceCreate()</code> returns a positive local identifier of the message space. Otherwise, a negative error code is returned.
<b>ERRORS</b>	<p>[<code>K_EINVAL</code>] <i>spaceGid</i> is not equal to <code>K_PRIVATEID</code>, and is already assigned to another existing message space.</p> <p>[<code>K_EOVERLAP</code>] There is not sufficient room in the addressing space of the actor to map the message pools of the created message space.</p> <p>[<code>K_EARGS</code>] The value of <i>msgQueueNb</i> is equal to zero, or the value of <i>msgPoolNb</i> is out of the range <math>[1 .. \text{K\_MSG\_POOLMAX}]</math>.</p> <p>[<code>K_EFAULT</code>] Some of the data provided are outside the current actor's address space.</p>

[K\_ENOMEM] The system was not able to allocate the memory needed to create the message space.

**ATTRIBUTES**

See attributes(5) for descriptions of the following attributes. This value is returned, for example, when the system attempts to exceed the limit defined by the tunable kern.proc.maxOpenSpaceNumber.

ATTRIBUTE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

msgAllocate(2K), msgFree(2K), msgGet(2K), msgPut(2K), msgRemove(2K), msgSpaceOpen(2K)

<b>NAME</b>	msgSpaceOpen – open a message space				
<b>SYNOPSIS</b>	#include <mipc/chMipc.h> int msgSpaceOpen(KnMsgSpaceId spaceGid);				
<b>FEATURES</b>	MIPC				
<b>DESCRIPTION</b>	The msgSpaceOpen( ) call assigns a local private identifier associated with the message space opened to the calling actor. It maps the message pools of the message to the calling actor's addressing space.  The <i>spaceGid</i> field is the global identifier of a message space created previously using msgSpaceCreate(2K).  The msgSpaceOpen( ) call transparently selects the address within the addressing space of the actor at which the set of message pools of the message space are to be mapped.				
<b>RETURN VALUES</b>	Upon successful completion, msgSpaceOpen( ) returns the local identifier of the message space assigned to the actor. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	[K_EINVAL] The message space does not exist. [K_EOVERLAP] There is insufficient room in the addressing space of the actor to map the message pools of the opened message space. [K_ENOMEM] The system is out of resources.				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes: This value is returned, for example, when the system attempts to exceed the limit defined by the tunable <code>ipc.maxOpenSpaceNum</code> .				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	msgAllocate(2K), msgFree(2K), msgGet(2K), msgPut(2K), msgRemove(2K), msgSpaceCreate(2K)				

<b>NAME</b>	mutexInit, mutexGet, mutexRel, mutexTry – initialize a mutex; acquire a mutex; release a mutex; try to acquire a mutex
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chMutex.h&gt; int mutexInit(KnMutex * mutex);  int mutexGet(KnMutex * mutex);  int mutexRel(KnMutex * mutex);  int mutexTry(KnMutex * mutex);</pre>
<b>FEATURES</b>	MUTEX
<b>DESCRIPTION</b>	<p>Mutexes are binary semaphores used to protect shared data from concurrent access: a mutex is a two state-variable - free or locked.</p> <p>Mutexes are <i>KnMutex</i> structures allocated in the user memory.</p> <p><i>mutexInit</i> initializes the mutex the address of which is <i>mutex</i>. The mutex is initialized as free.</p> <p>Statically allocated mutexes can be initialized using the <i>K_KNMUTEX_INITIALIZER</i> macro, which initializes the mutex as free. This macro is used as follows:</p> <pre>KnMutex myMutex = K_KNMUTEX_INITIALIZER ;</pre> <p><i>mutexGet</i> is used to acquire a mutex. If the mutex is free, it becomes locked and the caller continues its execution normally. If the mutex is locked, the caller is blocked until the mutex is released.</p> <p><i>mutexRel</i> is used to release a mutex. If threads are blocked behind the mutex, one of them is awakened.</p> <p><i>mutexTry</i> is only an attempt to acquire the mutex: it has the same effect as <i>mutexGet</i>, except that if the mutex is locked, the thread is not blocked (a return code is provided).</p> <p>Semaphores initialized with a count value of 1 (see <i>semInit (2K)</i>) are functionally equivalent to mutexes. However, there are implementation differences between mutexes and semaphores which make them behave differently. First, mutex operations are implemented so that optimal performance is obtained when no thread state change is needed (i.e. when a thread acquires a free mutex, or when a mutex is released and no thread is waiting for it). A consequence of this optimization is that mutexes do not guarantee fairness.</p> <p>A blocking <i>mutexGet</i> is <i>NONABORTABLE</i> (see <i>threadAbort (2K)</i>).</p>
<b>RESTRICTIONS</b>	A user application and a supervisor application may not share a mutex.

On the contrary, two user applications may share a mutex by mapping it in both user address spaces. Such shared mutexes must be dynamically allocated mutexes.

**RETURN VALUE**

*mutexTry* returns 0 if the mutex was already locked, 1 if it was free. Upon successful completion of *mutexInit*, *mutexGet* and *mutexRel*, 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EFAULT]                      Some of the provided data are outside the current actor's address space.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*semInit(2K)*

<b>NAME</b>	mutexInit, mutexGet, mutexRel, mutexTry – initialize a mutex; acquire a mutex; release a mutex; try to acquire a mutex
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chMutex.h&gt; int mutexInit(KnMutex * mutex);  int mutexGet(KnMutex * mutex);  int mutexRel(KnMutex * mutex);  int mutexTry(KnMutex * mutex);</pre>
<b>FEATURES</b>	MUTEX
<b>DESCRIPTION</b>	<p>Mutexes are binary semaphores used to protect shared data from concurrent access: a mutex is a two state-variable - free or locked.</p> <p>Mutexes are <i>KnMutex</i> structures allocated in the user memory.</p> <p><i>mutexInit</i> initializes the mutex the address of which is <i>mutex</i>. The mutex is initialized as free.</p> <p>Statically allocated mutexes can be initialized using the <i>K_KNMUTEX_INITIALIZER</i> macro, which initializes the mutex as free. This macro is used as follows:</p> <pre>KnMutex myMutex = K_KNMUTEX_INITIALIZER ;</pre> <p><i>mutexGet</i> is used to acquire a mutex. If the mutex is free, it becomes locked and the caller continues its execution normally. If the mutex is locked, the caller is blocked until the mutex is released.</p> <p><i>mutexRel</i> is used to release a mutex. If threads are blocked behind the mutex, one of them is awakened.</p> <p><i>mutexTry</i> is only an attempt to acquire the mutex: it has the same effect as <i>mutexGet</i>, except that if the mutex is locked, the thread is not blocked (a return code is provided).</p> <p>Semaphores initialized with a count value of 1 (see <i>semInit (2K)</i>) are functionally equivalent to mutexes. However, there are implementation differences between mutexes and semaphores which make them behave differently. First, mutex operations are implemented so that optimal performance is obtained when no thread state change is needed (i.e. when a thread acquires a free mutex, or when a mutex is released and no thread is waiting for it). A consequence of this optimization is that mutexes do not guarantee fairness.</p> <p>A blocking <i>mutexGet</i> is <i>NONABORTABLE</i> (see <i>threadAbort (2K)</i>).</p>
<b>RESTRICTIONS</b>	A user application and a supervisor application may not share a mutex.

On the contrary, two user applications may share a mutex by mapping it in both user address spaces. Such shared mutexes must be dynamically allocated mutexes.

**RETURN VALUE**

*mutexTry* returns 0 if the mutex was already locked, 1 if it was free. Upon successful completion of *mutexInit*, *mutexGet* and *mutexRel*, 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EFAULT]                      Some of the provided data are outside the current actor's address space.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*semInit(2K)*

<b>NAME</b>	mutexInit, mutexGet, mutexRel, mutexTry – initialize a mutex; acquire a mutex; release a mutex; try to acquire a mutex
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chMutex.h&gt; int mutexInit(KnMutex * mutex);  int mutexGet(KnMutex * mutex);  int mutexRel(KnMutex * mutex);  int mutexTry(KnMutex * mutex);</pre>
<b>FEATURES</b>	MUTEX
<b>DESCRIPTION</b>	<p>Mutexes are binary semaphores used to protect shared data from concurrent access: a mutex is a two state-variable - free or locked.</p> <p>Mutexes are <i>KnMutex</i> structures allocated in the user memory.</p> <p><i>mutexInit</i> initializes the mutex the address of which is <i>mutex</i>. The mutex is initialized as free.</p> <p>Statically allocated mutexes can be initialized using the <i>K_KNMUTEX_INITIALIZER</i> macro, which initializes the mutex as free. This macro is used as follows:</p> <pre>KnMutex myMutex = K_KNMUTEX_INITIALIZER ;</pre> <p><i>mutexGet</i> is used to acquire a mutex. If the mutex is free, it becomes locked and the caller continues its execution normally. If the mutex is locked, the caller is blocked until the mutex is released.</p> <p><i>mutexRel</i> is used to release a mutex. If threads are blocked behind the mutex, one of them is awakened.</p> <p><i>mutexTry</i> is only an attempt to acquire the mutex: it has the same effect as <i>mutexGet</i>, except that if the mutex is locked, the thread is not blocked (a return code is provided).</p> <p>Semaphores initialized with a count value of 1 (see <i>semInit (2K)</i>) are functionally equivalent to mutexes. However, there are implementation differences between mutexes and semaphores which make them behave differently. First, mutex operations are implemented so that optimal performance is obtained when no thread state change is needed (i.e. when a thread acquires a free mutex, or when a mutex is released and no thread is waiting for it). A consequence of this optimization is that mutexes do not guarantee fairness.</p> <p>A blocking <i>mutexGet</i> is <i>NONABORTABLE</i> (see <i>threadAbort (2K)</i>).</p>
<b>RESTRICTIONS</b>	A user application and a supervisor application may not share a mutex.

On the contrary, two user applications may share a mutex by mapping it in both user address spaces. Such shared mutexes must be dynamically allocated mutexes.

**RETURN VALUE**

*mutexTry* returns 0 if the mutex was already locked, 1 if it was free. Upon successful completion of *mutexInit*, *mutexGet* and *mutexRel*, 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EFAULT]                      Some of the provided data are outside the current actor's address space.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*semInit(2K)*

<b>NAME</b>	mutexInit, mutexGet, mutexRel, mutexTry – initialize a mutex; acquire a mutex; release a mutex; try to acquire a mutex
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chMutex.h&gt; int mutexInit(KnMutex * mutex);  int mutexGet(KnMutex * mutex);  int mutexRel(KnMutex * mutex);  int mutexTry(KnMutex * mutex);</pre>
<b>FEATURES</b>	MUTEX
<b>DESCRIPTION</b>	<p>Mutexes are binary semaphores used to protect shared data from concurrent access: a mutex is a two state-variable - free or locked.</p> <p>Mutexes are <i>KnMutex</i> structures allocated in the user memory.</p> <p><i>mutexInit</i> initializes the mutex the address of which is <i>mutex</i>. The mutex is initialized as free.</p> <p>Statically allocated mutexes can be initialized using the <i>K_KNMUTEX_INITIALIZER</i> macro, which initializes the mutex as free. This macro is used as follows:</p> <pre>KnMutex myMutex = K_KNMUTEX_INITIALIZER ;</pre> <p><i>mutexGet</i> is used to acquire a mutex. If the mutex is free, it becomes locked and the caller continues its execution normally. If the mutex is locked, the caller is blocked until the mutex is released.</p> <p><i>mutexRel</i> is used to release a mutex. If threads are blocked behind the mutex, one of them is awakened.</p> <p><i>mutexTry</i> is only an attempt to acquire the mutex: it has the same effect as <i>mutexGet</i>, except that if the mutex is locked, the thread is not blocked (a return code is provided).</p> <p>Semaphores initialized with a count value of 1 (see <i>semInit (2K)</i>) are functionally equivalent to mutexes. However, there are implementation differences between mutexes and semaphores which make them behave differently. First, mutex operations are implemented so that optimal performance is obtained when no thread state change is needed (i.e. when a thread acquires a free mutex, or when a mutex is released and no thread is waiting for it). A consequence of this optimization is that mutexes do not guarantee fairness.</p> <p>A blocking <i>mutexGet</i> is <i>NONABORTABLE</i> (see <i>threadAbort (2K)</i>).</p>
<b>RESTRICTIONS</b>	A user application and a supervisor application may not share a mutex.

On the contrary, two user applications may share a mutex by mapping it in both user address spaces. Such shared mutexes must be dynamically allocated mutexes.

**RETURN VALUE**

*mutexTry* returns 0 if the mutex was already locked, 1 if it was free. Upon successful completion of *mutexInit*, *mutexGet* and *mutexRel*, 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EFAULT]                      Some of the provided data are outside the current actor's address space.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*semInit*(2K)

<b>NAME</b>	padGet – return actor-specific values associated with keys				
<b>SYNOPSIS</b>	<pre>#include &lt;pd/chPd.h&gt; void *padGet(KnCap *actorcap, PdKey key);</pre>				
<b>FEATURES</b>	PRIVATE-DATA				
<b>DESCRIPTION</b>	<p>The <i>padGet</i> function returns the value currently bound to <i>key</i> for the actor whose capability is given by <i>actorcap</i>.</p> <p>If <i>actorcap</i> is K_MYACTOR, the current actor is used.</p> <p>You can only call <i>padGet</i> from a supervisor actor, but <i>actorcap</i> can designate any local actor.</p> <p>Calling <i>padGet</i> with a <i>key</i> value not obtained from <i>padKeyCreate</i> or after <i>key</i> has been deleted using <i>padKeyDelete</i> can produce unpredictable results.</p>				
<b>RETURN VALUE</b>	The <i>padGet</i> function returns the actor-specific data value associated with the <i>key</i> specified for the relevant actor. If no actor-specific data value has been associated with <i>key</i> , a NULL value is returned.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>padKeyCreate(2K)</i> , <i>padKeyDelete(2K)</i> , <i>padSet(2K)</i>				

<b>NAME</b>	padKeyCreate – create a private key for an actor
<b>SYNOPSIS</b>	<pre>#include &lt;pd/chPd.h&gt; int padKeyCreate(PdKey *pkey, KnPdHdl destructor);</pre>
<b>FEATURES</b>	PRIVATE-DATA
<b>DESCRIPTION</b>	<p>The <i>padKeyCreate</i> function creates an actor-specific data key visible to all actors on a site.</p> <p>The key created is returned in <i>*pkey</i>.</p> <p>You can only call <i>padKeyCreate</i> from a supervisor actor.</p> <p>The keys provided by <i>padKeyCreate</i> are opaque indices used to locate actor-specific data. Although the same key may be used by different actors, the values bound to the key by <i>padSet</i> (2K) are maintained on a per-actor basis and persist for the life of the designated actor.</p> <p>Upon key creation, the value NULL is associated with the new key in all active actors. Upon actor creation, the value NULL is associated with all keys defined in the new actor.</p> <p>An optional <i>destructor</i> function may be associated with each key. At actor exit, if a key has a non-NULL destructor pointer, and the actor has a non-NULL value associated with that key, the function <code>destructor( )</code> is called with the current value as its sole argument.</p> <p>By default, the Private Data Manager invokes the destructor function through the LAP invocation mechanism (see <i>IsvLapCreate</i>(2K)) to ensure that the destructor function has the same execution actor than the thread which called <code>padKeyCreate( )</code> originally. However, when an actor associates a destructor function to itself, the destructor function cannot be invoked by LAP (as the actor is being deleted) and is executed in the context of the Private Data Manager. As a consequence, an actor self deletion handler should not use any API which assumes that the current execution actor is the actor being deleted.</p> <p>The order of destructor calls is unspecified if there is more than one destructor for an actor when it exits.</p> <p>Destructor calls are repeated for a maximum of <i>PD_DESTRUCTOR_ITERATIONS</i> (as defined in <i>pd/chPd.h</i>) times as necessary to reclaim all actor-specific data in the case that a destructor causes more actor-specific data to be created.</p>
<b>RETURN VALUE</b>	Upon successful completion, the newly created key is stored in <i>*pkey</i> and a value of 0 is returned. Otherwise, a positive error code is returned.
<b>ERRORS</b>	[PD_ENOKEY]                      The system-imposed limit on the total number of keys has been exceeded (see <i>PAD_KEYS_MAX</i> in <i>pd/chPd.h</i> ).

[PD\_ENOMEM]            There is insufficient memory to create the key.  
[PD\_ESERVER]           The Private Data Manager is unreachable.  
[PD\_EUSER]             This call is reserved for supervisor actors.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`padSet(2K)`, `padKeyDelete(2K)`

<b>NAME</b>	padKeyDelete – delete an actor private key				
<b>SYNOPSIS</b>	<pre>#include &lt;pd/chPd.h&gt; int padKeyDelete(PdKey key);</pre>				
<b>FEATURES</b>	PRIVATE-DATA				
<b>DESCRIPTION</b>	<p>The <i>padKeyDelete</i> function deletes an actor-specific data key previously returned by <i>padKeyCreate</i> (2K).</p> <p>This function can only be called from a supervisor actor.</p> <p>The actor-specific data values associated with <i>key</i> need not be NULL at the time <i>padKeyDelete</i> is called. The application frees storage and performs any cleanup operations needed for data structures related to the deleted key or to associated actor-specific data. The cleanup can be performed before or after invoking <i>padKeyDelete</i>. Any attempt to use <i>key</i> after the call to <i>padKeyDelete</i> can produce unpredictable results There are no destructor functions invoked by <i>padKeyDelete</i>.</p> <p>At actor destruction, all keys still allocated by an actor will be deleted.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a positive error code is returned.				
<b>ERRORS</b>	<p>[PD_EINVAL]                   The <i>key</i> value is invalid.</p> <p>[PD_ESERVER]                 The Private Data Manager is unreachable.</p> <p>[PD_EUSER]                   This call is reserved for supervisor actors.</p>				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>padKeyCreate</i> (2K)				

**NAME** padSet – set the actor’s key to a specific value

**SYNOPSIS** #include <pd/chPd.h>  
int padSet(KnCap \*actorcap, PdKey key, const void \*value);

**FEATURES** PRIVATE-DATA

**DESCRIPTION** The *padSet* function associates an actor-specific *value* with a *key* (obtained via a previous call to *padKeyCreate* (2K)) for the actor whose capability is given by *actorcap*.  
If *actorcap* is K\_MYACTOR, the current actor is used.  
You can only call *padSet* from a supervisor actor, but *actorcap* can designate any local actor.  
Different actors may bind different values to the same key. These values are typically pointers to blocks of dynamically allocated memory that have been reserved for use by a controlling actor.  
Calling *padSet* from a destructor may result in lost storage.

**RETURN VALUE** Upon successful completion a value of 0 is returned. Otherwise, a positive error code is returned.

**ERRORS**

[PD_EINVAL]	The <i>key</i> value is invalid, or <i>actorcap</i> is an inconsistent actor capability.
[PD_ENOMEM]	There is insufficient memory to associate the value with the key.
[PD_ESERVER]	The Private Data Manager is unreachable.
[PD_EUSER]	This call is reserved for SUPERVISOR actors.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *padKeyCreate*(2K)

<b>NAME</b>	portCreate, portDeclare – create a port; declare a port										
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int portCreate(KnCap * actorcap, KnUniqueId * portui); int portDeclare(KnCap * actorcap, KnUniqueId * portui);</pre>										
<b>FEATURES</b>	IPC										
<b>DESCRIPTION</b>	<p><i>portCreate</i> creates a port which is attached to the actor the capability of which (see <i>actorCreate</i> (2K)) is given by <i>actorcap</i>. If <i>actorcap</i> is K_MYACTOR, the port is attached to the current actor.</p> <p>A unique identifier is allocated to the port. This identifier is returned in the <i>KnUniqueId</i> structure given by <i>portui</i>.</p> <p>The port is created in the <i>DISABLED</i> state (see <i>portEnable</i> (2K)).</p> <p>The port message queue capacity is limited to K_CPORTQUEUE.</p> <p>The maximum number of ports that may be created on a site is limited to a value set at kernel generation.</p> <p><i>portDeclare</i> is equivalent to <i>portCreate</i>, except that the unique identifier of the port is given by the user (in <i>portui</i>) instead of being generated by the kernel. The unique identifier provided by the user must have been previously built using <i>uiBuild</i> (2K). As the user is responsible for insuring the unicity of the port identifier, <i>portDeclare</i> is restricted to <i>SUPERVISOR</i> threads (see <i>threadCreate</i> (2K)), or threads belonging to <i>SYSTEM</i> actors (see <i>actorCreate</i> (2K)).</p>										
<b>RETURN VALUE</b>	Upon successful completion, a positive local identifier of the port is returned. Otherwise, a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EINVAL]</td> <td><i>actorcap</i> is an inconsistent actor capability.</td> </tr> <tr> <td>[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td>[K_EFAULT]</td> <td>Some of the provided data are outside the current actor's address space.</td> </tr> <tr> <td>[K_EPRIV]</td> <td>The current thread is neither a supervisor thread nor a thread of a system actor (<i>portDeclare</i> only).</td> </tr> <tr> <td>[K_ENOMEM]</td> <td>The system is out of resources.</td> </tr> </table>	[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.	[K_EFAULT]	Some of the provided data are outside the current actor's address space.	[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor ( <i>portDeclare</i> only).	[K_ENOMEM]	The system is out of resources.
[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.										
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.										
[K_EFAULT]	Some of the provided data are outside the current actor's address space.										
[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor ( <i>portDeclare</i> only).										
[K_ENOMEM]	The system is out of resources.										
<b>RESTRICTIONS</b>	The target actor and the current actor must be located on the same site.										
<b>ATTRIBUTES</b>	See <i>attributes</i> (5) for descriptions of the following attributes:										

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

portDelete(2K), portEnable(2K), portDisable(2K), actorCreate(2K)

<b>NAME</b>	portCreate, portDeclare – create a port; declare a port										
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int portCreate(KnCap * actorcap, KnUniqueId * portui); int portDeclare(KnCap * actorcap, KnUniqueId * portui);</pre>										
<b>FEATURES</b>	IPC										
<b>DESCRIPTION</b>	<p><i>portCreate</i> creates a port which is attached to the actor the capability of which (see <i>actorCreate</i> (2K)) is given by <i>actorcap</i>. If <i>actorcap</i> is K_MYACTOR, the port is attached to the current actor.</p> <p>A unique identifier is allocated to the port. This identifier is returned in the <i>KnUniqueId</i> structure given by <i>portui</i>.</p> <p>The port is created in the <i>DISABLED</i> state (see <i>portEnable</i> (2K)).</p> <p>The port message queue capacity is limited to K_CPORTQUEUE.</p> <p>The maximum number of ports that may be created on a site is limited to a value set at kernel generation.</p> <p><i>portDeclare</i> is equivalent to <i>portCreate</i>, except that the unique identifier of the port is given by the user (in <i>portui</i>) instead of being generated by the kernel. The unique identifier provided by the user must have been previously built using <i>uiBuild</i> (2K). As the user is responsible for insuring the unicity of the port identifier, <i>portDeclare</i> is restricted to <i>SUPERVISOR</i> threads (see <i>threadCreate</i> (2K)), or threads belonging to <i>SYSTEM</i> actors (see <i>actorCreate</i> (2K)).</p>										
<b>RETURN VALUE</b>	Upon successful completion, a positive local identifier of the port is returned. Otherwise, a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EINVAL]</td> <td><i>actorcap</i> is an inconsistent actor capability.</td> </tr> <tr> <td>[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td>[K_EFAULT]</td> <td>Some of the provided data are outside the current actor's address space.</td> </tr> <tr> <td>[K_EPRIV]</td> <td>The current thread is neither a supervisor thread nor a thread of a system actor (<i>portDeclare</i> only).</td> </tr> <tr> <td>[K_ENOMEM]</td> <td>The system is out of resources.</td> </tr> </table>	[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.	[K_EFAULT]	Some of the provided data are outside the current actor's address space.	[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor ( <i>portDeclare</i> only).	[K_ENOMEM]	The system is out of resources.
[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.										
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.										
[K_EFAULT]	Some of the provided data are outside the current actor's address space.										
[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor ( <i>portDeclare</i> only).										
[K_ENOMEM]	The system is out of resources.										
<b>RESTRICTIONS</b>	The target actor and the current actor must be located on the same site.										
<b>ATTRIBUTES</b>	See <i>attributes</i> (5) for descriptions of the following attributes:										

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

portDelete(2K), portEnable(2K), portDisable(2K), actorCreate(2K)

<b>NAME</b>	portDelete – delete a port				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int portDelete(KnCap *actorcap, int portli);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p><i>portDelete</i> deletes the port the local identifier of which is <i>portli</i> in the actor the capability of which is specified by the <i>KnCap</i> structure given by <i>actorcap</i> (see <i>actorCreate(2K)</i>). If <i>actorcap</i> is <i>K_MYACTOR</i>, the current actor is used.</p> <p>When the port is deleted, the queued messages are deleted (the <i>ipcCall(2K)</i> transactions are aborted), and the waiting threads exit their waiting status with the <i>K_ENOPORT</i> error code.</p>				
<b>RETURN VALUE</b>	Upon successful completion, a value of 0 is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[<i>K_EINVAL</i>] <i>actorcap</i> is an inconsistent actor capability, or <i>portli</i> is not a valid port identifier within the named actor.</p> <p>[<i>K_EUNKNOWN</i>] <i>actorcap</i> does not specify a reachable actor.</p> <p>[<i>K_EFAULT</i>] Some of the provided data are outside the current actor's address space.</p>				
<b>RESTRICTIONS</b>	The port and the current actor must be located on the same site.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>portCreate(2K)</i> , <i>actorCreate(2K)</i> , <i>ipcCall(2K)</i>				

<b>NAME</b>	portEnable – Enable a port portDisable; Disable a port				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int portEnable(KnCap *actorcap, int portli, int priority);  int portDisable(KnCap *actorcap, int portli);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p>The <i>portEnable</i> utility is used to put ports into the <i>ENABLED</i> state. The port's local identifier is <i>portli</i> in the actor whose capability is specified by the <i>KnCap</i> structure given by <i>actorcap</i> (see <i>actorCreate(2K)</i>). If <i>actorcap</i> is <i>K_MYACTOR</i>, the current actor is used. The <i>priority</i> field defines the priority of the port in the set of <i>ENABLED</i> ports (see below).</p> <p>The <i>portDisable</i> utility puts the port in the <i>DISABLED</i> state.</p> <p>The state of a port determines whether or not the port is used if a thread performs a multiple receive call (<i>ipcReceive(2K)</i> with <i>K_ANYENABLED</i> option set). Only a message delivered to an <i>ENABLED</i> port may be received during a multiple receive call. When several ports in the set have messages queued behind them, the <i>priority</i> parameter given when the port was <i>ENABLED</i> is used (a low value means a high priority). During a multiple receive, the oldest message queued behind the highest priority <i>ENABLED</i> port is delivered. If none of the <i>ENABLED</i> ports have a queued message, the first message received by any of the <i>ENABLED</i> ports is delivered.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EINVAL] <i>portli</i> is not a valid port identifier within the current actor.</p> <p>[K_EUNKNOWN] <i>actorcap</i> does not specify a reachable actor.</p> <p>[K_EFAULT] Some of the data provided are outside the current actor's address space.</p>				
<b>RESTRICTIONS</b>	The port and the current actor must be located on the same site.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>portCreate(2K)</i> , <i>ipcReceive(2K)</i>				

<b>NAME</b>	portMigrate, portGetSeqNum – Migrate a port; Get a port sequence number				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int portMigrate(int options, KnCap * srcactorcap, int portli, KnCap * dstactorcap, KnEvtNum * seqnum);  int portGetSeqNum(KnCap * srcactorcap, int portli, KnEvtNum * seqnum);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p>The <i>portMigrate</i> utility forces the port whose local identifier is <i>portli</i> in the actor whose capability is given by <i>srcactorcap</i> to migrate to the actor whose capability is given by <i>dstactorcap</i>. After the call, the port is attached to this destination actor. The local identifier of the port is returned by <i>portMigrate</i> to the destination actor.</p> <p>If <i>srcactorcap</i> and/or <i>dstactorcap</i> is K_MYACTOR, the current actor is used as the source and/or destination of the migration.</p> <p>The <i>options</i>, which describe the way in which the migration is performed may take one of the following values:</p> <table border="0"> <tr> <td style="padding-right: 20px;">K_WITHMSGS</td> <td>The messages queued behind the port follow the port during the migration.</td> </tr> <tr> <td style="padding-right: 20px;">K_KILLMSGS</td> <td>The queued messages are deleted (and <i>ipcCall</i> (2K) transactions are aborted) before the migration.</td> </tr> </table> <p>The <i>seqnum</i> out parameter is a pointer to a count of messages that have been received on the source port, plus one additional increment for the migrate itself. The port message count is reset to 0 on the destination port.</p> <p>If the port was inserted into port groups (using <i>grpPortInsert</i> (2K)) before the migration, the port remains a member of the same groups after the migration.</p> <p>Ports that had been enabled (using <i>portEnable</i> (2K)) before the migration are disabled after the migration.</p> <p>When a port migrates, any thread waiting for a message on that port exits its waiting state and returns the K_ENOPORT error code.</p> <p>A port may not be migrated if a message handler (see <i>svMsgHandler</i> (2K)) is attached to it. If such an attempt is made, the K_EBUSYPORT error code is returned.</p> <p>The <i>portGetSeqNum</i> function gets the count of messages that have been received on the port whose local identifier is <i>portli</i> in the actor whose capability is given by <i>srcactorcap</i>. This count is stored by the kernel in the data pointed to by <i>seqnum</i>.</p>	K_WITHMSGS	The messages queued behind the port follow the port during the migration.	K_KILLMSGS	The queued messages are deleted (and <i>ipcCall</i> (2K) transactions are aborted) before the migration.
K_WITHMSGS	The messages queued behind the port follow the port during the migration.				
K_KILLMSGS	The queued messages are deleted (and <i>ipcCall</i> (2K) transactions are aborted) before the migration.				
<b>RETURN VALUE</b>	Upon successful completion, <i>portMigrate</i> returns the local identifier of the port in the destination actor. Otherwise, a negative error code is returned.				

<b>ERRORS</b>	[K_EINVAL]	<i>srcactorcap</i> or <i>dstactorcap</i> is an inconsistent actor capability or <i>portli</i> is not a port of the source actor.
	[K_EUNKNOWN]	<i>srcactorcap</i> or <i>dstactorcap</i> does not specify a reachable actor.
	[K_EFAULT]	Some of the data provided are outside the current actor's address space.
	[K_ENOMEM]	The system is out of resources.
	[K_EBUSYPORT]	Attempted to migrate a port to which a message handler is attached.

**RESTRICTIONS** Messages that are sent using *ipcSend* are not reliably forwarded to the new location of the port.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `portCreate(2K)`, `portEnable(2K)`, `grpPortInsert(2K)`, `ipcCall(2K)`, `svMsgHandler(2K)`

<b>NAME</b>	portUi, portLi – Get the unique identifier of a port, given its local identifier; Get the local identifier of a port, given its unique identifier				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int portUi(KnUniqueId * portui, int portli);  int portLi(KnUniqueId * portui);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p>The <i>portUi</i> call returns to the <i>KnUniqueId</i> structure given by <i>portui</i>, the unique identifier of the port whose local identifier is <i>portli</i> in the current actor (the port must be attached to the current actor).</p> <p>The <i>portLi</i> call returns the local identifier in the current actor of the port whose unique identifier is pointed by <i>portui</i> (the port must be attached to the current actor).</p>				
<b>RETURN VALUE</b>	Upon successful completion, <i>portui</i> returns 0; <i>portLi</i> returns the local identifier of the port. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EINVAL] For <i>portUi</i>: <i>portli</i> is not a valid port identifier within the current actor. For <i>portLi</i>: <i>portui</i> is not the unique identifier of a port attached to the current actor.</p> <p>[K_EFAULT] Some of the data provided are outside the current actor's address space.</p>				
<b>ATTRIBUTES</b>	See <a href="#">attributes(5)</a> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>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				
<b>SEE ALSO</b>	<a href="#">portCreate(2K)</a>				

<b>NAME</b>	portMigrate, portGetSeqNum – Migrate a port; Get a port sequence number				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int portMigrate(int options, KnCap * srcactorcap, int portli, KnCap * dstactorcap, KnEvtNum * seqnum);  int portGetSeqNum(KnCap * srcactorcap, int portli, KnEvtNum * seqnum);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p>The <i>portMigrate</i> utility forces the port whose local identifier is <i>portli</i> in the actor whose capability is given by <i>srcactorcap</i> to migrate to the actor whose capability is given by <i>dstactorcap</i>. After the call, the port is attached to this destination actor. The local identifier of the port is returned by <i>portMigrate</i> to the destination actor.</p> <p>If <i>srcactorcap</i> and/or <i>dstactorcap</i> is K_MYACTOR, the current actor is used as the source and/or destination of the migration.</p> <p>The <i>options</i>, which describe the way in which the migration is performed may take one of the following values:</p> <table border="0"> <tr> <td style="padding-right: 20px;">K_WITHMSGS</td> <td>The messages queued behind the port follow the port during the migration.</td> </tr> <tr> <td style="padding-right: 20px;">K_KILLMSGS</td> <td>The queued messages are deleted (and <i>ipcCall</i> (2K) transactions are aborted) before the migration.</td> </tr> </table> <p>The <i>seqnum</i> out parameter is a pointer to a count of messages that have been received on the source port, plus one additional increment for the migrate itself. The port message count is reset to 0 on the destination port.</p> <p>If the port was inserted into port groups (using <i>grpPortInsert</i> (2K)) before the migration, the port remains a member of the same groups after the migration.</p> <p>Ports that had been enabled (using <i>portEnable</i> (2K)) before the migration are disabled after the migration.</p> <p>When a port migrates, any thread waiting for a message on that port exits its waiting state and returns the K_ENOPORT error code.</p> <p>A port may not be migrated if a message handler (see <i>svMsgHandler</i> (2K)) is attached to it. If such an attempt is made, the K_EBUSYPORT error code is returned.</p> <p>The <i>portGetSeqNum</i> function gets the count of messages that have been received on the port whose local identifier is <i>portli</i> in the actor whose capability is given by <i>srcactorcap</i>. This count is stored by the kernel in the data pointed to by <i>seqnum</i>.</p>	K_WITHMSGS	The messages queued behind the port follow the port during the migration.	K_KILLMSGS	The queued messages are deleted (and <i>ipcCall</i> (2K) transactions are aborted) before the migration.
K_WITHMSGS	The messages queued behind the port follow the port during the migration.				
K_KILLMSGS	The queued messages are deleted (and <i>ipcCall</i> (2K) transactions are aborted) before the migration.				
<b>RETURN VALUE</b>	Upon successful completion, <i>portMigrate</i> returns the local identifier of the port in the destination actor. Otherwise, a negative error code is returned.				

**ERRORS**

[K_EINVAL]	<i>srcactorcap</i> or <i>dstactorcap</i> is an inconsistent actor capability or <i>portli</i> is not a port of the source actor.
[K_EUNKNOWN]	<i>srcactorcap</i> or <i>dstactorcap</i> does not specify a reachable actor.
[K_EFAULT]	Some of the data provided are outside the current actor's address space.
[K_ENOMEM]	The system is out of resources.
[K_EBUSYPORT]	Attempted to migrate a port to which a message handler is attached.

**RESTRICTIONS**

Messages that are sent using *ipcSend* are not reliably forwarded to the new location of the port.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*portCreate(2K)* , *portEnable(2K)* , *grpPortInsert(2K)* , *ipcCall(2K)* , *svMsgHandler(2K)*

<b>NAME</b>	portPi – get and/or set the protection identifier of a port										
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int portPi(KnCap *actorcap, int portli, KnProtId *oldpi, KnProtId *newpi);</pre>										
<b>FEATURES</b>	IPC										
<b>DESCRIPTION</b>	<p>The <i>portPi</i> call gets and/or sets the protection identifier of the port whose local identifier is <i>portli</i> in the actor whose capability is given by <i>actorcap</i> (see <i>actorCreate(2K)</i>).</p> <p>If <i>actorcap</i> is <i>K_MYACTOR</i>, the operation is applied to the current actor.</p> <p>If <i>portli</i> is <i>K_DEFAULTPORT</i>, the actor's default port is used.</p> <p>The <i>oldpi</i> and <i>newpi</i> fields are pointers to <i>KnProtId</i> structures whose members are the following:</p> <pre>unsigned short uid ; unsigned short gid ; unsigned short did ;</pre> <p>The user identifier is <i>uid</i>, the group identifier is <i>gid</i>, and the domain identifier is <i>did</i>.</p> <p>The corresponding port's protection identifier is copied into the structure pointed to by <i>oldpi</i> in the client address space. The new corresponding context is copied from the structure pointed to by <i>newpi</i>.</p> <p>If <i>oldpi</i> is a NULL pointer, the port's protection identifier is not copied to the caller's address space. If <i>newpi</i> is a NULL pointer, the port's protection identifier is not modified. The caller thread must be a <i>SUPERVISOR</i> thread (see <i>threadCreate(2K)</i>) or must belong to a <i>SYSTEM</i> actor (see <i>actorCreate(2K)</i>).</p>										
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the data provided are outside the current actor's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td><i>actorcap</i> is an inconsistent actor capability.</td> </tr> <tr> <td style="vertical-align: top;">[K_EPRIV]</td> <td>The current thread is neither a supervisor thread nor a thread of a system actor.</td> </tr> <tr> <td style="vertical-align: top;">[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td style="vertical-align: top;">[K_ENOPORT]</td> <td><i>portli</i> is not a valid port identifier in the actor.</td> </tr> </table>	[K_EFAULT]	Some of the data provided are outside the current actor's address space.	[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.	[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.	[K_ENOPORT]	<i>portli</i> is not a valid port identifier in the actor.
[K_EFAULT]	Some of the data provided are outside the current actor's address space.										
[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.										
[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor.										
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.										
[K_ENOPORT]	<i>portli</i> is not a valid port identifier in the actor.										
<b>RESTRICTIONS</b>	The target actor and the current actor must be located on the same site.										
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:										

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

portCreate(2K), threadCreate(2K), actorCreate(2K)

<b>NAME</b>	portUi, portLi – Get the unique identifier of a port, given its local identifier; Get the local identifier of a port, given its unique identifier				
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int portUi(KnUniqueId * portui, int portli);  int portLi(KnUniqueId * portui);</pre>				
<b>FEATURES</b>	IPC				
<b>DESCRIPTION</b>	<p>The <i>portUi</i> call returns to the <i>KnUniqueId</i> structure given by <i>portui</i>, the unique identifier of the port whose local identifier is <i>portli</i> in the current actor (the port must be attached to the current actor).</p> <p>The <i>portLi</i> call returns the local identifier in the current actor of the port whose unique identifier is pointed by <i>portui</i> (the port must be attached to the current actor).</p>				
<b>RETURN VALUE</b>	Upon successful completion, <i>portui</i> returns 0; <i>portLi</i> returns the local identifier of the port. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EINVAL] For <i>portUi</i>: <i>portli</i> is not a valid port identifier within the current actor. For <i>portLi</i>: <i>portui</i> is not the unique identifier of a port attached to the current actor.</p> <p>[K_EFAULT] Some of the data provided are outside the current actor's address space.</p>				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>portCreate(2K)</i>				

<b>NAME</b>	ptdErrnoAddr – return a thread-specific errno address
<b>SYNOPSIS</b>	<pre>#include &lt;pd/chPd.h&gt; int *ptdErrnoAddr(void);</pre>
<b>FEATURES</b>	PRIVATE-DATA
<b>DESCRIPTION</b>	<p>The <i>ptdErrnoAddr</i> call returns the address of the private errno word of the calling thread.</p> <p>It may be called either explicitly or implicitly from a thread-specific data destructor function.</p> <p>This function is implemented as a macro.</p> <p>A typical usage (in a multi-threaded POSIX environment) of <i>ptdErrnoAddr</i> is to define <i>errno</i> as (<i>*ptdErrnoAddr</i>()).</p>
<b>RETURN VALUE</b>	The function <i>ptdErrnoAddr</i> returns the thread-specific errno address.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**NAME** ptdGet – return thread-specific value associated with key

**SYNOPSIS** #include <pd/chPd.h>  
void\* ptdGet(PdKey key);

**FEATURES** PRIVATE-DATA

**DESCRIPTION** *ptdGet* returns the value currently bound to *key* in the calling thread.  
The effect of calling *ptdGet* with a *key* value not obtained from *ptdKeyCreate* or after *key* has been deleted with *ptdKeyDelete* is undefined.  
*ptdGet* may be called either explicitly or implicitly from a thread-specific data destructor function.  
This function is implemented as a macro.

**RETURN VALUE** The function *ptdGet* returns the thread-specific data value associated with the given *key*. If no thread-specific data value has been associated with *key*, then the value NULL is returned.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** ptdKeyCreate(2K), ptdKeyDelete(2K), ptdSet(2K)

<b>NAME</b>	ptdKeyCreate – create a thread-specific data key				
<b>SYNOPSIS</b>	<pre>#include &lt;pd/chPd.h&gt; int ptdKeyCreate(PdKey *pkey, KnPdHdl destructor);</pre>				
<b>FEATURES</b>	PRIVATE-DATA				
<b>DESCRIPTION</b>	<p><i>ptdKeyCreate</i> creates a thread-specific data key visible to all threads in the actor. The created key is returned in <i>*pkey</i>.</p> <p>Keys provided by <i>ptdKeyCreate</i> are opaque indices used to locate thread-specific data. Although the same key may be used by different threads, the values bound to the key by <i>ptdSet</i> (2K) are maintained on a per-thread basis, and persist through the life of the calling thread.</p> <p>Upon key creation, the value NULL shall be associated with the new key in all active threads. Upon thread creation, the value NULL shall be associated with all defined keys in the new thread.</p> <p>An optional destructor function <i>destructor</i> may be associated with each key. At thread exit, if a key has a non-NULL destructor pointer, and the thread has a non-NULL value associated with that key, the function that is pointed to is called with the current value as its sole argument. The order of destructor calls is unspecified if more than one destructor exists for a thread when it exits.</p> <p>Destructor calls are repeated at as most <i>PD_DESTRUCTOR_ITERATIONS</i> (as defined in <i>pd/chPd.h</i>) times as necessary to reclaim all thread-specific data, in the case a destructor causes more thread-specific data to be created.</p> <p><i>destructor</i> will be called by <i>ptdThreadDelete</i>(2K) in user mode or by the Private Data Manager at thread termination in supervisor mode. Calling <i>ptdThreadDelete</i>(2K) in supervisor mode is useless (and harmless), since it is defined as an empty function.</p>				
<b>RETURN VALUE</b>	Upon successful completion, the newly created key is stored in <i>*pkey</i> , and a value of 0 is returned. Otherwise, a positive error code is returned.				
<b>ERRORS</b>	<p>[PD_ENOKEY]                   The system-imposed limit on the total number of keys per actor has been exceeded (see <i>PTD_KEYS_MAX</i> in <i>pd/chPd.h</i>).</p> <p>[PD_ENOMEM]                 Insufficient memory exists to create the key.</p> <p>[PD_ESERVER]                The Private Data Manager is unreachable.</p>				
<b>ATTRIBUTES</b>	See <i>attributes</i> (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>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**

ptdSet(2K), ptdKeyDelete(2K)

<b>NAME</b>	ptdKeyDelete – delete a thread-specific data key				
<b>SYNOPSIS</b>	<pre>#include &lt;pd/chPd.h&gt; int ptdKeyDelete(PdKey key);</pre>				
<b>FEATURES</b>	PRIVATE-DATA				
<b>DESCRIPTION</b>	<p>The <i>ptdKeyDelete</i> call deletes a thread-specific data key previously returned by <i>ptdKeyCreate (2K)</i>. The thread-specific data values associated with <i>key</i> need not be NULL at the time <i>ptdKeyDelete</i> is called. The application frees storage and performs any cleanup operations needed for data structures related to the deleted key or to associated thread-specific data. The cleanup can be performed before or after invoking <i>ptdKeyDelete</i>. Any attempt to use <i>key</i> after the call to <i>ptdKeyDelete</i> can produce unpredictable results.. There are no destructor functions invoked by <i>ptdKeyDelete</i>.</p> <p>At actor destruction, all keys still allocated by an actor will silently be deleted.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a positive error code is returned.				
<b>ERRORS</b>	<p>[PD_EINVAL]                   The <i>key</i> value is invalid.</p> <p>[PD_ESERVER]                 The Private Data Manager is unreachable.</p>				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>ptdKeyCreate(2K)</i>				

**NAME** ptdRemoteGet – return a thread-specific data value for another thread

**SYNOPSIS** #include <pd/chPd.h>  
 int ptdRemoteGet(KnCap \*actorcap, int threadli, PdKey key, void \*\*pvalue);

**FEATURES** PRIVATE-DATA

**DESCRIPTION** The *ptdRemoteGet* call gets the value currently bound to the specific *key* for the thread whose local identifier is *threadli* in the actor whose capability is given by *actorcap*.

If *actorcap* is K\_MYACTOR, the target thread must belong to the current actor. In this case, if *threadli* is K\_MYSELF, the current thread is used.

The value is returned to *\*pvalue*. If no thread-specific data value is associated with *key*, then the value NULL is returned to *\*pvalue*.

In user mode, *actorcap* must represent the current actor.

**RETURN VALUE** Upon successful completion a value of 0 is returned. Otherwise, a positive error code is returned.

**ERRORS** [PD\_EINVAL] The *key* value is invalid, or *threadli* does not belong to the target actor, or *actorcap* is an inconsistent actor capability.

[PD\_ESERVER] The Private Data Manager is unreachable.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** ptdKeyCreate(2K), ptdGet(2K), ptdRemoteSet(2K)

<b>NAME</b>	ptdRemoteSet – set a thread-specific data value for another thread					
<b>SYNOPSIS</b>	<pre>#include &lt;pd/chPd.h&gt; int ptdRemoteSet(KnCap *actorcap, int threadli, PdKey key, const void *value);</pre>					
<b>FEATURES</b>	PRIVATE-DATA					
<b>DESCRIPTION</b>	<p>The <i>ptdRemoteSet</i> call associates a thread-specific <i>value</i> with a <i>key</i> (obtained via a previous call to <i>ptdKeyCreate</i> (2K)) for the thread whose local identifier is <i>threadli</i> in the actor whose capability is given by <i>actorcap</i>.</p> <p>If <i>actorcap</i> is K_MYACTOR, the thread used must belong to the current actor. In this case, if <i>threadli</i> is K_MYSELF, the current thread is used.</p> <p>Different threads may bind different values to the same key. These values are typically pointers to blocks of dynamically allocated memory that have been reserved for use by the calling thread.</p> <p>In user mode, <i>actorcap</i> must represent the current actor.</p>					
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a positive error code is returned.					
<b>ERRORS</b>	[PD_EINVAL]	The <i>key</i> value is invalid, or <i>threadli</i> does not belong to the target actor, or <i>actorcap</i> is an inconsistent actor capability.				
	[PD_ENOMEM]	There is insufficient memory to associate the value with the key.				
	[PD_ESERVER]	The Private Data Manager is unreachable.				
<b>RESTRICTIONS</b>	The target actor and the current actor must be located on the same site.					
<b>ATTRIBUTES</b>	See <i>attributes</i> (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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Interface Stability	Evolving					
<b>SEE ALSO</b>	ptdKeyCreate(2K), ptdSet(2K)					

**NAME** | ptdSet – set a thread-specific value

**SYNOPSIS** | #include <pd/chPd.h>  
 int ptdset(PdKey key, const void \*value);

**FEATURES** | PRIVATE-DATA

**DESCRIPTION** | *ptdSet* associates a thread-specific *value* with a *key* obtained via a previous call to *ptdKeyCreate* (2K). Different threads may bind different values to the same key. These values are typically pointers to blocks of dynamically allocated memory that have been reserved for use by the calling thread.

Calling *ptdSet* from a destructor may result in lost storage.

**RETURN VALUE** | Upon successful completion, a value of 0 is returned. Otherwise, a positive error code is returned.

**ERRORS** | [PD\_EINVAL]                      The *key* value is invalid.  
 [PD\_ENOMEM]                      Insufficient memory exists to associate the value with the key.  
 [PD\_ESERVER]                      The Private Data Manager is unreachable.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | ptdKeyCreate(2K)

<b>NAME</b>	ptdThreadDelete – delete all thread-specific values and call destructors				
<b>SYNOPSIS</b>	<pre>#include &lt;pd/chPd.h&gt; void ptdThreadDelete(void);</pre>				
<b>FEATURES</b>	PRIVATE-DATA				
<b>DESCRIPTION</b>	<p>The <i>ptdThreadDelete</i> call deletes the thread-specific values associated with all valid keys in the current user actor for the current thread.</p> <p>For each non-NULL value, <i>ptdThreadDelete</i> calls the appropriate destructor before clearing the value.</p> <p>In user mode, <i>ptdThreadDelete</i> should be called before thread deletion to enable destructors to be invoked.</p> <p>In supervisor mode, <i>ptdThreadDelete</i> is an empty function. Destructors will be called at thread termination by the Private Data Manager.</p>				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>ptdKeyCreate(2K)</code> , <code>ptdSet(2K)</code>				

**NAME** ptdThreadId – return the thread ID

**SYNOPSIS** #include <pd/chPd.h>  
int ptdThreadId(void);

**FEATURES** PRIVATE-DATA

**DESCRIPTION** The *ptdThreadId* function returns the unique thread ID for the calling thread. The thread ID is equal to the thread local identifier as returned by *threadSelf(2K)*.  
It may be called either explicitly or implicitly from a thread-specific data destructor function.  
This function is implemented as a macro.

**RETURN VALUE** The *ptdThreadId* function returns the thread-specific unique ID.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** threadSelf(2K)

<b>NAME</b>	rgnAllocate – allocate a region in an actor address space				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnAllocate(KnCap *actorcap, KnRgnDesc *Rgndesc);</pre>				
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL				
<b>DESCRIPTION</b>	<p><i>rgnAllocate</i> creates a region in the target actor address space and maps some newly allocated volatile memory to the region.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is <code>K_MYACTOR</code>, the address space of the current actor is used. If <i>actorcap</i> is <code>K_SVACTOR</code>, the region is allocated in the supervisor address space and is not attached to any particular supervisor actor; therefore, it cannot be implicitly deallocated by an <i>actorDelete(2K)</i> of a supervisor actor.</p> <p><i>rgndesc</i> points to a <i>KnRgnDesc</i> structure which contains the specification of the required region, and the members of which are the following:</p> <pre>VmAddr  startAddr ; VmSize  size ; VmFlags options ; VmAddr  endAddr ; void*   opaque1 ; VmFlags opaque2 ;</pre> <p>If neither <code>K_ANYWHERE</code> nor <code>K_RESTRICTIVE</code> options (see below) are specified, <i>startAddr</i> is the region start address in the target address space. The region start address must be aligned with a virtual page boundary (see <i>vmPageSize(2K)</i>).</p> <p><i>size</i> is the size of the region in bytes. If the size is not page-aligned, the kernel rounds it up to the next virtual page boundary.</p> <p><i>endAddr</i> is only used with the <code>K_RESTRICTIVE</code> option (see below).</p> <p><i>options</i> is a combination of following flags:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>K_ANYWHERE</code></td> <td>The input value of <i>startAddr</i> is ignored, and the region is allocated wherever possible in the limits of the target address space as returned by <i>vmStat(2K)</i>. In this case the region's start address is returned in <i>startAddr</i>.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>K_RESTRICTIVE</code></td> <td>The region is allocated wherever possible between <i>startAddr</i> and <i>endAddr</i>. If <i>startAddr</i> is not page-aligned, the kernel rounds it up to the next virtual page boundary. If <i>endAddr</i> is not page-aligned the kernel rounds it down to the preceding virtual page boundary. The created</td> </tr> </table>	<code>K_ANYWHERE</code>	The input value of <i>startAddr</i> is ignored, and the region is allocated wherever possible in the limits of the target address space as returned by <i>vmStat(2K)</i> . In this case the region's start address is returned in <i>startAddr</i> .	<code>K_RESTRICTIVE</code>	The region is allocated wherever possible between <i>startAddr</i> and <i>endAddr</i> . If <i>startAddr</i> is not page-aligned, the kernel rounds it up to the next virtual page boundary. If <i>endAddr</i> is not page-aligned the kernel rounds it down to the preceding virtual page boundary. The created
<code>K_ANYWHERE</code>	The input value of <i>startAddr</i> is ignored, and the region is allocated wherever possible in the limits of the target address space as returned by <i>vmStat(2K)</i> . In this case the region's start address is returned in <i>startAddr</i> .				
<code>K_RESTRICTIVE</code>	The region is allocated wherever possible between <i>startAddr</i> and <i>endAddr</i> . If <i>startAddr</i> is not page-aligned, the kernel rounds it up to the next virtual page boundary. If <i>endAddr</i> is not page-aligned the kernel rounds it down to the preceding virtual page boundary. The created				

	region start address is returned in <i>startAddr</i> . If this flag is set, the K_ANYWHERE flag is ignored.
K_INHERITCOPY	If a <i>rgnDup(2K)</i> is subsequently applied, taking the target actor as the source, a copy of the region will be allocated in the destination address space by the <i>rgnDup</i> operation.
K_INHERITSHARE	If a <i>rgnDup(2K)</i> is subsequently applied, taking the target actor as the source, the region will be made shared by the <i>rgnDup</i> operation.
K_FILLZERO	The memory eventually allocated for the region is zero-filled. Otherwise, the initial state of the memory is undefined.
K_NODEMAND	Guarantees that no page faults occur during subsequent (legal) access to data within the region.
K_NOWAITFORMEMORY	Do not wait for memory to become available. If the K_NOWAITFORMEMORY and K_NODEMAND flags are set, the <i>rgnAllocate</i> fails when there is not enough physical memory, and returns the K_ENOMEM error. If only K_NOWAITFORMEMORY is set, possible memory underflow during a future on demand allocation provokes a memory underflow exception. Otherwise memory underflow blocks and waits for memory to become available.
K_NOSWAPOUT	Guarantees that no pageouts occur for the data within the region.
K_RESERVED	The region is a reserved region. If K_RESERVED option is set all other options, except of K_ANYWHERE or K_RESTRICTIVE ones, are ignored. Virtual-to-physical mappings (attributes as well as addresses) of a reserved region are managed from the kernel using <i>svPhysMap(9DKI)</i> call. A thread can access to any effectively mapped (valid) address within a reserved region, but any fault provokes an exception processing (see <i>svExcHandler(2K)</i> ). Mapped portions of reserved regions can also be used as data buffers for IPC messages, <i>vmCopy</i> , <i>sgRead</i> , etc. A reserved region can be deallocated by a

	regular <i>rgnFree</i> or <i>actorDelete</i> operation. Any other kernel operation performed on a reserved region either has no effect (eg. <i>rgnSetProtect(2K)</i> ), or returns K_ENOTIMP (eg. <i>vmFlush(2K)</i> or <i>rgnInitFromActor(2K)</i> ).
K_STACK	Optimize future region extension toward low addresses (by default extension toward high addresses is optimized).
K_READABLE	Read access is permitted. The current implementation always permits read access for backward compatibility reasons.
K_WRITABLE	Write access is permitted.
K_EXECUTABLE	Execute access is permitted.
K_SUPERVISOR	The region is accessible only by supervisor threads.
K_NODUMP	The region will not be dumped if the target actor is core dumped.

Any access to a region conflicting with the region protections can potentially be reported by an exception (see *svExcHandler(2K)*) but the guarantees of such exceptions are machine-dependent. For instance, on any available paged memory hardware the kernel guarantees that a user-level write access to a read only region provokes a protection exception, but only on some of them the kernel guarantees to detect illegal access when a thread executes a region without K\_EXECUTABLE attribute.

All flags specified in the *options* field as well as values specified by *opaque1* and *opaque2* fields compose region attributes. The attributes are stored in the kernel state associated with the region and can be consulted using *rgnStat(2K)* kernel call.

**RETURN VALUE**

If successful K\_OK, is returned. Otherwise a negative error code is returned.

**ERRORS**

[K_EFAULT]	Some of the provided arguments are outside the caller's address space.
[K_EINVAL]	An inconsistent actor capability was given.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EROUND]	<i>rgndesc-&gt;startAddr</i> is not page-aligned.

[K_ESPACE]	Try to allocate a region outside the valid range for the address space of an actor as returned by <i>vmStat(2K)</i> .
[K_ESPACE]	<i>rgndesc-&gt;size</i> is 0.
[K_EOVERLAP]	The K_ANYWHERE option was specified and there is not enough room available in the address space to allocate the region.
[K_EOVERLAP]	The K_RESTRICTIVE option was specified and there is not enough room in the target address range.
[K_EOVERLAP]	The allocated region overlaps an existing region.
[K_ENOMEM]	The system is out of resources.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

The K\_RESTRICTIVE option is valid only in the MEM\_PROTECTED and MEM\_VIRTUAL modules. This option is ignored by the MEM\_FLAT module.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnDup(2K)*, *rgnFree(2K)*, *rgnSetProtect(2K)*, *rgnStat(2K)*, *svExcHandler(2K)*

<b>NAME</b>	rgnDup – duplicate an actor address space												
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnDup(KnCap *tgtactorcap, KnCap *srcactorcap, VmFlags flags);</pre>												
<b>FEATURES</b>	MEM_VIRTUAL												
<b>DESCRIPTION</b>	<p>The <i>rgnDup</i> function duplicates regions from the address space specified by <i>srcactorcap</i> into the actor address space specified by <i>tgtactorcap</i>. If <i>srcactorcap</i> and/or <i>tgtactorcap</i> is K_MYACTOR, the address space of the current actor is used as the source and/or the destination, respectively.</p> <p>For each region from the source actor <i>rgnDup</i> duplicates the region in the target actor address space with the same address, the same size and the same attributes (options).</p> <p>The duplicate region can either share the data with the source region (map to the same segment of volatile memory as the source), or take a private volatile copy. The data is shared if the <i>flags</i> argument is K_INHERITSHARE or if <i>flags</i> is zero and the source region has the K_INHERITSHARE attribute. The data is copied if the <i>flags</i> argument is K_INHERITCOPY or if <i>flags</i> is zero and the source region has the K_INHERITCOPY attribute.</p> <p>The region isn't duplicated if the <i>flag</i> argument is zero and the region has either the K_INHERITSHARE nor the K_INHERITCOPY attribute.</p> <p>Before the operation, the target actor address space must be empty.</p>												
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.												
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the provided arguments are outside the caller's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td>An inconsistent actor capability.</td> </tr> <tr> <td style="vertical-align: top;">[K_EUNKNOWN]</td> <td><i>tgtactorcap</i> or <i>srcactorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td style="vertical-align: top;">[K_ENOEMPTY]</td> <td>The target actor address space is not empty.</td> </tr> <tr> <td style="vertical-align: top;">[K_ENOMEM]</td> <td>The system is out of resources.</td> </tr> <tr> <td style="vertical-align: top;">[K_EMAPPER]</td> <td>The segment mapper doesn't respect the kernel/mapper protocol.</td> </tr> </table>	[K_EFAULT]	Some of the provided arguments are outside the caller's address space.	[K_EINVAL]	An inconsistent actor capability.	[K_EUNKNOWN]	<i>tgtactorcap</i> or <i>srcactorcap</i> does not specify a reachable actor.	[K_ENOEMPTY]	The target actor address space is not empty.	[K_ENOMEM]	The system is out of resources.	[K_EMAPPER]	The segment mapper doesn't respect the kernel/mapper protocol.
[K_EFAULT]	Some of the provided arguments are outside the caller's address space.												
[K_EINVAL]	An inconsistent actor capability.												
[K_EUNKNOWN]	<i>tgtactorcap</i> or <i>srcactorcap</i> does not specify a reachable actor.												
[K_ENOEMPTY]	The target actor address space is not empty.												
[K_ENOMEM]	The system is out of resources.												
[K_EMAPPER]	The segment mapper doesn't respect the kernel/mapper protocol.												
<b>RESTRICTIONS</b>	The target actor, the source actor and the current actor must be located on the same site.												
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:												

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

rgnAllocate(2K), rgnInit(2K), rgnMap(2K), rgnSetInherit(2K),  
rgnSetProtect(2K), rgnStat(2K)

<b>NAME</b>	rgnFree – deallocate regions of an actor address space												
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnFree(KnCap *actorcap, KnRgnDes *rgndesc);</pre>												
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL												
<b>DESCRIPTION</b>	<p><i>rgnFree</i> deallocates all regions from a given address range of the actor address space specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is <code>K_MYACTOR</code>, the address space of the current actor is used. If <i>actorcap</i> is <code>K_SVACTOR</code>, the supervisor address space is used.</p> <p><i>rgndesc</i> points to a <i>KnRgnDesc</i> structure described in <i>rgnAllocate (2K)</i>. If <i>options</i> is equal to <code>K_FREEALL</code>, all the actor's allocated regions are deallocated. Otherwise, <i>option</i> must be set to 0, and in that case the <i>startAddr</i> field specifies the starting address, and <i>size</i> specifies the size of the address range to be deallocated. The <i>endAddr</i> field is ignored.</p> <p>If <i>actorcap</i> is equal to <code>K_SVACTOR</code>, the <code>K_FREEALL</code> flag cannot be used.</p> <p><i>rgndesc-&gt;startAddr</i> must be aligned with a page boundary (see <i>vmPageSize(2K)</i>). If <i>rgndesc-&gt;size</i> is not page-aligned, the kernel rounds it up to the next page boundary.</p> <p>Since <i>rgndesc-&gt;startAddr</i> and <i>rgndesc-&gt;startAddr + rgnFree-&gt;size</i> do not necessarily match the boundaries of a region, a <i>rgnFree</i> can shrink the address ranges of some regions and even, in the case of a <i>rgnFree</i> in the middle of a region, create a new one. It explains why the <code>K_ENOMEM</code> error can be returned by a <i>rgnFree</i>.</p> <p>Besides, the deallocated address range can contain holes, i.e. sub-ranges without regions mapped to.</p>												
<b>RETURN VALUE</b>	If successful, <code>K_OK</code> is returned. Otherwise a negative error code is returned.												
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;"><code>[K_EFAULT]</code></td> <td>Some of the provided arguments are outside the caller's address space.</td> </tr> <tr> <td style="vertical-align: top;"><code>[K_EINVAL]</code></td> <td>An inconsistent actor capability was given. The value of <i>option</i> is not <code>K_FREEALL</code> or 0.</td> </tr> <tr> <td style="vertical-align: top;"><code>[K_EUNKNOWN]</code></td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td style="vertical-align: top;"><code>[K_EROUND]</code></td> <td><i>rgndesc-&gt;startAddr</i> is not page-aligned.</td> </tr> <tr> <td style="vertical-align: top;"><code>[K_ENOMEM]</code></td> <td>The system is out of resources.</td> </tr> <tr> <td style="vertical-align: top;"><code>[K_EMAPPER]</code></td> <td>The segment mapper does not respect the kernel/mapper protocol</td> </tr> </table>	<code>[K_EFAULT]</code>	Some of the provided arguments are outside the caller's address space.	<code>[K_EINVAL]</code>	An inconsistent actor capability was given. The value of <i>option</i> is not <code>K_FREEALL</code> or 0.	<code>[K_EUNKNOWN]</code>	<i>actorcap</i> does not specify a reachable actor.	<code>[K_EROUND]</code>	<i>rgndesc-&gt;startAddr</i> is not page-aligned.	<code>[K_ENOMEM]</code>	The system is out of resources.	<code>[K_EMAPPER]</code>	The segment mapper does not respect the kernel/mapper protocol
<code>[K_EFAULT]</code>	Some of the provided arguments are outside the caller's address space.												
<code>[K_EINVAL]</code>	An inconsistent actor capability was given. The value of <i>option</i> is not <code>K_FREEALL</code> or 0.												
<code>[K_EUNKNOWN]</code>	<i>actorcap</i> does not specify a reachable actor.												
<code>[K_EROUND]</code>	<i>rgndesc-&gt;startAddr</i> is not page-aligned.												
<code>[K_ENOMEM]</code>	The system is out of resources.												
<code>[K_EMAPPER]</code>	The segment mapper does not respect the kernel/mapper protocol												
<b>RESTRICTIONS</b>	The target actor and the current actor must be located on the same site.												

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`rgnAllocate(2K)`, `rgnMap(2K)`, `rgnInit(2K)`, `MpCreate(2SEG)`,  
`MpRelease(2SEG)`

<b>NAME</b>	rgnInitFromActor – allocate a region in an actor address space and initialise it from another region						
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnInitFromActor(KnCap *tgtactorcap, KnRgnDesc *tgtrgndesc, KnCap *srcactorcap, KnRgnDesc *srcrgndesc);</pre>						
<b>FEATURES</b>	MEM_PROTECTED, MEM_VIRTUAL						
<b>DESCRIPTION</b>	<p>The <i>rgnInitFromActor</i> function creates regions in a contiguous address range of the target actor address space. It then maps to these regions a copy of the segments and/or volatile memory which are already mapped to an address range of the source actor's address space.</p> <p>The target and source actors are specified by <i>tgtactorcap</i> and <i>srcactorcap</i> - pointers to the capabilities of the target and source actors respectively. If <i>srcactorcap</i> and/or <i>tgtactorcap</i> is K_MYACTOR, the address space of the current actor is used as the source and/or the destination, respectively. If <i>srcactorcap</i> and/or <i>tgtactorcap</i> is K_SVACTOR, the supervisor address space is used as the source and/or the destination, respectively. Note that, if <i>tgtactorcap</i> is K_SVACTOR, the allocated regions aren't attached to any particular supervisor actor; therefore, they can't be implicitly deallocated by performing an <i>actorDelete(2K)</i> on a supervisor actor.</p> <p>The <i>tgtrgndesc</i> field points to a <i>KnRgnDesc</i> structure containing the specification for the target address range and the attributes (options) for the regions to be created as described in <i>rgnAllocate(2K)</i>.</p> <p>The <i>srcrgndesc</i> field points to another <i>KnRgnDesc</i> structure where only the <i>startAddr</i> and <i>size</i> fields are significant. It specifies the source range starting address and size respectively. The <i>srcrgndesc-&gt;size</i> must to be less than or equal to <i>tgtrgndesc-&gt;size</i>. The source address range cannot contain holes: any address from the source address range has to belong to a region.</p> <p>The caller can specify any of the <i>tgtrgndesc-&gt;options</i> as described in <i>rgnAllocate(2K)</i> apart from K_RESERVED, which is prohibited. Note that the K_FILLZERO option is interpreted differently here: only the range from <i>tgtrgndesc-&gt;startAddr + srcrgndesc-&gt;size</i> to <i>tgtrgndesc-&gt;startAddr + tgtrgndesc-&gt;size - 1</i> will be zero-filled. All regions in the target address range will be created with the same attributes and will be differentiated by the objects (copies of particular segments or of particular volatile memory) mapped to them.</p>						
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.						
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EFAULT]</td> <td>Some of the provided arguments are outside the caller's address space.</td> </tr> <tr> <td>[K_EINVAL]</td> <td>An inconsistent actor capability was given.</td> </tr> <tr> <td>[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> </table>	[K_EFAULT]	Some of the provided arguments are outside the caller's address space.	[K_EINVAL]	An inconsistent actor capability was given.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EFAULT]	Some of the provided arguments are outside the caller's address space.						
[K_EINVAL]	An inconsistent actor capability was given.						
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.						

- [K\_EROUND] *tgtrgndesc->startAddr* is not page-aligned.
- [K\_EROUND] *srcrgndesc->startAddr* is not page-aligned.
- [K\_ESPACE] Tried to create a region outside the valid range for the address space of an actor as returned by *vmStat(2K)*.
- [K\_ESPACE] *tgtrgndesc->size* is zero.
- [K\_EOVERLAP] The K\_ANYWHERE option was specified and there is not sufficient room available in the address space to create the regions.
- [K\_EOVERLAP] The K\_RESTRICTIVE option was specified and there is not sufficient room in the target address range.
- [K\_EOVERLAP] The regions created overlap an existing region.
- [K\_ESPACE] The source region is outside the valid range for the address space of an actor as returned by *vmStat(2K)*.
- [K\_ESPACE] *srcrgndesc->size* is zero.
- [K\_EADDR] The source address range contains holes.
- [K\_ESIZE] *srcrgndesc->size* is greater than *tgtrgndesc->size*.
- [K\_ENOMEM] The system is out of resources.
- [K\_EMAPPER] A segment mapper doesn't respect the kernel/mapper protocol.

**RESTRICTIONS**

The target actor, the source actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnAllocate(2K)*, *rgnMap(2K)*, *rgnStat(2K)*

<b>NAME</b>	rgnMapFromActor – create a region in an actor address space and map another region to it										
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnMapFromActor(KnCap *tgtactorcap, KnRgnDesc *tgtrgndesc, KnCap *srcactorcap, KnRgnDesc *srcrgndesc);</pre>										
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL										
<b>DESCRIPTION</b>	<p>The <i>rgnMapFromActor</i> system call creates regions in a contiguous address range of the target actor address space. It then maps the segments and/or volatile memory already mapped to an address range of the source actor address space to the newly created regions.</p> <p>The target and source actors are specified by <i>tgtactorcap</i> and <i>srcactorcap</i> these are pointers to the capabilities of the target and source actors, respectively. If <i>srcactorcap</i> and/or <i>tgtactorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>tgtactorcap</i> is K_SVACTOR, the regions are allocated in the supervisor address space and are not attached to any particular supervisor actor. They cannot therefore be implicitly deallocated by performing an <i>actorDelete(2K)</i> of a supervisor actor.</p> <p>The <i>tgtrgndesc</i> pointer points to a <i>KnRgnDesc</i> structure containing the specification for the target address range and the attributes (options) for the regions to be created. This is the same as described in <i>rgnAllocate(2K)</i>, with the exception that K_FILLZERO is ignored and the K_RESERVED option is prohibited. All regions in the target address range will be created with the same attributes, and will be differentiated by the objects (particular segments or particular volatile memory) mapped to them.</p> <p>The <i>srcrgndesc</i> pointer points to another <i>KnRgnDesc</i> structure, where only the <i>startAddr</i> field is significant. It specifies the starting address of the source address range. This starting address must be aligned on a virtual page boundary. The size of the source address range is equal to the size of the target address range.</p>										
<b>RETURN VALUE</b>	If successful, K_OK is returned. Otherwise a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space.</td> </tr> <tr> <td>[K_EINVAL]</td> <td>An inconsistent actor capability was given.</td> </tr> <tr> <td>[K_EUNKNOWN]</td> <td><i>srcactorcap</i> or <i>tgtactorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td>[K_EROUND]</td> <td><i>tgtrgndesc-&gt;startAddr</i> is not page-aligned.</td> </tr> <tr> <td>[K_EROUND]</td> <td><i>srcrgndesc-&gt;startAddr</i> is not page-aligned.</td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.	[K_EINVAL]	An inconsistent actor capability was given.	[K_EUNKNOWN]	<i>srcactorcap</i> or <i>tgtactorcap</i> does not specify a reachable actor.	[K_EROUND]	<i>tgtrgndesc-&gt;startAddr</i> is not page-aligned.	[K_EROUND]	<i>srcrgndesc-&gt;startAddr</i> is not page-aligned.
[K_EFAULT]	Some of the arguments provided are outside the caller's address space.										
[K_EINVAL]	An inconsistent actor capability was given.										
[K_EUNKNOWN]	<i>srcactorcap</i> or <i>tgtactorcap</i> does not specify a reachable actor.										
[K_EROUND]	<i>tgtrgndesc-&gt;startAddr</i> is not page-aligned.										
[K_EROUND]	<i>srcrgndesc-&gt;startAddr</i> is not page-aligned.										

- [K\_ESPACE] Tried to create a region outside the valid range for the address space of an actor as returned by *vmStat(2K)*.
- [K\_ESPACE] *tgtrgndesc->size* is 0.
- [K\_EOVERLAP] The K\_ANYWHERE option was specified and there is not enough room available in the address space to create the regions.
- [K\_EOVERLAP] The K\_RESTRICTIVE option was specified and there is not enough room in the target address range.
- [K\_EOVERLAP] The regions created overlap an existing region.
- [K\_ESPACE] The source region is outside the valid range for the address space of an actor as returned by *vmStat(2K)*.
- [K\_ESIZE] If one of these three options are wrong K\_WRITABLE,K\_READABLE,K\_executable
- [K\_ESPACE] *srcrgndesc->size* is 0.
- [K\_EADDR] The source address range contains holes.
- [K\_ENOMEM] The system is out of resources.
- [K\_EMAPPER] A segment mapper does not respect the kernel/mapper protocol.
- [K\_EMAPPER] A segment mapper has detected a mapper-level error.

**RESTRICTIONS**

The target actor, the source actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnAllocate(2K)*, *rgnInit(2K)*, *rgnMap(2K)*, *rgnStat(2K)*, *vmStat(2K)*, *MpGetAccess(2SEG)*

<b>NAME</b>	rgnPhysMap – create a region in an actor address space and map (on demand) to it physical memory specified by the caller
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnPhysMap(KnCap *actorcap, KnRgnDesc *rgndesc, KnMemoryHandler handler, void *opaque);</pre>
<b>FEATURES</b>	MEM_VIRTUAL
<b>DESCRIPTION</b>	<p>The <i>rgnPhysMap</i> function creates a region in the address space of the target actor. The target actor is specified by <i>actorcap</i> - a pointer to the target actor capability. If <i>actorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>actorcap</i> is K_SVACTOR, the region is allocated in the supervisor address space and isn't attached to any particular supervisor actor; therefore, it can't be implicitly deallocated by performing an <i>actorDelete(2K)</i> on a supervisor actor.</p> <p>The <i>rgndesc</i> field points to a <i>KnRgnDesc</i> structure containing the specification of the region to be created as described in <i>rgnAllocate (2K)</i>.</p> <p>The caller can specify any of the <i>rgndesc-&gt;options</i> as described in <i>rgnAllocate (2K)</i>; with the exceptions that the K_FILLZERO flag is ignored, and the K_RESERVED flag is prohibited.</p> <p>The <i>handler</i> argument specifies a memory handler which is attached to the newly created region. When access to a virtual page of the region is required (for example, a fault happens), the kernel invokes the memory handler. The handler returns the physical address to be mapped and the mapping attributes to the required virtual address.</p> <p>The memory handler interface is machine-dependent. The kernel on i386/i486 hardware defines the following handler interface:</p> <pre>#include &lt;mem/chMem.h&gt; #include &lt;mem/f_chMem.h&gt;  int handler (opaque, offset, paddr, access, cntlBits) void*      opaque ; VmOffset   offset ; PhysAddr   *paddr ; VmFlags    access ; PteCntlBits *cntlBits ;</pre> <p>Where <i>opaque</i> is an input argument specified by the <i>opaque</i> argument of <i>rgnPhysMap</i>; <i>offset</i> is another input argument which specifies the required page offset within the region. The <i>paddr</i> output argument specifies the physical address to be mapped. The last input argument is <i>access</i>, which defines the type of access rights required. If the K_WRITABLE bit is set in the <i>access</i> argument, write access is required. The second output argument, <i>cntlBits</i> specifies the</p>

mapping control bits as defined in the Intel386/Intel486 Programmer's Reference Manual. The valid bits are the following:

- PTE\_PRESENT (0x01)
- PTE\_READ\_WRITE (0x02)
- PTE\_USER\_SUPERVISOR (0x04)
- PTE\_WRITE\_TRANSPARENT (0x08)
- PTE\_CACHE\_DISABLE (0x10)

As the the memory handler is designed for mapping I/O devices it doesn't include a feature to detect and write back the data modified in the region.

**RETURN VALUE**

If successful K\_OK is returned, otherwise a negative error code is returned.

**ERRORS**

- [K\_EFAULT] Some of the provided arguments are outside the caller's address space.
- [K\_EINVAL] An inconsistent actor capability was given.
- [K\_EUNKNOWN] *actorcap* does not specify a reachable actor.
- [K\_EROUND] *rgndesc->startAddr* is not page-aligned.
- [K\_ESPACE] Tried to create a region outside the valid range for the address space of an actor as returned by *vmStat(2K)*.
- [K\_ESPACE] *rgndesc->size* is zero.
- [K\_EOVERLAP] The K\_ANYWHERE option was specified and there is insufficient room available in the address space to create the region.
- [K\_EOVERLAP] The K\_RESTRICTIVE option was specified and there is insufficient room in the target address range.
- [K\_EOVERLAP] The created region overlaps an existing region.
- [K\_ENOMEM] The system is out of resources.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnAllocate(2K)*, *rgnInit(2K)*, *rgnMap(2K)*, *rgnStat(2K)*

<b>NAME</b>	rgnSetInherit, rgnSetPaging, rgnSetOpaque – Change inheritance options associated with a region; Change paging options associated with a region; Change inheritance options associated with a region; Change opaque values associated with a region
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnSetInherit(KnCap * actorcap, KnRgnDesc * rgndesc);  int rgnSetPaging(KnCap * actorcap, KnRgnDesc * rgndesc);  int rgnSetOpaque(KnCap * actorcap, KnRgnDesc * rgndesc);</pre>
<b>FEATURES</b>	MEM_VIRTUAL
<b>DESCRIPTION</b>	<p>These calls allow you to change the attributes of all regions from the target address range of the target actor address space.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>actorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The <i>rgndesc</i> pointer points to a <i>KnRgnDesc</i> structure which contains the specification of the target address range and the new values of the attributes. Only the following members of the structure are used:</p> <pre>VmAddr    startAddr ; VmSize    size ; VmFlags   options ; void*     opaque1 ; VmFlags   opaque2 ;</pre> <p>The <i>startAddr</i> field gives the start of the target address range and the <i>size</i> field gives its size; <i>rgndesc-&gt;startAddr</i> must be aligned with a virtual page boundary (see <i>vmPageSize</i> (2K)). If <i>rgndesc-&gt;size</i> is not page-aligned the kernel rounds it up to the next virtual page boundary.</p> <p>The <i>options</i> field gives the new attributes of the regions from the target address range.</p> <p>The target address range can contain holes; sub-ranges without regions mapped to them.</p> <p>If <i>rgndesc-&gt;startAddr</i> is inside a region, the implementation replaces the original region with two regions. The first region conserves the old attributes and terminates at <i>rgndesc-&gt;startAddr - 1</i>. The second region takes the new attributes and starts immediately after the first one.</p> <p>If <i>rgndesc-&gt;startAddr + rgndesc-&gt;size</i> is inside a region, the implementation replaces the original region with three regions. The first one conserves the old attributes and terminates at <i>rgndesc-&gt;startAddr - 1</i>. The second region takes the</p>

new attributes and terminates at *rgndesc->startAddr + rgndesc->size - 1*. The third region conserves the old attributes and starts immediately after the second one.

An attribute modification can sometimes replace a number of consecutive regions with a single one.

The *rgnSetInherit* command replaces the values of the K\_INHERITCOPY and K\_INHERITSHARE attributes with the corresponding values from the *options* field of *rgndesc*.

The *rgnSetPaging* command replaces the values of the K\_NODEMAND and K\_NOWAITFORMEMORY attributes with the corresponding values from the *options* field of *rgndesc*.

The *rgnSetOpaque* command replaces the opaque values associated with the region with the new opaque values of *rgndesc*.

**RETURN VALUE**

If successful K\_OK is returned, otherwise a negative error code is returned.

**ERRORS**

- [K\_EFAULT]                   Some of the arguments provided are outside the caller's address space.
- [K\_EINVAL]                   An inconsistent actor capability was given.
- [K\_EUNKNOWN]                *actorcap* does not specify a reachable actor.
- [K\_EROUND]                   *rgndesc->startAddr* is not page-aligned.
- [K\_ENOMEM]                   The system is out of resources.
- [K\_ENOMEM]                   Some or all of the memory identified by the operation could not be locked when K\_NODEMAND and K\_NOWAITFORMEMORY had both been specified.
- [K\_EMAPPER]                  The segment mapper has detected a mapper-level error.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnAllocate(2K)*

<b>NAME</b>	rgnSetInherit, rgnSetPaging, rgnSetOpaque – Change inheritance options associated with a region; Change paging options associated with a region; Change inheritance options associated with a region; Change opaque values associated with a region
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnSetInherit(KnCap * actorcap, KnRgnDesc * rgndesc);  int rgnSetPaging(KnCap * actorcap, KnRgnDesc * rgndesc);  int rgnSetOpaque(KnCap * actorcap, KnRgnDesc * rgndesc);</pre>
<b>FEATURES</b>	MEM_VIRTUAL
<b>DESCRIPTION</b>	<p>These calls allow you to change the attributes of all regions from the target address range of the target actor address space.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>actorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The <i>rgndesc</i> pointer points to a <i>KnRgnDesc</i> structure which contains the specification of the target address range and the new values of the attributes. Only the following members of the structure are used:</p> <pre>VmAddr    startAddr ; VmSize    size ; VmFlags   options ; void*     opaque1 ; VmFlags   opaque2 ;</pre> <p>The <i>startAddr</i> field gives the start of the target address range and the <i>size</i> field gives its size; <i>rgndesc-&gt;startAddr</i> must be aligned with a virtual page boundary (see <i>vmPageSize</i> (2K)). If <i>rgndesc-&gt;size</i> is not page-aligned the kernel rounds it up to the next virtual page boundary.</p> <p>The <i>options</i> field gives the new attributes of the regions from the target address range.</p> <p>The target address range can contain holes; sub-ranges without regions mapped to them.</p> <p>If <i>rgndesc-&gt;startAddr</i> is inside a region, the implementation replaces the original region with two regions. The first region conserves the old attributes and terminates at <i>rgndesc-&gt;startAddr - 1</i>. The second region takes the new attributes and starts immediately after the first one.</p> <p>If <i>rgndesc-&gt;startAddr + rgndesc-&gt;size</i> is inside a region, the implementation replaces the original region with three regions. The first one conserves the old attributes and terminates at <i>rgndesc-&gt;startAddr - 1</i>. The second region takes the</p>

new attributes and terminates at *rgndesc->startAddr + rgndesc->size - 1*. The third region conserves the old attributes and starts immediately after the second one.

An attribute modification can sometimes replace a number of consecutive regions with a single one.

The *rgnSetInherit* command replaces the values of the K\_INHERITCOPY and K\_INHERITSHARE attributes with the corresponding values from the *options* field of *rgndesc*.

The *rgnSetPaging* command replaces the values of the K\_NODEMAND and K\_NOWAITFORMEMORY attributes with the corresponding values from the *options* field of *rgndesc*.

The *rgnSetOpaque* command replaces the opaque values associated with the region with the new opaque values of *rgndesc*.

**RETURN VALUE**

If successful K\_OK is returned, otherwise a negative error code is returned.

**ERRORS**

- [K\_EFAULT]                   Some of the arguments provided are outside the caller's address space.
- [K\_EINVAL]                   An inconsistent actor capability was given.
- [K\_EUNKNOWN]                *actorcap* does not specify a reachable actor.
- [K\_EROUND]                   *rgndesc->startAddr* is not page-aligned.
- [K\_ENOMEM]                   The system is out of resources.
- [K\_ENOMEM]                   Some or all of the memory identified by the operation could not be locked when K\_NODEMAND and K\_NOWAITFORMEMORY had both been specified.
- [K\_EMAPPER]                  The segment mapper has detected a mapper-level error.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnAllocate(2K)*

<b>NAME</b>	rgnSetInherit, rgnSetPaging, rgnSetOpaque – Change inheritance options associated with a region; Change paging options associated with a region; Change inheritance options associated with a region; Change opaque values associated with a region
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnSetInherit(KnCap * actorcap, KnRgnDesc * rgndesc);  int rgnSetPaging(KnCap * actorcap, KnRgnDesc * rgndesc);  int rgnSetOpaque(KnCap * actorcap, KnRgnDesc * rgndesc);</pre>
<b>FEATURES</b>	MEM_VIRTUAL
<b>DESCRIPTION</b>	<p>These calls allow you to change the attributes of all regions from the target address range of the target actor address space.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>actorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The <i>rgndesc</i> pointer points to a <i>KnRgnDesc</i> structure which contains the specification of the target address range and the new values of the attributes. Only the following members of the structure are used:</p> <pre>VmAddr    startAddr ; VmSize    size ; VmFlags   options ; void*     opaque1 ; VmFlags   opaque2 ;</pre> <p>The <i>startAddr</i> field gives the start of the target address range and the <i>size</i> field gives its size; <i>rgndesc-&gt;startAddr</i> must be aligned with a virtual page boundary (see <i>vmPageSize</i> (2K)). If <i>rgndesc-&gt;size</i> is not page-aligned the kernel rounds it up to the next virtual page boundary.</p> <p>The <i>options</i> field gives the new attributes of the regions from the target address range.</p> <p>The target address range can contain holes; sub-ranges without regions mapped to them.</p> <p>If <i>rgndesc-&gt;startAddr</i> is inside a region, the implementation replaces the original region with two regions. The first region conserves the old attributes and terminates at <i>rgndesc-&gt;startAddr - 1</i>. The second region takes the new attributes and starts immediately after the first one.</p> <p>If <i>rgndesc-&gt;startAddr + rgndesc-&gt;size</i> is inside a region, the implementation replaces the original region with three regions. The first one conserves the old attributes and terminates at <i>rgndesc-&gt;startAddr - 1</i>. The second region takes the</p>

new attributes and terminates at *rgndesc->startAddr + rgndesc->size - 1*. The third region conserves the old attributes and starts immediately after the second one.

An attribute modification can sometimes replace a number of consecutive regions with a single one.

The *rgnSetInherit* command replaces the values of the K\_INHERITCOPY and K\_INHERITSHARE attributes with the corresponding values from the *options* field of *rgndesc*.

The *rgnSetPaging* command replaces the values of the K\_NODEMAND and K\_NOWAITFORMEMORY attributes with the corresponding values from the *options* field of *rgndesc*.

The *rgnSetOpaque* command replaces the opaque values associated with the region with the new opaque values of *rgndesc*.

**RETURN VALUE**

If successful K\_OK is returned, otherwise a negative error code is returned.

**ERRORS**

[K_EFAULT]	Some of the arguments provided are outside the caller's address space.
[K_EINVAL]	An inconsistent actor capability was given.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EROUND]	<i>rgndesc-&gt;startAddr</i> is not page-aligned.
[K_ENOMEM]	The system is out of resources.
[K_ENOMEM]	Some or all of the memory identified by the operation could not be locked when K_NODEMAND and K_NOWAITFORMEMORY had both been specified.
[K_EMAPPER]	The segment mapper has detected a mapper-level error.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnAllocate(2K)*

<b>NAME</b>	rgnSetProtect – change protection options associated with a region
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnSetProtect(KnCap *actorcap, KnRgnDesc *rgndesc);</pre>
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL
<b>DESCRIPTION</b>	<p>The <i>rgnSetProtect</i> function allows you to change the protection attributes of all regions of the target address range of the target actor address space.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is <code>K_MYACTOR</code>, the address space of the current actor is used. If <i>actorcap</i> is <code>K_SVACTOR</code>, the supervisor address space is used.</p> <p>The <i>rgndesc</i> pointer points to a <i>KnRgnDesc</i> structure which contains the specification of the target address range and the new values of the attributes. Only the following members of the structure are used:</p> <pre>VmAddr  startAddr ; VmSize  size ; VmFlags options ;</pre> <p>The <i>startAddr</i> field gives the start of the target address range and the <i>size</i> field gives its size; <i>rgndesc-&gt;startAddr</i> must be aligned with a page boundary (see <i>vmPageSize(2K)</i>). If <i>rgndesc-&gt;size</i> is not page-aligned the kernel rounds it up to the next virtual page boundary.</p> <p>The <i>options</i> field gives the new attributes of the regions from the target address range.</p> <p>The target address range can contain holes; sub-ranges without regions mapped to them.</p> <p>If <i>rgndesc-&gt;startAddr</i> is inside a region, the implementation replaces the original region with two regions. The first region conserves the old attributes and terminates at <i>rgndesc-&gt;startAddr - 1</i>. The second region takes the new attributes and starts immediately after the first one.</p> <p>If <i>rgndesc-&gt;startAddr + rgndesc-&gt;size</i> is inside a region, the implementation replaces the original region with three regions. The first one conserves the old attributes and terminates at <i>rgndesc-&gt;startAddr - 1</i>. The second region takes the new attributes and terminates at <i>rgndesc-&gt;startAddr + rgndesc-&gt;size - 1</i>. The third region conserves the old attributes and starts immediately after the second one.</p> <p>An attribute modification can sometimes replace a number of consecutive regions with a single one.</p> <p>The values for the <i>options</i> field of <i>rgndesc</i> parameter must be among <code>K_READABLE</code>, <code>K_WRITABLE</code>, <code>K_EXECUTABLE</code> and <code>K_SUPERVISOR</code>.</p>

The `rgnSetProtect` command replaces the values of the `K_READABLE`, `K_WRITABLE`, `K_EXECUTABLE` and `K_SUPERVISOR` attributes with the corresponding values from the *options* field of *rgndesc*.

**RETURN VALUE**

If successful `K_OK` is returned, otherwise a negative error code is returned.

**ERRORS**

[K_EFAULT]	Some of the arguments provided are outside the caller's address space.
[K_EINVAL]	An inconsistent actor capability or invalid options were given.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EROUND]	<i>rgndesc-&gt;startAddr</i> is not page-aligned.
[K_ENOMEM]	The system is out of resources.
[K_EMAPPER]	The segment mapper has detected a mapper-level error.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`rgnAllocate(2K)`, `MpGetAccess(2SEG)`

<b>NAME</b>	rgnStat – get the statistics of a region of an actor address space
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnStat(KnCap *actorcap, KnRgnDesc *rgndesc, KnRgnStat *stat, unsigned bufsize);</pre>
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL
<b>DESCRIPTION</b>	<p>The <i>rgnStat</i> function gets the status of regions from the given address range of the target actor address space.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is <code>K_MYACTOR</code>, the address space of the current actor is used. If the target actor is a user actor <i>rgnStat</i> takes all regions of the target address range into account. If the target actor is a supervisor actor <i>rgnStat</i> only takes into account the regions attached to the target supervisor actor. If <i>actorcap</i> is <code>K_SVACTOR</code>, <i>rgnStat</i> takes all regions of the target supervisor address range into account.</p> <p>The <i>rgndesc</i> pointer points to a <i>KnRgnDesc</i> structure which contains the specification of the target address range. Only the following members of the structure are used:</p> <pre>VmAddr  startAddr ; VmSize  size ; VmFlags options</pre> <p>The <i>startAddr</i> field gives the start of the target address range and the <i>size</i> field gives its size.</p> <p>The <i>stat</i> argument points to a buffer in which the description of the region will be returned. The <i>bufsize</i> argument gives the size of the buffer. If the <code>K_USERSPACE</code> flag is set in <i>rgndesc-&gt;option</i> the buffer is in the current user address space; otherwise the buffer is in the caller (current user or supervisor) address space.</p> <p>The statuses are returned as a number of items, each one describing a particular sub-range of the target address range. The <i>KnRgnStat</i> structure defines the fields of each item, as follows:</p> <pre>VmAddr  startAddr ; VmSize  size ; VmFlags options ; void*   opaque1 ; VmFlags opaque2 ; VmOffset offset ; KnLcStat lcstat ;</pre> <p>The <i>startAddr</i> field is the starting address of the sub-range described by the item and <i>size</i> is size of the sub-range. There is exactly one item per region which has a non-empty intersection with the target address range. The items are sorted by increasing order of <i>startAddr</i>. If <i>rgndesc-&gt;startAddr</i> is inside a region, the <i>startAddr</i></p>

of the first item is equal to *rgndesc->startAddr*. If *rgndesc->startAddr + rgndesc->size* is inside a region, the *size* of the last item is equal to  $(rgndesc->startAddr + rgndesc->size) - rgn\_start\_addr$ , where *rgn\_start\_addr* is the start address of the last region of the target range. Otherwise, an item describes precisely one region.

The *options* field defines the region options as specified in *rgnAllocate(2K)*.

*(The next four fields are only meaningful if virtual memory management is being used)*

The values *opaque1* and *opaque2* define the opaque values associated with the region.

The *offset* value defines the offset of the item's *startAddr* within the mapped segment.

If the *K\_RGNFULLSTAT* flag is set in the *rgndesc->options* field, the *lcstat* field contains the following status information about the range of the local cache mapped to the item's address range:

```
VmSize  physMem ;
VmSize  lockMem ;
KnCap   segcap ;
KnCap   lccap  ;
```

The *physMem* field contains the size (in bytes) of the physical memory mapped to the address range.

The *lockMem* field contains the size (in bytes) of no-demand physical memory mapped to the address range.

If the target local cache corresponds to an external segment, the *segcap* field contains the capability of the segment; otherwise the field contains zero.

If the target local cache can be designated by a capability, the *lccap* field contains the capability of the local cache; otherwise the field contains zero.

## RETURN VALUE

If successful, *rgnStat* returns the number of items needed to describe the target address range, otherwise a negative error code is returned. When the *MEM\_FLAT* module is enabled and the caller is a supervisor actor (*K\_SVACTOR*), this primitive returns 0. In this case no regions are available.

## ERRORS

[K_EFAULT]	Some of the arguments provided are outside the caller's address space.
[K_EINVAL]	An inconsistent actor capability was given.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.

## RESTRICTIONS

The target actor and the current actor must be located on the same site.

## ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

rgnAllocate(2K)

<b>NAME</b>	rtMutexInit, rtMutexGet, rtMutexRel, rtMutexTry – Initialize a realtime mutex; Acquire a realtime mutex; Release a realtime mutex; Try to acquire a realtime mutex
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chRtMutex.h&gt; int rtMutexInit(KnRtMutex * mutex);  int rtMutexGet(KnRtMutex * mutex);  int rtMutexRel(KnRtMutex * mutex);  int rtMutexTry(KnRtMutex * mutex);</pre>
<b>FEATURES</b>	RTMUTEX
<b>DESCRIPTION</b>	<p>Realtime mutexes are binary semaphores used to protect shared data from concurrent access: a realtime mutex is a two state variable - free or locked. When a thread is blocking higher priority threads because it owns one or more realtime mutexes, it executes either its highest priority, or the priority of the highest priority thread waiting on any of the realtime mutexes owned by the thread.</p> <p>Realtime mutexes are <i>KnRtMutex</i> structures allocated in user memory.</p> <p>The <i>rtMutexInit</i> function initializes the realtime mutex whose address is <i>mutex</i>. The mutex is initialized as free.</p> <p>The <i>rtMutexGet</i> function is used to acquire a realtime mutex. If the mutex is free, it becomes locked and the caller continues to execute normally. If the mutex is locked, the caller is blocked. The owner thread inherits the priority level of the calling thread as long as it continues to own the mutex. The implementation updates its execution priority to the maximum of its assigned priority plus all its inherited priority. Furthermore, if the owner thread itself becomes blocked on another realtime mutex, the same priority inheritance effect is propagated recursively to this other owner thread.</p> <p>The <i>rtMutexRel</i> function is used to release a realtime mutex. If threads are blocked behind the mutex, the thread with the highest priority is awakened. The priority of the calling thread can be reduced by using <i>rtMutexRel</i>.</p> <p>The <i>rtMutexTry</i> function is an attempt to acquire a realtime mutex: it has the same effect as <i>rtMutexGet</i>, except that if the mutex is locked, the thread is not blocked (a return code is provided).</p> <p>Any priority adjustments caused by priority inheritance are transparent and are not visible through the <i>threadScheduler</i> interface. In other words, the priority value returned from <i>threadScheduler</i> will be the value specified by the most recent <i>threadScheduler</i> or <i>threadCreate</i> call.</p> <p>A blocking <i>rtMutexGet</i> is <i>NONABORTABLE</i> (see <i>threadAbort</i> (2K)).</p>

**RETURN VALUE** The *rtMutexTry* function returns 0 if the mutex was already locked, 1 if it was free. Upon successful completion of *rtMutexInit*, *rtMutexGet* and *rtMutexRel*, 0 is returned. Otherwise, a negative error code is returned.

**ERRORS** [K\_EFAULT] Some of the data provided are outside the current actor's address space.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *mutexInit(2K)*

<b>NAME</b>	rtMutexInit, rtMutexGet, rtMutexRel, rtMutexTry – Initialize a realtime mutex; Acquire a realtime mutex; Release a realtime mutex; Try to acquire a realtime mutex
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chRtMutex.h&gt; int rtMutexInit(KnRtMutex * mutex);  int rtMutexGet(KnRtMutex * mutex);  int rtMutexRel(KnRtMutex * mutex);  int rtMutexTry(KnRtMutex * mutex);</pre>
<b>FEATURES</b>	RTMUTEX
<b>DESCRIPTION</b>	<p>Realtime mutexes are binary semaphores used to protect shared data from concurrent access: a realtime mutex is a two state variable - free or locked. When a thread is blocking higher priority threads because it owns one or more realtime mutexes, it executes either its highest priority, or the priority of the highest priority thread waiting on any of the realtime mutexes owned by the thread.</p> <p>Realtime mutexes are <i>KnRtMutex</i> structures allocated in user memory.</p> <p>The <i>rtMutexInit</i> function initializes the realtime mutex whose address is <i>mutex</i>. The mutex is initialized as free.</p> <p>The <i>rtMutexGet</i> function is used to acquire a realtime mutex. If the mutex is free, it becomes locked and the caller continues to execute normally. If the mutex is locked, the caller is blocked. The owner thread inherits the priority level of the calling thread as long as it continues to own the mutex. The implementation updates its execution priority to the maximum of its assigned priority plus all its inherited priority. Furthermore, if the owner thread itself becomes blocked on another realtime mutex, the same priority inheritance effect is propagated recursively to this other owner thread.</p> <p>The <i>rtMutexRel</i> function is used to release a realtime mutex. If threads are blocked behind the mutex, the thread with the highest priority is awakened. The priority of the calling thread can be reduced by using <i>rtMutexRel</i>.</p> <p>The <i>rtMutexTry</i> function is an attempt to acquire a realtime mutex: it has the same effect as <i>rtMutexGet</i>, except that if the mutex is locked, the thread is not blocked (a return code is provided).</p> <p>Any priority adjustments caused by priority inheritance are transparent and are not visible through the <i>threadScheduler</i> interface. In other words, the priority value returned from <i>threadScheduler</i> will be the value specified by the most recent <i>threadScheduler</i> or <i>threadCreate</i> call.</p> <p>A blocking <i>rtMutexGet</i> is <i>NONABORTABLE</i> (see <i>threadAbort</i> (2K)).</p>

**RETURN VALUE** The *rtMutexTry* function returns 0 if the mutex was already locked, 1 if it was free. Upon successful completion of *rtMutexInit*, *rtMutexGet* and *rtMutexRel*, 0 is returned. Otherwise, a negative error code is returned.

**ERRORS** [K\_EFAULT] Some of the data provided are outside the current actor's address space.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *mutexInit(2K)*

<b>NAME</b>	rtMutexInit, rtMutexGet, rtMutexRel, rtMutexTry – Initialize a realtime mutex; Acquire a realtime mutex; Release a realtime mutex; Try to acquire a realtime mutex
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chRtMutex.h&gt; int rtMutexInit(KnRtMutex * mutex);  int rtMutexGet(KnRtMutex * mutex);  int rtMutexRel(KnRtMutex * mutex);  int rtMutexTry(KnRtMutex * mutex);</pre>
<b>FEATURES</b>	RTMUTEX
<b>DESCRIPTION</b>	<p>Realtime mutexes are binary semaphores used to protect shared data from concurrent access: a realtime mutex is a two state variable - free or locked. When a thread is blocking higher priority threads because it owns one or more realtime mutexes, it executes either its highest priority, or the priority of the highest priority thread waiting on any of the realtime mutexes owned by the thread.</p> <p>Realtime mutexes are <i>KnRtMutex</i> structures allocated in user memory.</p> <p>The <i>rtMutexInit</i> function initializes the realtime mutex whose address is <i>mutex</i>. The mutex is initialized as free.</p> <p>The <i>rtMutexGet</i> function is used to acquire a realtime mutex. If the mutex is free, it becomes locked and the caller continues to execute normally. If the mutex is locked, the caller is blocked. The owner thread inherits the priority level of the calling thread as long as it continues to own the mutex. The implementation updates its execution priority to the maximum of its assigned priority plus all its inherited priority. Furthermore, if the owner thread itself becomes blocked on another realtime mutex, the same priority inheritance effect is propagated recursively to this other owner thread.</p> <p>The <i>rtMutexRel</i> function is used to release a realtime mutex. If threads are blocked behind the mutex, the thread with the highest priority is awakened. The priority of the calling thread can be reduced by using <i>rtMutexRel</i>.</p> <p>The <i>rtMutexTry</i> function is an attempt to acquire a realtime mutex: it has the same effect as <i>rtMutexGet</i>, except that if the mutex is locked, the thread is not blocked (a return code is provided).</p> <p>Any priority adjustments caused by priority inheritance are transparent and are not visible through the <i>threadScheduler</i> interface. In other words, the priority value returned from <i>threadScheduler</i> will be the value specified by the most recent <i>threadScheduler</i> or <i>threadCreate</i> call.</p> <p>A blocking <i>rtMutexGet</i> is <i>NONABORTABLE</i> (see <i>threadAbort</i> (2K)).</p>

**RETURN VALUE** The *rtMutexTry* function returns 0 if the mutex was already locked, 1 if it was free. Upon successful completion of *rtMutexInit*, *rtMutexGet* and *rtMutexRel*, 0 is returned. Otherwise, a negative error code is returned.

**ERRORS** [K\_EFAULT] Some of the data provided are outside the current actor's address space.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *mutexInit(2K)*

<b>NAME</b>	rtMutexInit, rtMutexGet, rtMutexRel, rtMutexTry – Initialize a realtime mutex; Acquire a realtime mutex; Release a realtime mutex; Try to acquire a realtime mutex
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chRtMutex.h&gt; int rtMutexInit(KnRtMutex * mutex);  int rtMutexGet(KnRtMutex * mutex);  int rtMutexRel(KnRtMutex * mutex);  int rtMutexTry(KnRtMutex * mutex);</pre>
<b>FEATURES</b>	RTMUTEX
<b>DESCRIPTION</b>	<p>Realtime mutexes are binary semaphores used to protect shared data from concurrent access: a realtime mutex is a two state variable - free or locked. When a thread is blocking higher priority threads because it owns one or more realtime mutexes, it executes either its highest priority, or the priority of the highest priority thread waiting on any of the realtime mutexes owned by the thread.</p> <p>Realtime mutexes are <i>KnRtMutex</i> structures allocated in user memory.</p> <p>The <i>rtMutexInit</i> function initializes the realtime mutex whose address is <i>mutex</i>. The mutex is initialized as free.</p> <p>The <i>rtMutexGet</i> function is used to acquire a realtime mutex. If the mutex is free, it becomes locked and the caller continues to execute normally. If the mutex is locked, the caller is blocked. The owner thread inherits the priority level of the calling thread as long as it continues to own the mutex. The implementation updates its execution priority to the maximum of its assigned priority plus all its inherited priority. Furthermore, if the owner thread itself becomes blocked on another realtime mutex, the same priority inheritance effect is propagated recursively to this other owner thread.</p> <p>The <i>rtMutexRel</i> function is used to release a realtime mutex. If threads are blocked behind the mutex, the thread with the highest priority is awakened. The priority of the calling thread can be reduced by using <i>rtMutexRel</i>.</p> <p>The <i>rtMutexTry</i> function is an attempt to acquire a realtime mutex: it has the same effect as <i>rtMutexGet</i>, except that if the mutex is locked, the thread is not blocked (a return code is provided).</p> <p>Any priority adjustments caused by priority inheritance are transparent and are not visible through the <i>threadScheduler</i> interface. In other words, the priority value returned from <i>threadScheduler</i> will be the value specified by the most recent <i>threadScheduler</i> or <i>threadCreate</i> call.</p> <p>A blocking <i>rtMutexGet</i> is <i>NONABORTABLE</i> (see <i>threadAbort</i> (2K)).</p>

**RETURN VALUE**

The *rtMutexTry* function returns 0 if the mutex was already locked, 1 if it was free. Upon successful completion of *rtMutexInit*, *rtMutexGet* and *rtMutexRel*, 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EFAULT]                      Some of the data provided are outside the current actor's address space.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*mutexInit(2K)*

**NAME** schedAdmin – scheduling classes administration

**SYNOPSIS**

```
#include <sched/chSched.h>
#include <sched/chRt.h>
#include <sched/chTs.h>
int schedAdmin(void *adminparam);
```

**FEATURES** SCHED\_CLASS

**DESCRIPTION** **Caution** - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.

The *schedAdmin* function is used to administer the attributes of the scheduling classes CLASS\_RT and CLASS\_TS, within the SCHED\_CLASS scheduler.

It provides the necessary functionality to implement the CHORUS/MiX V.4 Process Manager. It is not intended for general use by other applications.

**RETURN VALUE** Upon successful completion, K\_OK is returned. Otherwise, a negative error code is returned.

**ERRORS** [K\_EINVAL] Invalid argument.

**RESTRICTION** The *schedAdmin* function is restricted to supervisor threads.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** threadScheduler(2K)

<b>NAME</b>	semInit, semP, semV – initialize a semaphore; wait on a semaphore; signal a semaphore
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSem.h&gt; int semInit(KnSem * sem, unsigned int count);  int semP(KnSem * sem, KnTimeVal * waitLimit);  int semV(KnSem * sem);</pre>
<b>FEATURES</b>	SEM
<b>DESCRIPTION</b>	<p>Semaphores are <i>KnSem</i> structures allocated in the user memory.</p> <p><i>semInit</i> initializes the semaphore the address of which is <i>sem</i>. <i>count</i> is the initial positive value given to the semaphore counter.</p> <p>Unless they are shared between two or more actors, statically allocated semaphores can be initialized using the <i>K_KNSEM_INITIALIZER(cnt)</i> macro, where <i>cnt</i> is the initial positive value given to the semaphore counter. This macro is used as follows:</p> <pre>KnSem mySem = K_KNSEM_INITIALIZER(cnt) ;</pre> <p><i>semInit</i> is used to initialize a semaphore that is shared between two or more actors.</p> <p><i>semP</i> decrements the semaphore counter; if the counter reaches a strictly negative value, the calling thread is blocked, according to the options described by <i>waitLimit</i> in <i>intro(2K)</i>. <i>waitLimit</i> is a pointer to a <i>KnTimeVal</i> structure containing a timeout interval as described in <i>sysTime(2K)</i>.</p> <p><i>semV</i> increments the counter. If the new value is less than or equal to 0, the thread that has been blocked behind the semaphore for the longest time is awakened.</p> <p>As semaphore structures are allocated in the client memory, the number of semaphores used by an application is not limited. Any modification to the semaphore structure while the semaphore is in use will cause unpredictable synchronization behavior in the application.</p> <p>A blocking <i>semP</i> is <i>ABORTABLE</i> (see <i>threadAbort(2K)</i>).</p>
<b>RESTRICTIONS</b>	<p>A user application and a supervisor application may not share a semaphore.</p> <p>On the contrary, two user applications may share a semaphore by mapping it in both user address spaces.</p>
<b>RETURN VALUE</b>	Upon successful completion, a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	[K_EINVAL]                      The <i>waitLimit</i> is not a valid <i>KnTimeVal</i> .

[K\_EABORT] *semP* has been aborted.

[K\_ETIMEOUT] The timeout occurred.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`mutexInit(2K)`

<b>NAME</b>	semInit, semP, semV – initialize a semaphore; wait on a semaphore; signal a semaphore
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSem.h&gt; int semInit(KnSem * sem, unsigned int count);  int semP(KnSem * sem, KnTimeVal * waitLimit);  int semV(KnSem * sem);</pre>
<b>FEATURES</b>	SEM
<b>DESCRIPTION</b>	<p>Semaphores are <i>KnSem</i> structures allocated in the user memory.</p> <p><i>semInit</i> initializes the semaphore the address of which is <i>sem</i>. <i>count</i> is the initial positive value given to the semaphore counter.</p> <p>Unless they are shared between two or more actors, statically allocated semaphores can be initialized using the <i>K_KNSEM_INITIALIZER(cnt)</i> macro, where <i>cnt</i> is the initial positive value given to the semaphore counter. This macro is used as follows:</p> <pre>KnSem mySem = K_KNSEM_INITIALIZER(cnt) ;</pre> <p><i>semInit</i> is used to initialize a semaphore that is shared between two or more actors.</p> <p><i>semP</i> decrements the semaphore counter; if the counter reaches a strictly negative value, the calling thread is blocked, according to the options described by <i>waitLimit</i> in <i>intro(2K)</i>. <i>waitLimit</i> is a pointer to a <i>KnTimeVal</i> structure containing a timeout interval as described in <i>sysTime(2K)</i>.</p> <p><i>semV</i> increments the counter. If the new value is less than or equal to 0, the thread that has been blocked behind the semaphore for the longest time is awakened.</p> <p>As semaphore structures are allocated in the client memory, the number of semaphores used by an application is not limited. Any modification to the semaphore structure while the semaphore is in use will cause unpredictable synchronization behavior in the application.</p> <p>A blocking <i>semP</i> is <i>ABORTABLE</i> (see <i>threadAbort(2K)</i>).</p>
<b>RESTRICTIONS</b>	<p>A user application and a supervisor application may not share a semaphore.</p> <p>On the contrary, two user applications may share a semaphore by mapping it in both user address spaces.</p>
<b>RETURN VALUE</b>	Upon successful completion, a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	[K_EINVAL]                      The <i>waitLimit</i> is not a valid <i>KnTimeVal</i> .

[K\_EABORT] *semP* has been aborted.

[K\_ETIMEOUT] The timeout occurred.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`mutexInit(2K)`

<b>NAME</b>	semInit, semP, semV – initialize a semaphore; wait on a semaphore; signal a semaphore
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSem.h&gt; int semInit(KnSem * sem, unsigned int count);  int semP(KnSem * sem, KnTimeVal * waitLimit);  int semV(KnSem * sem);</pre>
<b>FEATURES</b>	SEM
<b>DESCRIPTION</b>	<p>Semaphores are <i>KnSem</i> structures allocated in the user memory.</p> <p><i>semInit</i> initializes the semaphore the address of which is <i>sem</i>. <i>count</i> is the initial positive value given to the semaphore counter.</p> <p>Unless they are shared between two or more actors, statically allocated semaphores can be initialized using the <i>K_KNSEM_INITIALIZER(cnt)</i> macro, where <i>cnt</i> is the initial positive value given to the semaphore counter. This macro is used as follows:</p> <pre>KnSem mySem = K_KNSEM_INITIALIZER(cnt) ;</pre> <p><i>semInit</i> is used to initialize a semaphore that is shared between two or more actors.</p> <p><i>semP</i> decrements the semaphore counter; if the counter reaches a strictly negative value, the calling thread is blocked, according to the options described by <i>waitLimit</i> in <i>intro(2K)</i>. <i>waitLimit</i> is a pointer to a <i>KnTimeVal</i> structure containing a timeout interval as described in <i>sysTime(2K)</i>.</p> <p><i>semV</i> increments the counter. If the new value is less than or equal to 0, the thread that has been blocked behind the semaphore for the longest time is awakened.</p> <p>As semaphore structures are allocated in the client memory, the number of semaphores used by an application is not limited. Any modification to the semaphore structure while the semaphore is in use will cause unpredictable synchronization behavior in the application.</p> <p>A blocking <i>semP</i> is <i>ABORTABLE</i> (see <i>threadAbort(2K)</i>).</p>
<b>RESTRICTIONS</b>	<p>A user application and a supervisor application may not share a semaphore.</p> <p>On the contrary, two user applications may share a semaphore by mapping it in both user address spaces.</p>
<b>RETURN VALUE</b>	Upon successful completion, a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	[K_EINVAL]                      The <i>waitLimit</i> is not a valid <i>KnTimeVal</i> .

[K\_EABORT] *semP* has been aborted.

[K\_ETIMEOUT] The timeout occurred.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`mutexInit(2K)`

<b>NAME</b>	svExcHandler, svAbortHandler – Define an exception handler; Define an abort handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svExcHandler(KnCap * actorcap, KnExcHdl routine);  int svAbortHandler(KnCap * actorcap, KnAbortHdl routine);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svExcHandler</i> function defines an exception handler for the actor whose capability is pointed to by <i>actorcap</i>. If <i>actorcap</i> is <code>K_MYACTOR</code>, the current actor is used. The <i>routine</i> field defines a handler that will be called every time a thread whose execution actor matches the specified actor encounters an exception (see <i>svTrapConnect</i> (2K) for details on execution actors). An exception handler is a function which takes three arguments:</p> <pre>int handler (ctx, excno, ptr)      KnThreadCtx*   ctx ;     int            excno ;     void*          ptr ;</pre> <p>The <i>ctx</i> pointer is the register context of the faulty thread. It points to a <i>KnThreadCtx</i> structure. The fields of this structure are machine-dependent.</p> <p>The <i>excno</i> field is a machine-dependent exception number. It can only be used for exception handlers.</p> <p>The <i>ptr</i> pointer highlight any relevant machine-dependent data that may be needed to process the exception (for example,. the faulty address for a page fault). It may not be used on certain hardware architectures.</p> <p>The exception handler may modify the faulty thread's register context before returning.</p> <p>The <i>svAbortHandler</i> function defines a handler to be called each time a thread of the specified actor returns from executing within its home actor environment, while in the <i>ABORTED</i> state. When the handler is called, the thread exits its <i>ABORTED</i> state (just as if <i>threadAborted</i> (2K) were called). An abort handler has the same form as an exception handler, except that the exception number is the <code>K_EABORT</code> constant, and no return value is taken into account.</p> <p>These calls are restricted to <i>SUPERVISOR</i> threads. Handlers are executed in <i>SUPERVISOR</i> execution mode. The handlers' code and data accessed must be parts of the locked-in-memory regions of a <i>SUPERVISOR</i> actor.</p> <p>These calls are used by subsystem process managers. The <i>svExcHandler</i> function is used to catch exceptions, and to translate them into subsystem-specific events</p>

(for example, . UNIX signals); *svAbortHandler* is used to force abort requests to be taken into account by controlled threads.

**RETURN VALUE**

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EINVAL] *actorcap* is an inconsistent actor capability.

[K\_EUNKNOWN] *actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*svTrapConnect(2K)*

<b>NAME</b>	svActorAbortHandler, svActorAbortHandlerConnect, svActorAbortHandlerDisconnect, svActorAbortHandlerGetConnected – Connect an actor abort handler; Disconnect an actor abort handler; Get an actor abort handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorAbortHandlerConnect(KnCap * actorcap, int vectindex, KnLapDesc * newhandler);  int svActorAbortHandlerDisconnect(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);  int svActorAbortHandlerGetConnected(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svActorAbortHandler</i> calls are used to manage abort handlers associated with the actor whose capability is pointed to by <i>actorcap</i> . These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>Abort handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>As several subsystem managers may need to control the aborted state of a thread, the Nucleus keeps a vector of <code>K_ACTOR_ABORTVECT_MAX</code> abort handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this abort handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is left to the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant <code>K_ACTOR_ABORTVECT_LAST</code> can be used independently of the system configuration to designate the index <code>K_ACTOR_ABORTVECT_MAX - 1</code>.</p> <p>When a thread returns from executing within its home actor environment while in the <i>ABORTED</i> state, the system invokes all abort handlers which have been associated with the thread's home actor (see <i>lapInvoke(2K)</i> ). The system clears the thread <i>ABORTED</i> state before calling the handlers (just as if <i>threadAborted(2K)</i> were called). Abort handlers are invoked in sequence, starting from index 0, until one of them succeeds in processing the abort event (sets the status to <code>K_ABORT_COMPLETED</code>).</p> <p>The argument of the lap abort handler is a pointer to a <code>KnActorAbortDesc</code> data structure. This structure has the following fields:</p> <pre>KnAbortStatus *abortStatus ; KnThreadCtx *threadCtx ;</pre>

The *abortStatus* field can be used by the abort handler to indicate that the abort has been successfully processed. The abort handler is called with *abortStatus* pointing to a status variable set to `K_ABORT_INPROGRESS`. The handler can indicate that it has completed the abort processing by setting this status variable to `K_ABORT_COMPLETED`. In this case no other abort handlers will be invoked.

The *threadCtx* field gives access to the register context of the aborted thread. This context is processor— dependent. Abort handlers can modify the register context of the aborted thread to recover from the abort.

The *svActorAbortHandlerConnect* function duplicates the lap descriptor pointed to by *newhandler* into the specified entry of the abort handlers vector.

The *svActorAbortHandlerDisconnect* function clears the lap descriptor stored in the specified entry of the abort handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

The *svActorAbortHandlerGetConnected* function copies the lap descriptor currently installed in the abort handlers vector at the location pointed to by *curhandler*.

**RETURN VALUE**

On success, the calls return `K_OK` otherwise a negative error code is returned.

**ERRORS**

- [`K_EBUSY`] *svActorAbortHandlerConnect* is called and there is already a valid lap descriptor installed in the specified entry of the abort handlers vector.
- [`K_EINVAL`] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorAbortHandlerDisconnect* is not equal to `K_CONNECTED_LAP` and does not match the lap descriptor currently installed.
- [`K_EUNKNOWN`] *actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*svLapCreate(2K)* , *lapInvoke(2K)*

<b>NAME</b>	svActorAbortHandler, svActorAbortHandlerConnect, svActorAbortHandlerDisconnect, svActorAbortHandlerGetConnected – Connect an actor abort handler; Disconnect an actor abort handler; Get an actor abort handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorAbortHandlerConnect(KnCap * actorcap, int vectindex, KnLapDesc * newhandler);  int svActorAbortHandlerDisconnect(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);  int svActorAbortHandlerGetConnected(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svActorAbortHandler</i> calls are used to manage abort handlers associated with the actor whose capability is pointed to by <i>actorcap</i> . These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>Abort handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>As several subsystem managers may need to control the aborted state of a thread, the Nucleus keeps a vector of <code>K_ACTOR_ABORTVECT_MAX</code> abort handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this abort handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is left to the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant <code>K_ACTOR_ABORTVECT_LAST</code> can be used independently of the system configuration to designate the index <code>K_ACTOR_ABORTVECT_MAX - 1</code>.</p> <p>When a thread returns from executing within its home actor environment while in the <i>ABORTED</i> state, the system invokes all abort handlers which have been associated with the thread's home actor (see <i>lapInvoke(2K)</i> ). The system clears the thread <i>ABORTED</i> state before calling the handlers (just as if <i>threadAborted(2K)</i> were called). Abort handlers are invoked in sequence, starting from index 0, until one of them succeeds in processing the abort event (sets the status to <code>K_ABORT_COMPLETED</code>).</p> <p>The argument of the lap abort handler is a pointer to a <code>KnActorAbortDesc</code> data structure. This structure has the following fields:</p> <pre>KnAbortStatus *abortStatus ; KnThreadCtx *threadCtx ;</pre>

The *abortStatus* field can be used by the abort handler to indicate that the abort has been successfully processed. The abort handler is called with *abortStatus* pointing to a status variable set to `K_ABORT_INPROGRESS`. The handler can indicate that it has completed the abort processing by setting this status variable to `K_ABORT_COMPLETED`. In this case no other abort handlers will be invoked.

The *threadCtx* field gives access to the register context of the aborted thread. This context is processor— dependent. Abort handlers can modify the register context of the aborted thread to recover from the abort.

The *svActorAbortHandlerConnect* function duplicates the lap descriptor pointed to by *newhandler* into the specified entry of the abort handlers vector.

The *svActorAbortHandlerDisconnect* function clears the lap descriptor stored in the specified entry of the abort handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

The *svActorAbortHandlerGetConnected* function copies the lap descriptor currently installed in the abort handlers vector at the location pointed to by *curhandler*.

**RETURN VALUE**

On success, the calls return `K_OK` otherwise a negative error code is returned.

**ERRORS**

- [`K_EBUSY`] *svActorAbortHandlerConnect* is called and there is already a valid lap descriptor installed in the specified entry of the abort handlers vector.
- [`K_EINVAL`] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorAbortHandlerDisconnect* is not equal to `K_CONNECTED_LAP` and does not match the lap descriptor currently installed.
- [`K_EUNKNOWN`] *actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*svLapCreate(2K)* , *lapInvoke(2K)*

<b>NAME</b>	svActorAbortHandler, svActorAbortHandlerConnect, svActorAbortHandlerDisconnect, svActorAbortHandlerGetConnected – Connect an actor abort handler; Disconnect an actor abort handler; Get an actor abort handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorAbortHandlerConnect(KnCap * actorcap, int vectindex, KnLapDesc * newhandler);  int svActorAbortHandlerDisconnect(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);  int svActorAbortHandlerGetConnected(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svActorAbortHandler</i> calls are used to manage abort handlers associated with the actor whose capability is pointed to by <i>actorcap</i> . These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>Abort handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>As several subsystem managers may need to control the aborted state of a thread, the Nucleus keeps a vector of <code>K_ACTOR_ABORTVECT_MAX</code> abort handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this abort handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is left to the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant <code>K_ACTOR_ABORTVECT_LAST</code> can be used independently of the system configuration to designate the index <code>K_ACTOR_ABORTVECT_MAX - 1</code>.</p> <p>When a thread returns from executing within its home actor environment while in the <i>ABORTED</i> state, the system invokes all abort handlers which have been associated with the thread's home actor (see <i>lapInvoke(2K)</i> ). The system clears the thread <i>ABORTED</i> state before calling the handlers (just as if <i>threadAborted(2K)</i> were called). Abort handlers are invoked in sequence, starting from index 0, until one of them succeeds in processing the abort event (sets the status to <code>K_ABORT_COMPLETED</code>).</p> <p>The argument of the lap abort handler is a pointer to a <code>KnActorAbortDesc</code> data structure. This structure has the following fields:</p> <pre>KnAbortStatus *abortStatus ; KnThreadCtx *threadCtx ;</pre>

The *abortStatus* field can be used by the abort handler to indicate that the abort has been successfully processed. The abort handler is called with *abortStatus* pointing to a status variable set to `K_ABORT_INPROGRESS`. The handler can indicate that it has completed the abort processing by setting this status variable to `K_ABORT_COMPLETED`. In this case no other abort handlers will be invoked.

The *threadCtx* field gives access to the register context of the aborted thread. This context is processor— dependent. Abort handlers can modify the register context of the aborted thread to recover from the abort.

The *svActorAbortHandlerConnect* function duplicates the lap descriptor pointed to by *newhandler* into the specified entry of the abort handlers vector.

The *svActorAbortHandlerDisconnect* function clears the lap descriptor stored in the specified entry of the abort handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

The *svActorAbortHandlerGetConnected* function copies the lap descriptor currently installed in the abort handlers vector at the location pointed to by *curhandler*.

**RETURN VALUE**

On success, the calls return `K_OK` otherwise a negative error code is returned.

**ERRORS**

- [`K_EBUSY`] *svActorAbortHandlerConnect* is called and there is already a valid lap descriptor installed in the specified entry of the abort handlers vector.
- [`K_EINVAL`] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorAbortHandlerDisconnect* is not equal to `K_CONNECTED_LAP` and does not match the lap descriptor currently installed.
- [`K_EUNKNOWN`] *actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*svLapCreate(2K)* , *lapInvoke(2K)*

<b>NAME</b>	svActorAbortHandler, svActorAbortHandlerConnect, svActorAbortHandlerDisconnect, svActorAbortHandlerGetConnected – Connect an actor abort handler; Disconnect an actor abort handler; Get an actor abort handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorAbortHandlerConnect(KnCap * actorcap, int vectindex, KnLapDesc * newhandler);  int svActorAbortHandlerDisconnect(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);  int svActorAbortHandlerGetConnected(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svActorAbortHandler</i> calls are used to manage abort handlers associated with the actor whose capability is pointed to by <i>actorcap</i> . These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>Abort handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>As several subsystem managers may need to control the aborted state of a thread, the Nucleus keeps a vector of <code>K_ACTOR_ABORTVECT_MAX</code> abort handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this abort handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is left to the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant <code>K_ACTOR_ABORTVECT_LAST</code> can be used independently of the system configuration to designate the index <code>K_ACTOR_ABORTVECT_MAX - 1</code>.</p> <p>When a thread returns from executing within its home actor environment while in the <i>ABORTED</i> state, the system invokes all abort handlers which have been associated with the thread's home actor (see <i>lapInvoke(2K)</i> ). The system clears the thread <i>ABORTED</i> state before calling the handlers (just as if <i>threadAborted(2K)</i> were called). Abort handlers are invoked in sequence, starting from index 0, until one of them succeeds in processing the abort event (sets the status to <code>K_ABORT_COMPLETED</code>).</p> <p>The argument of the lap abort handler is a pointer to a <code>KnActorAbortDesc</code> data structure. This structure has the following fields:</p> <pre>KnAbortStatus *abortStatus ; KnThreadCtx *threadCtx ;</pre>

The *abortStatus* field can be used by the abort handler to indicate that the abort has been successfully processed. The abort handler is called with *abortStatus* pointing to a status variable set to `K_ABORT_INPROGRESS`. The handler can indicate that it has completed the abort processing by setting this status variable to `K_ABORT_COMPLETED`. In this case no other abort handlers will be invoked.

The *threadCtx* field gives access to the register context of the aborted thread. This context is processor— dependent. Abort handlers can modify the register context of the aborted thread to recover from the abort.

The *svActorAbortHandlerConnect* function duplicates the lap descriptor pointed to by *newhandler* into the specified entry of the abort handlers vector.

The *svActorAbortHandlerDisconnect* function clears the lap descriptor stored in the specified entry of the abort handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

The *svActorAbortHandlerGetConnected* function copies the lap descriptor currently installed in the abort handlers vector at the location pointed to by *curhandler*.

**RETURN VALUE**

On success, the calls return `K_OK` otherwise a negative error code is returned.

**ERRORS**

- [`K_EBUSY`] *svActorAbortHandlerConnect* is called and there is already a valid lap descriptor installed in the specified entry of the abort handlers vector.
- [`K_EINVAL`] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorAbortHandlerDisconnect* is not equal to `K_CONNECTED_LAP` and does not match the lap descriptor currently installed.
- [`K_EUNKNOWN`] *actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*svLapCreate(2K)* , *lapInvoke(2K)*

<b>NAME</b>	svActorExceptionHandler, svActorExceptionHandlerConnect, svActorExceptionHandlerDisconnect, svActorExceptionHandlerGetConnected – Connect an actor exception handler; Disconnect an actor exception handler; Get an actor exception handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorExceptionHandlerConnect(KnCap * actorcap, int vectindex, KnLapDesc * newhandler);  int svActorExceptionHandlerDisconnect(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);  int svActorExceptionHandlerGetConnected(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svActorExceptionHandler</i> calls are used to manage exception handlers associated with the actor whose capability is pointed to by <i>actorcap</i> . These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>Exception handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>As several subsystem managers may need to control exceptions which occur while a thread is executing within a given actor, the system keeps a vector of <code>K_ACTOR_EXCVECT_MAX</code> exception handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this exception handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant <code>K_ACTOR_EXCVECT_LAST</code> can be used independently of the system configuration to designate the index <code>K_ACTOR_EXCVECT_MAX - 1</code>.</p> <p>When a thread causes an exception, the system uses the exception handlers vector associated with the current execution actor of the faulting thread to recover from the exception. During this recovery phase, all valid exception handlers are invoked in sequence (see <i>lapInvoke(2K)</i> ), starting from index 0, until one of them succeeds in recovering from the exception (sets the status to <code>K_EXC_COMPLETED</code>).</p> <p>If the exception is not recovered after calling all the exception handlers, the system starts the termination phase where all exception handlers are called a second time in reverse order (starting from index <code>K_ACTOR_EXCVECT_MAX - 1</code>) until one of them takes a default action on the faulting thread (sets the <code>K_EXC_COMPLETED</code> status described below).</p>

The argument of the lap exception handler is a pointer to a `KnActorExcDesc` data structure. This structure has the following fields:

```
KnExcPhase excPhase ;
KnExcStatus *excStatus ;
int excNumber ;
KnThreadCtx *threadCtx ;
void *exceptionCtx ;
```

The *excPhase* field indicates the phase in which the handler was invoked, either `K_EXC_RECOVER` (recovery phase) or `K_EXC_TERMINATE` (termination phase).

The *excStatus* field can be used by the exception handler to indicate that the exception has been processed successfully. The exception handler is called with *excStatus* pointing to a status variable set to `K_EXC_INPROGRESS`. The handler can indicate that it has completed the exception processing by setting this status variable to `K_EXC_COMPLETED`. In this case, no other exception handlers will be invoked.

The *excNumber* field is a processor—dependent exception number.

The *threadCtx* field gives access to the register context of the faulting thread. This context is processor—dependent. Exception handlers may modify the register context of the faulting thread to recover from the exception.

The *exceptionCtx* field gives access to additional data specific to the exception (for example, the fault address for a page fault). This context is processor— and exception— dependent.

The *svActorExcHandlerConnect* function duplicates the lap descriptor pointed to by *newhandler* into the entry specified in the exception handlers vector.

The *svActorExcHandlerDisconnect* function clears the lap descriptor stored in the specified entry of the exception handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

The *svActorExcHandlerGetConnected* function copies the lap descriptor currently installed in the exception handlers vector to the location pointed to by *curhandler*.

## RETURN VALUE

On success, the calls return `K_OK` otherwise a negative error code is returned.

## ERRORS

[`K_EBUSY`] *svActorExcHandlerConnect* is called and there is already a valid lap descriptor installed in the entry specified of the exception handlers vector.

[K\_EINVAL] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorExcHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

[K\_EUNKNOWN] *actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svLapCreate(2K)` , `lapInvoke(2K)`

<b>NAME</b>	svActorExcHandler, svActorExcHandlerConnect, svActorExcHandlerDisconnect, svActorExcHandlerGetConnected – Connect an actor exception handler; Disconnect an actor exception handler; Get an actor exception handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorExcHandlerConnect(KnCap * actorcap, int vectindex, KnLapDesc * newhandler);  int svActorExcHandlerDisconnect(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);  int svActorExcHandlerGetConnected(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svActorExcHandler</i> calls are used to manage exception handlers associated with the actor whose capability is pointed to by <i>actorcap</i>. These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>Exception handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>As several subsystem managers may need to control exceptions which occur while a thread is executing within a given actor, the system keeps a vector of <i>K_ACTOR_EXCVECT_MAX</i> exception handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this exception handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant <i>K_ACTOR_EXCVECT_LAST</i> can be used independently of the system configuration to designate the index <i>K_ACTOR_EXCVECT_MAX - 1</i>.</p> <p>When a thread causes an exception, the system uses the exception handlers vector associated with the current execution actor of the faulting thread to recover from the exception. During this recovery phase, all valid exception handlers are invoked in sequence (see <i>lapInvoke(2K)</i>), starting from index 0, until one of them succeeds in recovering from the exception (sets the status to <i>K_EXC_COMPLETED</i>).</p> <p>If the exception is not recovered after calling all the exception handlers, the system starts the termination phase where all exception handlers are called a second time in reverse order (starting from index <i>K_ACTOR_EXCVECT_MAX - 1</i>) until one of them takes a default action on the faulting thread (sets the <i>K_EXC_COMPLETED</i> status described below).</p>

The argument of the lap exception handler is a pointer to a `KnActorExcDesc` data structure. This structure has the following fields:

```
KnExcPhase excPhase ;
KnExcStatus *excStatus ;
int excNumber ;
KnThreadCtx *threadCtx ;
void *exceptionCtx ;
```

The *excPhase* field indicates the phase in which the handler was invoked, either `K_EXC_RECOVER` (recovery phase) or `K_EXC_TERMINATE` (termination phase).

The *excStatus* field can be used by the exception handler to indicate that the exception has been processed successfully. The exception handler is called with *excStatus* pointing to a status variable set to `K_EXC_INPROGRESS`. The handler can indicate that it has completed the exception processing by setting this status variable to `K_EXC_COMPLETED`. In this case, no other exception handlers will be invoked.

The *excNumber* field is a processor—dependent exception number.

The *threadCtx* field gives access to the register context of the faulting thread. This context is processor—dependent. Exception handlers may modify the register context of the faulting thread to recover from the exception.

The *exceptionCtx* field gives access to additional data specific to the exception (for example, the fault address for a page fault). This context is processor— and exception— dependent.

The *svActorExcHandlerConnect* function duplicates the lap descriptor pointed to by *newhandler* into the entry specified in the exception handlers vector.

The *svActorExcHandlerDisconnect* function clears the lap descriptor stored in the specified entry of the exception handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

The *svActorExcHandlerGetConnected* function copies the lap descriptor currently installed in the exception handlers vector to the location pointed to by *curhandler*.

## RETURN VALUE

On success, the calls return `K_OK` otherwise a negative error code is returned.

## ERRORS

[`K_EBUSY`] *svActorExcHandlerConnect* is called and there is already a valid lap descriptor installed in the entry specified of the exception handlers vector.

[K\_EINVAL]

*actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorExcHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

[K\_EUNKNOWN]

*actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*svLapCreate(2K)* , *lapInvoke(2K)*

<b>NAME</b>	svActorExcHandler, svActorExcHandlerConnect, svActorExcHandlerDisconnect, svActorExcHandlerGetConnected – Connect an actor exception handler; Disconnect an actor exception handler; Get an actor exception handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorExcHandlerConnect(KnCap * actorcap, int vectindex, KnLapDesc * newhandler);  int svActorExcHandlerDisconnect(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);  int svActorExcHandlerGetConnected(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svActorExcHandler</i> calls are used to manage exception handlers associated with the actor whose capability is pointed to by <i>actorcap</i> . These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>Exception handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>As several subsystem managers may need to control exceptions which occur while a thread is executing within a given actor, the system keeps a vector of <code>K_ACTOR_EXCVECT_MAX</code> exception handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this exception handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant <code>K_ACTOR_EXCVECT_LAST</code> can be used independently of the system configuration to designate the index <code>K_ACTOR_EXCVECT_MAX - 1</code>.</p> <p>When a thread causes an exception, the system uses the exception handlers vector associated with the current execution actor of the faulting thread to recover from the exception. During this recovery phase, all valid exception handlers are invoked in sequence (see <i>lapInvoke(2K)</i> ), starting from index 0, until one of them succeeds in recovering from the exception (sets the status to <code>K_EXC_COMPLETED</code>).</p> <p>If the exception is not recovered after calling all the exception handlers, the system starts the termination phase where all exception handlers are called a second time in reverse order (starting from index <code>K_ACTOR_EXCVECT_MAX - 1</code>) until one of them takes a default action on the faulting thread (sets the <code>K_EXC_COMPLETED</code> status described below).</p>

The argument of the lap exception handler is a pointer to a `KnActorExcDesc` data structure. This structure has the following fields:

```
KnExcPhase excPhase ;
KnExcStatus *excStatus ;
int excNumber ;
KnThreadCtx *threadCtx ;
void *exceptionCtx ;
```

The *excPhase* field indicates the phase in which the handler was invoked, either `K_EXC_RECOVER` (recovery phase) or `K_EXC_TERMINATE` (termination phase).

The *excStatus* field can be used by the exception handler to indicate that the exception has been processed successfully. The exception handler is called with *excStatus* pointing to a status variable set to `K_EXC_INPROGRESS`. The handler can indicate that it has completed the exception processing by setting this status variable to `K_EXC_COMPLETED`. In this case, no other exception handlers will be invoked.

The *excNumber* field is a processor—dependent exception number.

The *threadCtx* field gives access to the register context of the faulting thread. This context is processor—dependent. Exception handlers may modify the register context of the faulting thread to recover from the exception.

The *exceptionCtx* field gives access to additional data specific to the exception (for example, the fault address for a page fault). This context is processor— and exception— dependent.

The *svActorExcHandlerConnect* function duplicates the lap descriptor pointed to by *newhandler* into the entry specified in the exception handlers vector.

The *svActorExcHandlerDisconnect* function clears the lap descriptor stored in the specified entry of the exception handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

The *svActorExcHandlerGetConnected* function copies the lap descriptor currently installed in the exception handlers vector to the location pointed to by *curhandler*.

## RETURN VALUE

On success, the calls return `K_OK` otherwise a negative error code is returned.

## ERRORS

[`K_EBUSY`] *svActorExcHandlerConnect* is called and there is already a valid lap descriptor installed in the entry specified of the exception handlers vector.

[K\_EINVAL] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorExcHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

[K\_EUNKNOWN] *actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svLapCreate(2K)` , `lapInvoke(2K)`

<b>NAME</b>	svActorExcHandler, svActorExcHandlerConnect, svActorExcHandlerDisconnect, svActorExcHandlerGetConnected – Connect an actor exception handler; Disconnect an actor exception handler; Get an actor exception handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorExcHandlerConnect(KnCap * actorcap, int vectindex, KnLapDesc * newhandler);  int svActorExcHandlerDisconnect(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);  int svActorExcHandlerGetConnected(KnCap * actorcap, int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svActorExcHandler</i> calls are used to manage exception handlers associated with the actor whose capability is pointed to by <i>actorcap</i>. These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>Exception handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>As several subsystem managers may need to control exceptions which occur while a thread is executing within a given actor, the system keeps a vector of <i>K_ACTOR_EXCVECT_MAX</i> exception handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this exception handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant <i>K_ACTOR_EXCVECT_LAST</i> can be used independently of the system configuration to designate the index <i>K_ACTOR_EXCVECT_MAX - 1</i>.</p> <p>When a thread causes an exception, the system uses the exception handlers vector associated with the current execution actor of the faulting thread to recover from the exception. During this recovery phase, all valid exception handlers are invoked in sequence (see <i>lapInvoke(2K)</i>), starting from index 0, until one of them succeeds in recovering from the exception (sets the status to <i>K_EXC_COMPLETED</i>).</p> <p>If the exception is not recovered after calling all the exception handlers, the system starts the termination phase where all exception handlers are called a second time in reverse order (starting from index <i>K_ACTOR_EXCVECT_MAX - 1</i>) until one of them takes a default action on the faulting thread (sets the <i>K_EXC_COMPLETED</i> status described below).</p>

The argument of the lap exception handler is a pointer to a `KnActorExcDesc` data structure. This structure has the following fields:

```
KnExcPhase excPhase ;
KnExcStatus *excStatus ;
int excNumber ;
KnThreadCtx *threadCtx ;
void *exceptionCtx ;
```

The *excPhase* field indicates the phase in which the handler was invoked, either `K_EXC_RECOVER` (recovery phase) or `K_EXC_TERMINATE` (termination phase).

The *excStatus* field can be used by the exception handler to indicate that the exception has been processed successfully. The exception handler is called with *excStatus* pointing to a status variable set to `K_EXC_INPROGRESS`. The handler can indicate that it has completed the exception processing by setting this status variable to `K_EXC_COMPLETED`. In this case, no other exception handlers will be invoked.

The *excNumber* field is a processor—dependent exception number.

The *threadCtx* field gives access to the register context of the faulting thread. This context is processor—dependent. Exception handlers may modify the register context of the faulting thread to recover from the exception.

The *exceptionCtx* field gives access to additional data specific to the exception (for example, the fault address for a page fault). This context is processor— and exception— dependent.

The *svActorExcHandlerConnect* function duplicates the lap descriptor pointed to by *newhandler* into the entry specified in the exception handlers vector.

The *svActorExcHandlerDisconnect* function clears the lap descriptor stored in the specified entry of the exception handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

The *svActorExcHandlerGetConnected* function copies the lap descriptor currently installed in the exception handlers vector to the location pointed to by *curhandler*.

## RETURN VALUE

On success, the calls return `K_OK` otherwise a negative error code is returned.

## ERRORS

[`K_EBUSY`] *svActorExcHandlerConnect* is called and there is already a valid lap descriptor installed in the entry specified of the exception handlers vector.

[K\_EINVAL]

*actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorExcHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

[K\_EUNKNOWN]

*actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*svLapCreate(2K)* , *lapInvoke(2K)*

<b>NAME</b>	svActorStopHandler, svActorStopHandlerConnect, svActorStopHandlerDisconnect, svActorStopHandlerGetConnected – Actor stop handler management: Connect an actor stop handler; Disconnect an actor stop handler; Get an actor stop handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorStopHandlerConnect(KnCap * actorcap, unsigned int vectindex, KnLapDesc * newhandler);  int svActorStopHandlerDisconnect(KnCap * actorcap, unsigned int vectindex, KnLapDesc * curhandler);  int svActorStopHandlerGetConnected(KnCap * actorcap, unsigned int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>svActorStopHandler</i> calls are used to manage stop handlers associated with the actor whose capability is pointed to by <i>actorcap</i> . These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>The stop handlers are invoked as a consequence of a <i>threadStop</i> (2K) or an <i>actorStop</i> (2K) system call. In the last case, they are called for all the threads that appear in the actor being stopped. If a thread has performed a safe lap in this actor, it has at least one lap frame in this actor; in such conditions, the stop handlers are called for every lap frame that appears in the actor.</p> <p>Stop handlers are specified in the form of lap descriptors (see <i>svLapCreate</i>(2K)).</p> <p>As several subsystem managers may need to control the stopped state of a thread, the Nucleus keeps a vector of K_ACTOR_STOPVECT_MAX stop handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this stop handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is left to the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant K_ACTOR_STOPVECT_LAST can be used independently of the system configuration to designate the index K_ACTOR_STOPVECT_MAX - 1.</p> <p>The stop handlers are usually called by the stopped thread just before it actually stops (K_STOP_INCTX). However, if the thread being stopped is blocked in a blocking system call, or if it has left the current stopped actor through a <i>lapInvoke</i> (2K) system call, stop handlers are called by the thread which performs the stop action (K_STOP_OUTCTX).</p>

In any case stop handlers are invoked in sequence, starting from index 0, until one of them succeeds to process the stop event (for example, sets the `K_STOP_COMPLETED` status).

The argument of the lap stop handler is a pointer to a `KnActorStopDesc` data structure. This structure has the following fields:

```
KnStopMode      stopMode ;
KnStopStatus*   stopStatus ;
int             topFrameLevel ;
int             curFrameLevel ;
KnStopThStatus* threadStatus ;
KnThreadLid     threadLid ;
KnThreadCtx*    threadCtx ;
```

The `stopMode` field indicates whether or not the stop handlers are called in the context of the stopping thread (`K_STOP_INCTX` vs `K_STOP_OUTCTX`).

The `stopStatus` field may be used by the stop handler to indicate that the stop has been successfully processed. The stop handler is called with `stopStatus` pointing to a status variable set to `K_STOP_INPROGRESS`. The handler can indicate that it has completed the stop processing by setting this status variable to `K_STOP_COMPLETED`. In this case no other stop handlers will be invoked.

The `topFrameLevel` and the `curFrameLevel` fields are used to represent respectively the top and the current lap frames. The top lap frame is the one identified by the highest lap frame level (see `svLapCreate (2K)`).

The `threadStatus` field is used to specify whether or not the thread must be stopped when returning from the stop handler invocation. The corresponding values are `K_STOP_THREADSTART` and `K_STOP_THREADSTOP`.

The `threadLid` field describes the thread for which the stopping action must be performed.

The `threadCtx` field gives access to the register context of the stopped thread. This context is processor dependent. The given context does not depend on the invocation mode (`K_STOP_INCTX` vs `K_STOP_OUTCTX`); this is always the valid thread's context in the stopped actor.

`svActorStopHandlerConnect` duplicates the lap descriptor pointed to by `newhandler` into the specified entry of the stop handlers vector.

`svActorStopHandlerDisconnect` clears the lap descriptor stored in the specified entry of the stop handlers vector. If `curhandler` is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

`svActorStopHandlerGetConnected` copies at the location pointed to by `curhandler` the lap descriptor currently installed in the stop handlers vector.

**RETURN VALUE** On success, the calls return K\_OK otherwise a negative error code is returned.

**ERRORS**

- [K\_EBUSY] *svActorStopHandlerConnect* is called and there is already a valid lap descriptor installed in the specified entry of the stop handlers vector.
- [K\_EINVAL] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorStopHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.
- [K\_EUNKNOWN] *actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`actorStop(2K)` , `lapInvoke(2K)` , `svLapCreate(2K)` , `threadStop(2K)`

<b>NAME</b>	svActorStopHandler, svActorStopHandlerConnect, svActorStopHandlerDisconnect, svActorStopHandlerGetConnected – Actor stop handler management: Connect an actor stop handler; Disconnect an actor stop handler; Get an actor stop handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorStopHandlerConnect(KnCap * actorcap, unsigned int vectindex, KnLapDesc * newhandler);  int svActorStopHandlerDisconnect(KnCap * actorcap, unsigned int vectindex, KnLapDesc * curhandler);  int svActorStopHandlerGetConnected(KnCap * actorcap, unsigned int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>svActorStopHandler</i> calls are used to manage stop handlers associated with the actor whose capability is pointed to by <i>actorcap</i>. These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>The stop handlers are invoked as a consequence of a <i>threadStop</i> (2K) or an <i>actorStop</i> (2K) system call. In the last case, they are called for all the threads that appear in the actor being stopped. If a thread has performed a safe lap in this actor, it has at least one lap frame in this actor; in such conditions, the stop handlers are called for every lap frame that appears in the actor.</p> <p>Stop handlers are specified in the form of lap descriptors (see <i>svLapCreate</i>(2K)).</p> <p>As several subsystem managers may need to control the stopped state of a thread, the Nucleus keeps a vector of K_ACTOR_STOPVECT_MAX stop handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this stop handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is left to the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant K_ACTOR_STOPVECT_LAST can be used independently of the system configuration to designate the index K_ACTOR_STOPVECT_MAX - 1.</p> <p>The stop handlers are usually called by the stopped thread just before it actually stops (K_STOP_INCTX). However, if the thread being stopped is blocked in a blocking system call, or if it has left the current stopped actor through a <i>lapInvoke</i> (2K) system call, stop handlers are called by the thread which performs the stop action (K_STOP_OUTCTX).</p>

In any case stop handlers are invoked in sequence, starting from index 0, until one of them succeeds to process the stop event (for example, sets the `K_STOP_COMPLETED` status).

The argument of the lap stop handler is a pointer to a `KnActorStopDesc` data structure. This structure has the following fields:

```
KnStopMode      stopMode ;
KnStopStatus*   stopStatus ;
int             topFrameLevel ;
int             curFrameLevel ;
KnStopThStatus* threadStatus ;
KnThreadLid     threadLid ;
KnThreadCtx*    threadCtx ;
```

The *stopMode* field indicates whether or not the stop handlers are called in the context of the stopping thread (`K_STOP_INCTX` vs `K_STOP_OUTCTX`).

The *stopStatus* field may be used by the stop handler to indicate that the stop has been successfully processed. The stop handler is called with *stopStatus* pointing to a status variable set to `K_STOP_INPROGRESS`. The handler can indicate that it has completed the stop processing by setting this status variable to `K_STOP_COMPLETED`. In this case no other stop handlers will be invoked.

The *topFrameLevel* and the *curFrameLevel* fields are used to represent respectively the top and the current lap frames. The top lap frame is the one identified by the highest lap frame level (see *svLapCreate* (2K)).

The *threadStatus* field is used to specify whether or not the thread must be stopped when returning from the stop handler invocation. The corresponding values are `K_STOP_THREADSTART` and `K_STOP_THREADSTOP`.

The *threadLid* field describes the thread for which the stopping action must be performed.

The *threadCtx* field gives access to the register context of the stopped thread. This context is processor dependent. The given context does not depend on the invocation mode (`K_STOP_INCTX` vs `K_STOP_OUTCTX`); this is always the valid thread's context in the stopped actor.

*svActorStopHandlerConnect* duplicates the lap descriptor pointed to by *newhandler* into the specified entry of the stop handlers vector.

*svActorStopHandlerDisconnect* clears the lap descriptor stored in the specified entry of the stop handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

*svActorStopHandlerGetConnected* copies at the location pointed to by *curhandler* the lap descriptor currently installed in the stop handlers vector.

**RETURN VALUE** | On success, the calls return K\_OK otherwise a negative error code is returned.

**ERRORS** | [K\_EBUSY] *svActorStopHandlerConnect* is called and there is already a valid lap descriptor installed in the specified entry of the stop handlers vector.

[K\_EINVAL] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorStopHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

[K\_EUNKNOWN] *actorcap* does not specify a local actor.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | `actorStop(2K)`, `lapInvoke(2K)`, `svLapCreate(2K)`, `threadStop(2K)`

<b>NAME</b>	svActorStopHandler, svActorStopHandlerConnect, svActorStopHandlerDisconnect, svActorStopHandlerGetConnected – Actor stop handler management: Connect an actor stop handler; Disconnect an actor stop handler; Get an actor stop handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorStopHandlerConnect(KnCap * actorcap, unsigned int vectindex, KnLapDesc * newhandler);  int svActorStopHandlerDisconnect(KnCap * actorcap, unsigned int vectindex, KnLapDesc * curhandler);  int svActorStopHandlerGetConnected(KnCap * actorcap, unsigned int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>svActorStopHandler</i> calls are used to manage stop handlers associated with the actor whose capability is pointed to by <i>actorcap</i> . These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>The stop handlers are invoked as a consequence of a <i>threadStop</i> (2K) or an <i>actorStop</i> (2K) system call. In the last case, they are called for all the threads that appear in the actor being stopped. If a thread has performed a safe lap in this actor, it has at least one lap frame in this actor; in such conditions, the stop handlers are called for every lap frame that appears in the actor.</p> <p>Stop handlers are specified in the form of lap descriptors (see <i>svLapCreate</i>(2K)).</p> <p>As several subsystem managers may need to control the stopped state of a thread, the Nucleus keeps a vector of K_ACTOR_STOPVECT_MAX stop handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this stop handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is left to the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant K_ACTOR_STOPVECT_LAST can be used independently of the system configuration to designate the index K_ACTOR_STOPVECT_MAX - 1.</p> <p>The stop handlers are usually called by the stopped thread just before it actually stops (K_STOP_INCTX). However, if the thread being stopped is blocked in a blocking system call, or if it has left the current stopped actor through a <i>lapInvoke</i> (2K) system call, stop handlers are called by the thread which performs the stop action (K_STOP_OUTCTX).</p>

In any case stop handlers are invoked in sequence, starting from index 0, until one of them succeeds to process the stop event (for example, sets the `K_STOP_COMPLETED` status).

The argument of the lap stop handler is a pointer to a `KnActorStopDesc` data structure. This structure has the following fields:

```
KnStopMode      stopMode ;
KnStopStatus*   stopStatus ;
int             topFrameLevel ;
int             curFrameLevel ;
KnStopThStatus* threadStatus ;
KnThreadLid     threadLid ;
KnThreadCtx*    threadCtx ;
```

The `stopMode` field indicates whether or not the stop handlers are called in the context of the stopping thread (`K_STOP_INCTX` vs `K_STOP_OUTCTX`).

The `stopStatus` field may be used by the stop handler to indicate that the stop has been successfully processed. The stop handler is called with `stopStatus` pointing to a status variable set to `K_STOP_INPROGRESS`. The handler can indicate that it has completed the stop processing by setting this status variable to `K_STOP_COMPLETED`. In this case no other stop handlers will be invoked.

The `topFrameLevel` and the `curFrameLevel` fields are used to represent respectively the top and the current lap frames. The top lap frame is the one identified by the highest lap frame level (see `svLapCreate (2K)`).

The `threadStatus` field is used to specify whether or not the thread must be stopped when returning from the stop handler invocation. The corresponding values are `K_STOP_THREADSTART` and `K_STOP_THREADSTOP`.

The `threadLid` field describes the thread for which the stopping action must be performed.

The `threadCtx` field gives access to the register context of the stopped thread. This context is processor dependent. The given context does not depend on the invocation mode (`K_STOP_INCTX` vs `K_STOP_OUTCTX`); this is always the valid thread's context in the stopped actor.

`svActorStopHandlerConnect` duplicates the lap descriptor pointed to by `newhandler` into the specified entry of the stop handlers vector.

`svActorStopHandlerDisconnect` clears the lap descriptor stored in the specified entry of the stop handlers vector. If `curhandler` is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

`svActorStopHandlerGetConnected` copies at the location pointed to by `curhandler` the lap descriptor currently installed in the stop handlers vector.

**RETURN VALUE** On success, the calls return K\_OK otherwise a negative error code is returned.

**ERRORS**

[K\_EBUSY] *svActorStopHandlerConnect* is called and there is already a valid lap descriptor installed in the specified entry of the stop handlers vector.

[K\_EINVAL] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorStopHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

[K\_EUNKNOWN] *actorcap* does not specify a local actor.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `actorStop(2K)` , `lapInvoke(2K)` , `svLapCreate(2K)` , `threadStop(2K)`

<b>NAME</b>	svActorStopHandler, svActorStopHandlerConnect, svActorStopHandlerDisconnect, svActorStopHandlerGetConnected – Actor stop handler management: Connect an actor stop handler; Disconnect an actor stop handler; Get an actor stop handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svActorStopHandlerConnect(KnCap * actorcap, unsigned int vectindex, KnLapDesc * newhandler);  int svActorStopHandlerDisconnect(KnCap * actorcap, unsigned int vectindex, KnLapDesc * curhandler);  int svActorStopHandlerGetConnected(KnCap * actorcap, unsigned int vectindex, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>svActorStopHandler</i> calls are used to manage stop handlers associated with the actor whose capability is pointed to by <i>actorcap</i>. These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>The stop handlers are invoked as a consequence of a <i>threadStop</i> (2K) or an <i>actorStop</i> (2K) system call. In the last case, they are called for all the threads that appear in the actor being stopped. If a thread has performed a safe lap in this actor, it has at least one lap frame in this actor; in such conditions, the stop handlers are called for every lap frame that appears in the actor.</p> <p>Stop handlers are specified in the form of lap descriptors (see <i>svLapCreate</i>(2K)).</p> <p>As several subsystem managers may need to control the stopped state of a thread, the Nucleus keeps a vector of K_ACTOR_STOPVECT_MAX stop handlers per actor.</p> <p>The <i>vectindex</i> parameter identifies which entry of this stop handlers vector should be used by the calls. The consistent allocation of vector indexes among subsystem managers is left to the responsibility of the subsystem managers. Depending on the subsystems, the last vector entry may be reserved for use by the application. For that purpose, the constant K_ACTOR_STOPVECT_LAST can be used independently of the system configuration to designate the index K_ACTOR_STOPVECT_MAX - 1.</p> <p>The stop handlers are usually called by the stopped thread just before it actually stops (K_STOP_INCTX). However, if the thread being stopped is blocked in a blocking system call, or if it has left the current stopped actor through a <i>lapInvoke</i> (2K) system call, stop handlers are called by the thread which performs the stop action (K_STOP_OUTCTX).</p>

In any case stop handlers are invoked in sequence, starting from index 0, until one of them succeeds to process the stop event (for example, sets the `K_STOP_COMPLETED` status).

The argument of the lap stop handler is a pointer to a `KnActorStopDesc` data structure. This structure has the following fields:

```
KnStopMode      stopMode ;
KnStopStatus*   stopStatus ;
int             topFrameLevel ;
int             curFrameLevel ;
KnStopThStatus* threadStatus ;
KnThreadLid     threadLid ;
KnThreadCtx*    threadCtx ;
```

The *stopMode* field indicates whether or not the stop handlers are called in the context of the stopping thread (`K_STOP_INCTX` vs `K_STOP_OUTCTX`).

The *stopStatus* field may be used by the stop handler to indicate that the stop has been successfully processed. The stop handler is called with *stopStatus* pointing to a status variable set to `K_STOP_INPROGRESS`. The handler can indicate that it has completed the stop processing by setting this status variable to `K_STOP_COMPLETED`. In this case no other stop handlers will be invoked.

The *topFrameLevel* and the *curFrameLevel* fields are used to represent respectively the top and the current lap frames. The top lap frame is the one identified by the highest lap frame level (see *svLapCreate* (2K)).

The *threadStatus* field is used to specify whether or not the thread must be stopped when returning from the stop handler invocation. The corresponding values are `K_STOP_THREADSTART` and `K_STOP_THREADSTOP`.

The *threadLid* field describes the thread for which the stopping action must be performed.

The *threadCtx* field gives access to the register context of the stopped thread. This context is processor dependent. The given context does not depend on the invocation mode (`K_STOP_INCTX` vs `K_STOP_OUTCTX`); this is always the valid thread's context in the stopped actor.

*svActorStopHandlerConnect* duplicates the lap descriptor pointed to by *newhandler* into the specified entry of the stop handlers vector.

*svActorStopHandlerDisconnect* clears the lap descriptor stored in the specified entry of the stop handlers vector. If *curhandler* is not equal to `K_CONNECTED_LAP`, it must point to a lap descriptor identical to the lap descriptor currently installed.

*svActorStopHandlerGetConnected* copies at the location pointed to by *curhandler* the lap descriptor currently installed in the stop handlers vector.

**RETURN VALUE** On success, the calls return K\_OK otherwise a negative error code is returned.

**ERRORS**

[K\_EBUSY] *svActorStopHandlerConnect* is called and there is already a valid lap descriptor installed in the specified entry of the stop handlers vector.

[K\_EINVAL] *actorcap* is an inconsistent actor capability or *vectindex* is invalid. The *curhandler* lap descriptor specified to *svActorStopHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

[K\_EUNKNOWN] *actorcap* does not specify a local actor.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `actorStop(2K)`, `lapInvoke(2K)`, `svLapCreate(2K)`, `threadStop(2K)`

<b>NAME</b>	svActorVirtualTimeout, svActorVirtualTimeoutSet, svActorVirtualTimeoutCancel – Set an actor’s virtual timeout; Cancel an actor’s virtual timeout
<b>SYNOPSIS</b>	<pre>#include &lt;vtimer/chVtimer.h&gt; int svActorVirtualTimeoutSet(KnCap * Iactor, KnVirtTimeout * vtimeout, KnTimeVal * cputime, int flag, KnLapDesc * lapdesc);  int svActorVirtualTimeoutCancel(KnVirtTimeout * vtimeout);</pre>
<b>FEATURES</b>	VTIMER
<b>DESCRIPTION</b>	<p>The <i>svActorVirtualTimeoutSet</i> function sets a timeout on the total execution time of all threads created within the actor specified by <i>actorcap</i> .</p> <p>Once the threads have consumed <i>cputime</i> of additional execution time, the lap handler designated by <i>lapdesc</i> is invoked with <i>vtimeout</i> as its argument (see <i>lapInvoke(2K)</i> ).</p> <p>If <i>flag</i> is set to <code>K_VTIME_INTERNAL</code>, only the execution time in the threads’ owning actor is taken into account for the timeout.</p> <p>If <i>flag</i> is set to <code>K_VTIME_TOTAL</code>, all execution time is counted, regardless of whether or not the threads are executing in cross-actor invocations.</p> <p>Note that in case of a cross-actor invocation, thread execution time is charged to the thread’s owning actor only, not to the execution actor.</p> <p>The <i>lapDesc</i> argument is a lap descriptor previously created by <i>svLapCreate</i> .</p> <p>The <i>KnVirtTimeout</i> object is opaque; it must be pre-allocated by the caller but its fields are inaccessible to the caller. Virtual timeouts are always relative.</p> <p>Virtual timeout handlers are not invoked at interrupt level. The execution environment is similar to that of an abort handler. The handler is executed by the next thread of that actor to run, and has no restrictions on use of system calls. A thread can be diverted to execute a handler only when executing in its owning actor; during any cross-actor invocation (trap, ipc, or lap call), virtual timeout handler execution is masked.</p> <p>The <i>svActorVirtualTimeoutCancel</i> function attempts to cancel a virtual timeout request. The <i>vtimeout</i> argument contains the address of the <i>KnVirtTimeout</i> object used in the corresponding call to <i>svActorVirtualTimeoutSet</i> . If the timeout request is still pending, it is cancelled and the call returns <code>K_EOK</code>. If the virtual timeout interval has passed and the handler invocation has been posted, the call takes no action and returns <code>K_ETIMEOUT</code>. In the latter case, there is no information available as to whether the handler execution has actually begun, nor whether the handler is still executing.</p>

A timeout cannot be reinstalled. To reinstall a timeout request, the timeout must be previously cancelled, otherwise an unpredictable behaviour may occur.

**RETURN VALUES**

Upon successful completion, *svActorVirtualTimeoutSet* and *svActorVirtualTimeoutCancel* return K\_OK. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EINVAL]                   The actor capability is not valid  
 [K\_EINVAL]                   The validity of the LAP descriptor is not checked.  
 [K\_EUNKNOWN]                *actor* is unreachable.  
 [K\_ETIMEOUT]                ( *svActorVirtualTimeoutCancel* only) the virtual interval had already expired, and the handler had been queued for execution.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke(2K)* , *svLapCreate(2K)* , *svThreadVirtualTimeoutSet(2K)* , *threadTimes(2K)* , *virtualTimeGetRes(2K)*

<b>NAME</b>	svActorVirtualTimeout, svActorVirtualTimeoutSet, svActorVirtualTimeoutCancel – Set an actor’s virtual timeout; Cancel an actor’s virtual timeout
<b>SYNOPSIS</b>	<pre>#include &lt;vtimer/chVtimer.h&gt; int svActorVirtualTimeoutSet(KnCap * Iactor, KnVirtTimeout * vtimeout, KnTimeVal * cputime, int flag, KnLapDesc * lapdesc);  int svActorVirtualTimeoutCancel(KnVirtTimeout * vtimeout);</pre>
<b>FEATURES</b>	VTIMER
<b>DESCRIPTION</b>	<p>The <i>svActorVirtualTimeoutSet</i> function sets a timeout on the total execution time of all threads created within the actor specified by <i>actorcap</i> .</p> <p>Once the threads have consumed <i>cputime</i> of additional execution time, the lap handler designated by <i>lapdesc</i> is invoked with <i>vtimeout</i> as its argument (see <i>lapInvoke(2K)</i> ).</p> <p>If <i>flag</i> is set to <code>K_VTIME_INTERNAL</code>, only the execution time in the threads’ owning actor is taken into account for the timeout.</p> <p>If <i>flag</i> is set to <code>K_VTIME_TOTAL</code>, all execution time is counted, regardless of whether or not the threads are executing in cross-actor invocations.</p> <p>Note that in case of a cross-actor invocation, thread execution time is charged to the thread’s owning actor only, not to the execution actor.</p> <p>The <i>lapDesc</i> argument is a lap descriptor previously created by <i>svLapCreate</i> .</p> <p>The <i>KnVirtTimeout</i> object is opaque; it must be pre-allocated by the caller but its fields are inaccessible to the caller. Virtual timeouts are always relative.</p> <p>Virtual timeout handlers are not invoked at interrupt level. The execution environment is similar to that of an abort handler. The handler is executed by the next thread of that actor to run, and has no restrictions on use of system calls. A thread can be diverted to execute a handler only when executing in its owning actor; during any cross-actor invocation (trap, ipc, or lap call), virtual timeout handler execution is masked.</p> <p>The <i>svActorVirtualTimeoutCancel</i> function attempts to cancel a virtual timeout request. The <i>vtimeout</i> argument contains the address of the <i>KnVirtTimeout</i> object used in the corresponding call to <i>svActorVirtualTimeoutSet</i> . If the timeout request is still pending, it is cancelled and the call returns <code>K_EOK</code>. If the virtual timeout interval has passed and the handler invocation has been posted, the call takes no action and returns <code>K_ETIMEOUT</code>. In the latter case, there is no information available as to whether the handler execution has actually begun, nor whether the handler is still executing.</p>

A timeout cannot be reinstalled. To reinstall a timeout request, the timeout must be previously cancelled, otherwise an unpredictable behaviour may occur.

**RETURN VALUES**

Upon successful completion, *svActorVirtualTimeoutSet* and *svActorVirtualTimeoutCancel* return K\_OK. Otherwise, a negative error code is returned.

**ERRORS**

[K_EINVAL]	The actor capability is not valid
[K_EINVAL]	The validity of the LAP descriptor is not checked.
[K_EUNKNOWN]	<i>actor</i> is unreachable.
[K_ETIMEOUT]	( <i>svActorVirtualTimeoutCancel</i> only) the virtual interval had already expired, and the handler had been queued for execution.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke(2K)* , *svLapCreate(2K)* , *svThreadVirtualTimeoutSet(2K)* , *threadTimes(2K)* , *virtualTimeGetRes(2K)*

<b>NAME</b>	svActorVirtualTimeout, svActorVirtualTimeoutSet, svActorVirtualTimeoutCancel – Set an actor’s virtual timeout; Cancel an actor’s virtual timeout
<b>SYNOPSIS</b>	<pre>#include &lt;vtimer/chVtimer.h&gt; int svActorVirtualTimeoutSet(KnCap * Iactor, KnVirtTimeout * vtimeout, KnTimeVal * cputime, int flag, KnLapDesc * lapdesc);  int svActorVirtualTimeoutCancel(KnVirtTimeout * vtimeout);</pre>
<b>FEATURES</b>	VTIMER
<b>DESCRIPTION</b>	<p>The <i>svActorVirtualTimeoutSet</i> function sets a timeout on the total execution time of all threads created within the actor specified by <i>actorcap</i> .</p> <p>Once the threads have consumed <i>cputime</i> of additional execution time, the lap handler designated by <i>lapdesc</i> is invoked with <i>vtimeout</i> as its argument (see <i>lapInvoke(2K)</i> ).</p> <p>If <i>flag</i> is set to <i>K_VTIME_INTERNAL</i>, only the execution time in the threads’ owning actor is taken into account for the timeout.</p> <p>If <i>flag</i> is set to <i>K_VTIME_TOTAL</i>, all execution time is counted, regardless of whether or not the threads are executing in cross-actor invocations.</p> <p>Note that in case of a cross-actor invocation, thread execution time is charged to the thread’s owning actor only, not to the execution actor.</p> <p>The <i>lapDesc</i> argument is a lap descriptor previously created by <i>svLapCreate</i> .</p> <p>The <i>KnVirtTimeout</i> object is opaque; it must be pre-allocated by the caller but its fields are inaccessible to the caller. Virtual timeouts are always relative.</p> <p>Virtual timeout handlers are not invoked at interrupt level. The execution environment is similar to that of an abort handler. The handler is executed by the next thread of that actor to run, and has no restrictions on use of system calls. A thread can be diverted to execute a handler only when executing in its owning actor; during any cross-actor invocation (trap, ipc, or lap call), virtual timeout handler execution is masked.</p> <p>The <i>svActorVirtualTimeoutCancel</i> function attempts to cancel a virtual timeout request. The <i>vtimeout</i> argument contains the address of the <i>KnVirtTimeout</i> object used in the corresponding call to <i>svActorVirtualTimeoutSet</i> . If the timeout request is still pending, it is cancelled and the call returns <i>K_EOK</i>. If the virtual timeout interval has passed and the handler invocation has been posted, the call takes no action and returns <i>K_ETIMEOUT</i>. In the latter case, there is no information available as to whether the handler execution has actually begun, nor whether the handler is still executing.</p>

A timeout cannot be reinstalled. To reinstall a timeout request, the timeout must be previously cancelled, otherwise an unpredictable behaviour may occur.

**RETURN VALUES**

Upon successful completion, *svActorVirtualTimeoutSet* and *svActorVirtualTimeoutCancel* return K\_OK. Otherwise, a negative error code is returned.

**ERRORS**

[K_EINVAL]	The actor capability is not valid
[K_EINVAL]	The validity of the LAP descriptor is not checked.
[K_EUNKNOWN]	<i>actor</i> is unreachable.
[K_ETIMEOUT]	( <i>svActorVirtualTimeoutCancel</i> only) the virtual interval had already expired, and the handler had been queued for execution.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke(2K)* , *svLapCreate(2K)* , *svThreadVirtualTimeoutSet(2K)* , *threadTimes(2K)* , *virtualTimeGetRes(2K)*

<b>NAME</b>	svCopyIn, svCopyInString, svCopyOut – Copy from trap caller space; Copy string from trap caller space; Copy to trap caller space				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int svCopyIn(VmAddr src, void * dest, int count);  int svCopyInString(VmAddr src, void * dest, int maxCount);  int svCopyOut(void * src, VmAddr dest, int count);</pre>				
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL				
<b>DESCRIPTION</b>	<p>These calls are only used by trap handlers (see <i>svTrapConnect</i> (2K)). They are used to copy the trap caller parameters.</p> <p>The <i>svCopyIn</i> function copies <i>count</i> bytes from the address <i>src</i> in the caller's address space, to the address <i>dest</i> in the kernel address space. If access to the area defined by <i>src</i> and <i>count</i> in the caller address space produces a memory access fault, <i>svCopyIn</i> will return an error code (K_EFAULT).</p> <p>The <i>svCopyInString</i> function copies characters from the address <i>src</i> in the caller's address space, to the address <i>dest</i> in the kernel address space until a maximum <i>maxCount</i> bytes have been transferred (including the NULL terminating character), or a NULL character is transferred. If access to the area defined by <i>src</i> and <i>maxCount</i> in the caller address space produces a memory access fault, <i>svCopyInString</i> will return an error code (K_EFAULT). If <i>src</i> string (including the NULL character) is bigger than <i>maxCount</i>, <i>svCopyInString</i> will return an error code (K_ENOMEM). If no error is found during string copy, <i>svCopyInString</i> will return the actual number of characters transferred(including the NULL character).</p> <p>The <i>svCopyOut</i> function copies <i>count</i> bytes from the address <i>src</i> in the kernel address space, to the address <i>dest</i> in the caller address space. If access to the area defined by <i>dest</i> and <i>count</i> in the caller address space produces a memory access fault, <i>svCopyOut</i> will return an error code (K_EFAULT).</p>				
<b>RETURN VALUE</b>	If successful, <i>svCopyIn</i> and <i>svCopyOut</i> return 0. Otherwise, they return K_EFAULT. If successful, <i>svCopyInString</i> returns the number of characters transferred, otherwise, it returns K_EFAULT, or K_ENOMEM.				
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EFAULT]</td> <td>Access to the provided caller's data produced a memory access fault.</td> </tr> <tr> <td>[K_ENOMEM]</td> <td>String is longer than requested.</td> </tr> </table>	[K_EFAULT]	Access to the provided caller's data produced a memory access fault.	[K_ENOMEM]	String is longer than requested.
[K_EFAULT]	Access to the provided caller's data produced a memory access fault.				
[K_ENOMEM]	String is longer than requested.				
<b>ATTRIBUTES</b>	See <i>attributes</i> (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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				



<b>NAME</b>	svCopyIn, svCopyInString, svCopyOut – Copy from trap caller space; Copy string from trap caller space; Copy to trap caller space				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int svCopyIn(VmAddr src, void * dest, int count);  int svCopyInString(VmAddr src, void * dest, int maxCount);  int svCopyOut(void * src, VmAddr dest, int count);</pre>				
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL				
<b>DESCRIPTION</b>	<p>These calls are only used by trap handlers (see <i>svTrapConnect</i> (2K)). They are used to copy the trap caller parameters.</p> <p>The <i>svCopyIn</i> function copies <i>count</i> bytes from the address <i>src</i> in the caller's address space, to the address <i>dest</i> in the kernel address space. If access to the area defined by <i>src</i> and <i>count</i> in the caller address space produces a memory access fault, <i>svCopyIn</i> will return an error code (K_EFAULT).</p> <p>The <i>svCopyInString</i> function copies characters from the address <i>src</i> in the caller's address space, to the address <i>dest</i> in the kernel address space until a maximum <i>maxCount</i> bytes have been transferred (including the NULL terminating character), or a NULL character is transferred. If access to the area defined by <i>src</i> and <i>maxCount</i> in the caller address space produces a memory access fault, <i>svCopyInString</i> will return an error code (K_EFAULT). If <i>src</i> string (including the NULL character) is bigger than <i>maxCount</i>, <i>svCopyInString</i> will return an error code (K_ENOMEM). If no error is found during string copy, <i>svCopyInString</i> will return the actual number of characters transferred(including the NULL character).</p> <p>The <i>svCopyOut</i> function copies <i>count</i> bytes from the address <i>src</i> in the kernel address space, to the address <i>dest</i> in the caller address space. If access to the area defined by <i>dest</i> and <i>count</i> in the caller address space produces a memory access fault, <i>svCopyOut</i> will return an error code (K_EFAULT).</p>				
<b>RETURN VALUE</b>	If successful, <i>svCopyIn</i> and <i>svCopyOut</i> return 0. Otherwise, they return K_EFAULT. If successful, <i>svCopyInString</i> returns the number of characters transferred, otherwise, it returns K_EFAULT, or K_ENOMEM.				
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EFAULT]</td> <td>Access to the provided caller's data produced a memory access fault.</td> </tr> <tr> <td>[K_ENOMEM]</td> <td>String is longer than requested.</td> </tr> </table>	[K_EFAULT]	Access to the provided caller's data produced a memory access fault.	[K_ENOMEM]	String is longer than requested.
[K_EFAULT]	Access to the provided caller's data produced a memory access fault.				
[K_ENOMEM]	String is longer than requested.				
<b>ATTRIBUTES</b>	See <i>attributes</i> (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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				



<b>NAME</b>	svCopyIn, svCopyInString, svCopyOut – Copy from trap caller space; Copy string from trap caller space; Copy to trap caller space				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int svCopyIn(VmAddr src, void * dest, int count);  int svCopyInString(VmAddr src, void * dest, int maxCount);  int svCopyOut(void * src, VmAddr dest, int count);</pre>				
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL				
<b>DESCRIPTION</b>	<p>These calls are only used by trap handlers (see <i>svTrapConnect</i> (2K)). They are used to copy the trap caller parameters.</p> <p>The <i>svCopyIn</i> function copies <i>count</i> bytes from the address <i>src</i> in the caller's address space, to the address <i>dest</i> in the kernel address space. If access to the area defined by <i>src</i> and <i>count</i> in the caller address space produces a memory access fault, <i>svCopyIn</i> will return an error code (K_EFAULT).</p> <p>The <i>svCopyInString</i> function copies characters from the address <i>src</i> in the caller's address space, to the address <i>dest</i> in the kernel address space until a maximum <i>maxCount</i> bytes have been transferred (including the NULL terminating character), or a NULL character is transferred. If access to the area defined by <i>src</i> and <i>maxCount</i> in the caller address space produces a memory access fault, <i>svCopyInString</i> will return an error code (K_EFAULT). If <i>src</i> string (including the NULL character) is bigger than <i>maxCount</i>, <i>svCopyInString</i> will return an error code (K_ENOMEM). If no error is found during string copy, <i>svCopyInString</i> will return the actual number of characters transferred(including the NULL character).</p> <p>The <i>svCopyOut</i> function copies <i>count</i> bytes from the address <i>src</i> in the kernel address space, to the address <i>dest</i> in the caller address space. If access to the area defined by <i>dest</i> and <i>count</i> in the caller address space produces a memory access fault, <i>svCopyOut</i> will return an error code (K_EFAULT).</p>				
<b>RETURN VALUE</b>	If successful, <i>svCopyIn</i> and <i>svCopyOut</i> return 0. Otherwise, they return K_EFAULT. If successful, <i>svCopyInString</i> returns the number of characters transferred, otherwise, it returns K_EFAULT, or K_ENOMEM.				
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EFAULT]</td> <td>Access to the provided caller's data produced a memory access fault.</td> </tr> <tr> <td>[K_ENOMEM]</td> <td>String is longer than requested.</td> </tr> </table>	[K_EFAULT]	Access to the provided caller's data produced a memory access fault.	[K_ENOMEM]	String is longer than requested.
[K_EFAULT]	Access to the provided caller's data produced a memory access fault.				
[K_ENOMEM]	String is longer than requested.				
<b>ATTRIBUTES</b>	See <i>attributes</i> (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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				



<b>NAME</b>	svExceptionHandler, svAbortHandler – Define an exception handler; Define an abort handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svExceptionHandler(KnCap * actorcap, KnExcHdl routine);  int svAbortHandler(KnCap * actorcap, KnAbortHdl routine);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svExceptionHandler</i> function defines an exception handler for the actor whose capability is pointed to by <i>actorcap</i>. If <i>actorcap</i> is <code>K_MYACTOR</code>, the current actor is used. The <i>routine</i> field defines a handler that will be called every time a thread whose execution actor matches the specified actor encounters an exception (see <i>svTrapConnect</i> (2K) for details on execution actors). An exception handler is a function which takes three arguments:</p> <pre>int handler (ctx, excno, ptr)      KnThreadCtx*   ctx ;     int            excno ;     void*          ptr ;</pre> <p>The <i>ctx</i> pointer is the register context of the faulty thread. It points to a <i>KnThreadCtx</i> structure. The fields of this structure are machine-dependent.</p> <p>The <i>excno</i> field is a machine-dependent exception number. It can only be used for exception handlers.</p> <p>The <i>ptr</i> pointer highlight any relevant machine-dependent data that may be needed to process the exception (for example,. the faulty address for a page fault). It may not be used on certain hardware architectures.</p> <p>The exception handler may modify the faulty thread's register context before returning.</p> <p>The <i>svAbortHandler</i> function defines a handler to be called each time a thread of the specified actor returns from executing within its home actor environment, while in the <i>ABORTED</i> state. When the handler is called, the thread exits its <i>ABORTED</i> state (just as if <i>threadAborted</i> (2K) were called). An abort handler has the same form as an exception handler, except that the exception number is the <code>K_EABORT</code> constant, and no return value is taken into account.</p> <p>These calls are restricted to <i>SUPERVISOR</i> threads. Handlers are executed in <i>SUPERVISOR</i> execution mode. The handlers' code and data accessed must be parts of the locked-in-memory regions of a <i>SUPERVISOR</i> actor.</p> <p>These calls are used by subsystem process managers. The <i>svExceptionHandler</i> function is used to catch exceptions, and to translate them into subsystem-specific events</p>

(for example, . UNIX signals); *svAbortHandler* is used to force abort requests to be taken into account by controlled threads.

**RETURN VALUE**

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EINVAL] *actorcap* is an inconsistent actor capability.

[K\_EUNKNOWN] *actorcap* does not specify a local actor.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*svTrapConnect(2K)*

<b>NAME</b>	svGetInvoker – get handler invoker				
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int svGetInvoker(KnCap *actorcap);</pre>				
<b>FEATURES</b>	CORE				
<b>DESCRIPTION</b>	<p>The <i>svGetInvoker</i> function returns the capability of the execution actor prior to the execution of the current handler to the value pointed to by <i>actorcap</i>.</p> <p>When called outside the execution of a handler (connected by <i>svTrapConnect</i>(2K), <i>svExcHandler</i>(2K), <i>svAbortHandler</i>(2K), <i>svMsgHandler</i>(2K), <i>svLapCreate</i>(2K)), or during execution of an IPC message handler for a remote <i>ipcCall</i>(2K), <i>svGetInvoker</i> will return K_EINVAL.</p> <p>The <i>svGetInvoker</i> function should not be called from within an interrupt handler.</p> <p>This call is restricted to <i>SUPERVISOR</i> threads.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	[K_EINVAL]                      The current thread is not executing a handler.				
<b>ATTRIBUTES</b>	See <i>attributes</i> (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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>svTrapConnect</i> (2K), <i>svExcHandler</i> (2K), <i>svAbortHandler</i> (2K), <i>svMsgHandler</i> (2K), <i>svLapCreate</i> (2K), <i>ipcCall</i> (2K), <i>lapInvoke</i> (2K)				

<b>NAME</b>	svLapBind, svLapUnbind, lapResolve – bind a symbolic name to a lap descriptor; unbind the symbolic name bound to a lap descriptor; get a lap descriptor from a lap symbolic name										
<b>SYNOPSIS</b>	<pre>#include &lt;lap/chLap.h&gt; int svLapBind(KnLapDesc * lapdesc, char * name, unsigned int options);  int svLapUnbind(char * name);  int lapResolve(KnLapDesc * lapdesc, char * name, unsigned int options);</pre>										
<b>FEATURES</b>	LAPBIND										
<b>DESCRIPTION</b>	<p>The <i>svLapBind</i> (2K) system call binds the lap descriptor pointed to by <i>lapdesc</i> with the symbolic name pointed to by <i>name</i> .</p> <p><i>name</i> points to a null-terminated string of K_LAPNAMEMAX characters at most (not including the null-terminating character).</p> <p>If the K_LAP_PROTECTED option is set in the <i>options</i> parameter, the binding will only be visible to trusted threads (e.g. threads executing in user actors will get an error from <i>lapResolve</i> (2K) on this name).</p> <p>The <i>svLapUnbind</i> (2K) system call removes the lap binding associated to <i>name</i> .</p> <p>The <i>lapResolve</i> (2K) system call initializes the lap descriptor pointed to by <i>lapdesc</i> with the lap descriptor bound to <i>name</i> . <i>lapResolve</i> (2K) will block until a lap descriptor is bound to <i>name</i> , except if the K_LAP_NOBLOCK option is set in <i>options</i> , in which case the call returns immediatly with an error.</p> <p><i>svLapBind</i> and <i>svLapUnbind</i> are restricted to SUPERVISOR threads.</p>										
<b>RETURN VALUE</b>	On success, these calls return K_OK. Otherwise, a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EABORT]</td> <td>The calling thread has been aborted in <i>lapResolve</i> (2K).</td> </tr> <tr> <td style="vertical-align: top;">[K_EBUSY]</td> <td>The name given to <i>svLapBind</i> (2K) is already in use.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>The <i>lapdesc</i> or <i>name</i> arguments of <i>lapResolve</i> points to the outside of the caller's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td>The <i>name</i> given to <i>svLapUnbind</i> (2K) is not bound to any lap, or the <i>name</i> given to <i>lapResolve</i> (2K) is not bound to any lap, and the K_LAP_NOBLOCK option is set.</td> </tr> <tr> <td style="vertical-align: top;">[K_ENOMEM]</td> <td>The <i>name</i> given to <i>svLapBind</i> (2K), <i>svLapUnbind</i> (2K) or <i>lapResolve</i> (2K) is too long. [K_ENOMEM] The system is out of resources.</td> </tr> </table>	[K_EABORT]	The calling thread has been aborted in <i>lapResolve</i> (2K).	[K_EBUSY]	The name given to <i>svLapBind</i> (2K) is already in use.	[K_EFAULT]	The <i>lapdesc</i> or <i>name</i> arguments of <i>lapResolve</i> points to the outside of the caller's address space.	[K_EINVAL]	The <i>name</i> given to <i>svLapUnbind</i> (2K) is not bound to any lap, or the <i>name</i> given to <i>lapResolve</i> (2K) is not bound to any lap, and the K_LAP_NOBLOCK option is set.	[K_ENOMEM]	The <i>name</i> given to <i>svLapBind</i> (2K), <i>svLapUnbind</i> (2K) or <i>lapResolve</i> (2K) is too long. [K_ENOMEM] The system is out of resources.
[K_EABORT]	The calling thread has been aborted in <i>lapResolve</i> (2K).										
[K_EBUSY]	The name given to <i>svLapBind</i> (2K) is already in use.										
[K_EFAULT]	The <i>lapdesc</i> or <i>name</i> arguments of <i>lapResolve</i> points to the outside of the caller's address space.										
[K_EINVAL]	The <i>name</i> given to <i>svLapUnbind</i> (2K) is not bound to any lap, or the <i>name</i> given to <i>lapResolve</i> (2K) is not bound to any lap, and the K_LAP_NOBLOCK option is set.										
[K_ENOMEM]	The <i>name</i> given to <i>svLapBind</i> (2K), <i>svLapUnbind</i> (2K) or <i>lapResolve</i> (2K) is too long. [K_ENOMEM] The system is out of resources.										

[K\_EPRIV]

*lapResolve* (2K) is called by a non trusted thread to resolve a name which has been bound with the K\_LAP\_PROTECTED option.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svLapCreate(2K)`

<b>NAME</b>	svLapCreate, lapDescZero, lapDescIsZero, lapDescDup – create a lap; reset a lap descriptor; test if a lap descriptor has been initialized; duplicate a lap descriptor
<b>SYNOPSIS</b>	<pre>#include &lt;lap/chLap.h&gt; int svLapCreate(KnCap * actorcap, KnLapHdl laphdl, void * cookie, unsigned int options, KnLapDesc * lapdesc);  void lapDescZero(KnLapDesc * lapdesc);  int lapDescIsZero(KnLapDesc * lapdesc);  void lapDescDup(KnLapDesc * olddesc, KnLapDesc * newdesc);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svLapCreate</i> (2K) system call creates a new local access point (lap) for the actor designated by the <i>actorcap</i> capability. If <i>actorcap</i> is equal to <i>K_MYACTOR</i>, the home actor of the calling thread is considered. If the call succeeds, the lap descriptor pointed to by <i>lapdesc</i> represents this new lap and can be passed to client threads that will use it to invoke the lap handler.</p> <p>On lap invocation (see <i>lapInvoke</i> (2K)), the <i>laphdl</i> handler is called. This handler is a function which takes two arguments:</p> <pre>void handler (arg, cookie) void          *arg ; void          *cookie ;</pre> <p>Where <i>arg</i> is the argument specified by <i>lapInvoke</i> (2K), and <i>cookie</i> is the value of the cookie specified by <i>svLapCreate</i> (2K).</p> <p>By default as a consequence of a lap invocation, the execution actor of the calling thread are set to the lap owning actor for the duration of the lap handler execution.</p> <p>Additional actions may take place during the lap invocation, depending on the combination of flags specified in the <i>options</i> parameter:</p> <p><b>K_LAP_SAFE</b>                    A "lap frame" descriptor is allocated to register the calling thread as a temporary resource of the invoked actor. Each "lap frame" has an associated level (lap frame level) which represents the number of lap frames for the considered thread.</p> <p>This option is only valid if the LAPS SAFE feature has been specified into the system. It will enforce a stronger checking during lap invocation. This guarantees in particular that the kernel will synchronize the <i>svLapDelete</i> (2K) operation with concurrent lap invocations.</p>

Furthermore, full context of the calling thread is saved prior to the lap invocation. This guarantees that the calling thread can return from its invocation even if a failure (exception, deletion of the lap owning actor, etc.) occurs during the execution of the lap handler.

This option is mandatory for laps to be called from user mode.

The two utility routines *lapDescZero* (2K) and *lapDescIsZero* (2K) are available to manipulate the state of a lap descriptor with the following semantics: if a lap descriptor has been initialized with *lapDescZero* (2K) or is implemented in a zero-filled memory region, *lapDescIsZero* (2K) returns a non-zero value until the lap descriptor has been successfully initialized by *svLapCreate* (2K) or *lapResolve* (2K).

*lapDescDup* (2K) must be used to duplicate the contents of a lap descriptor.

**RETURN VALUE**

On success, *svLapCreate* (2K) returns K\_OK. Otherwise, a negative error code is returned.

*lapDescIsZero* (2K) returns a non-zero value if the lap descriptor is not initialized.

*lapDescZero* (2K) and *lapDescDup* (2K) have no return values.

**ERRORS**

[K\_EINVAL] *actorcap* is an inconsistent actor capability.

[K\_ENOMEM] The system is out of resources.

**RESTRICTIONS**

The current implementation does not support the K\_LAP\_SETJMP option, which is silently ignored.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke*(2K) , *lapResolve*(2K) , *svLapBind*(2K) , *svLapDelete*(2K) , *svLapUnbind*(2K) , *threadStat*(2K)

<b>NAME</b>	svLapDelete – delete a lap				
<b>SYNOPSIS</b>	<pre>#include &lt;lap/chLap.h&gt; int svLapDelete(KnLapDesc *lapdesc);</pre>				
<b>DESCRIPTION</b>	<p>The <i>svLapDelete(2K)</i> system call deletes the local access point the descriptor of which is pointed to by <i>lapdesc</i>.</p> <p>If the corresponding lap has been created with the <code>K_LAP_SAFE</code> option set (see <i>svLapCreate(2K)</i>), <i>svLapDelete(2K)</i> will return only when all invocations currently in progress on this lap have completed. When a safe lap is deleted, the threads currently executing in the lap are aborted. These threads return with error from the invocation of the lap as described in (see <i>lapInvoke(2K)</i>).</p> <p>Symbolic name bindings associated with this lap (see <i>svLapBind(2K)</i>) are not affected by <i>svLapDelete(2K)</i> and must be explicitly deleted (see <i>svLapUnbind(2K)</i>).</p>				
<b>RETURN VALUE</b>	On success, <i>svLapDelete(2K)</i> returns <code>K_OK</code> . Otherwise, a negative error code is returned.				
<b>ERRORS</b>	[ <code>K_EINVAL</code> ] <i>lapdesc</i> points to an invalid lap descriptor.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>lapInvoke(2K)</i> , <i>svLapBind(2K)</i> , <i>svLapCreate(2K)</i> , <i>svLapUnbind(2K)</i>				

<b>NAME</b>	svLapBind, svLapUnbind, lapResolve – bind a symbolic name to a lap descriptor; unbind the symbolic name bound to a lap descriptor; get a lap descriptor from a lap symbolic name										
<b>SYNOPSIS</b>	<pre>#include &lt;lap/chLap.h&gt; int svLapBind(KnLapDesc * lapdesc, char * name, unsigned int options);  int svLapUnbind(char * name);  int lapResolve(KnLapDesc * lapdesc, char * name, unsigned int options);</pre>										
<b>FEATURES</b>	LAPBIND										
<b>DESCRIPTION</b>	<p>The <i>svLapBind</i> (2K) system call binds the lap descriptor pointed to by <i>lapdesc</i> with the symbolic name pointed to by <i>name</i> .</p> <p><i>name</i> points to a null-terminated string of K_LAPNAMEMAX characters at most (not including the null-terminating character).</p> <p>If the K_LAP_PROTECTED option is set in the <i>options</i> parameter, the binding will only be visible to trusted threads (e.g. threads executing in user actors will get an error from <i>lapResolve</i> (2K) on this name).</p> <p>The <i>svLapUnbind</i> (2K) system call removes the lap binding associated to <i>name</i> .</p> <p>The <i>lapResolve</i> (2K) system call initializes the lap descriptor pointed to by <i>lapdesc</i> with the lap descriptor bound to <i>name</i> . <i>lapResolve</i> (2K) will block until a lap descriptor is bound to <i>name</i> , except if the K_LAP_NOBLOCK option is set in <i>options</i> , in which case the call returns immediatly with an error.</p> <p><i>svLapBind</i> and <i>svLapUnbind</i> are restricted to <i>SUPERVISOR</i> threads.</p>										
<b>RETURN VALUE</b>	On success, these calls return K_OK. Otherwise, a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EABORT]</td> <td>The calling thread has been aborted in <i>lapResolve</i> (2K).</td> </tr> <tr> <td style="vertical-align: top;">[K_EBUSY]</td> <td>The name given to <i>svLapBind</i> (2K) is already in use.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>The <i>lapdesc</i> or <i>name</i> arguments of <i>lapResolve</i> points to the outside of the caller's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td>The <i>name</i> given to <i>svLapUnbind</i> (2K) is not bound to any lap, or the <i>name</i> given to <i>lapResolve</i> (2K) is not bound to any lap, and the K_LAP_NOBLOCK option is set.</td> </tr> <tr> <td style="vertical-align: top;">[K_ENOMEM]</td> <td>The <i>name</i> given to <i>svLapBind</i> (2K), <i>svLapUnbind</i> (2K) or <i>lapResolve</i> (2K) is too long. [K_ENOMEM] The system is out of resources.</td> </tr> </table>	[K_EABORT]	The calling thread has been aborted in <i>lapResolve</i> (2K).	[K_EBUSY]	The name given to <i>svLapBind</i> (2K) is already in use.	[K_EFAULT]	The <i>lapdesc</i> or <i>name</i> arguments of <i>lapResolve</i> points to the outside of the caller's address space.	[K_EINVAL]	The <i>name</i> given to <i>svLapUnbind</i> (2K) is not bound to any lap, or the <i>name</i> given to <i>lapResolve</i> (2K) is not bound to any lap, and the K_LAP_NOBLOCK option is set.	[K_ENOMEM]	The <i>name</i> given to <i>svLapBind</i> (2K), <i>svLapUnbind</i> (2K) or <i>lapResolve</i> (2K) is too long. [K_ENOMEM] The system is out of resources.
[K_EABORT]	The calling thread has been aborted in <i>lapResolve</i> (2K).										
[K_EBUSY]	The name given to <i>svLapBind</i> (2K) is already in use.										
[K_EFAULT]	The <i>lapdesc</i> or <i>name</i> arguments of <i>lapResolve</i> points to the outside of the caller's address space.										
[K_EINVAL]	The <i>name</i> given to <i>svLapUnbind</i> (2K) is not bound to any lap, or the <i>name</i> given to <i>lapResolve</i> (2K) is not bound to any lap, and the K_LAP_NOBLOCK option is set.										
[K_ENOMEM]	The <i>name</i> given to <i>svLapBind</i> (2K), <i>svLapUnbind</i> (2K) or <i>lapResolve</i> (2K) is too long. [K_ENOMEM] The system is out of resources.										

[K\_EPRIV]

*lapResolve* (2K) is called by a non trusted thread to resolve a name which has been bound with the K\_LAP\_PROTECTED option.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svLapCreate(2K)`

<b>NAME</b>	svMaskAll, svUnmaskAll, svUnmask – Disable interrupt processing; Enable interrupt processing; Reenable interrupt processing				
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSpinLock.h&gt; int svMaskAll(void);  void svUnmaskAll(void);  void svUnmask(int oldval);</pre>				
<b>DESCRIPTION</b>	<p>The <i>svMaskAll</i> function disables interrupts on the current processor by setting the processor interrupt mask to its maximum value. There is no effect on other processors. It returns the previous interrupt mask.</p> <p>The <i>svUnmaskAll</i> function enables interrupts on the current processor by setting the processor interrupt mask to its minimum value. There is no effect on other processors.</p> <p>The <i>svUnmask</i> function resets the interrupt mask to a state represented by <i>oldval</i>. The <i>oldval</i> value must have been returned by a previous invocation of <i>svMaskAll</i>. Pairs of <i>svMaskAll</i> and <i>svUnmask</i> may be nested. The interrupt mask is part of a thread's execution context, it is saved as part of the thread's status when the thread is suspended. If a thread performs a blocking call after calling <i>svMaskAll</i>, interrupt processing may be enabled by another scheduled thread. It will be disabled again when the thread is restarted. Such situations should be avoided.</p> <p>The <i>svMaskAll</i>, <i>svUnmaskAll</i> and <i>svUnmask</i> functions only relate to the interrupt mask from processor standpoint. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.</p> <p>These calls are restricted to <i>SUPERVISOR</i> threads. Their use must be limited to a very small number of instructions. They should not be used to ensure mutual exclusion. The routines <i>svMaskedLockGet</i> and <i>svMaskedLockRel</i> should be used for this purpose.</p>				
<b>RETURN VALUE</b>	<i>svMaskAll</i> returns the current interrupt masking status.				
<b>ERRORS</b>	No error messages are returned .				
<b>ATTRIBUTES</b>	See <a href="#">attributes(5)</a> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>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				
<b>SEE ALSO</b>	<a href="#">svMaskedLockInit(2K)</a>				

<b>NAME</b>	svMaskedLockInit, svMaskedLockGet, svMaskedLockTry, svMaskedLockRel, svSpinLockInit, svSpinLockGet, svSpinLockTry, svSpinLockRel – Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSpinLock.h&gt; void <b>svMaskedLockInit</b>(KnMaskedSpinLock * <i>lockaddr</i>, KnMaskedSpinLockInfo * <i>info</i>);  void <b>svMaskedLockGet</b>(KnMaskedSpinLock * <i>lockaddr</i>);  int <b>svMaskedLockTry</b>(KnMaskedSpinLock * <i>lockaddr</i>);  void <b>svMaskedLockRel</b>(KnMaskedSpinLock * <i>lockaddr</i>);  void <b>svSpinLockInit</b>(KnSpinLock * <i>lockaddr</i>, KnSpinLockInfo * <i>info</i>);  void <b>svSpinLockGet</b>(KnSpinLock * <i>lockaddr</i>);  int <b>svSpinLockTry</b>(KnSpinLock * <i>lockaddr</i>);  void <b>svSpinLockRel</b>(KnSpinLock * <i>lockaddr</i>);  int <b>svPreemptable</b>(void);  int <b>svIntrLevel</b>(void);</pre>
<b>DESCRIPTION</b>	<p>Spin locks are mutual exclusion objects, used to protect very short critical sections. These critical sections must not include kernel calls and the data they manipulate must be locked in memory to avoid page faults.</p> <p>In multiprocessor systems, threads loop actively when resources are not available. A spin lock is an object in memory, typically modified and checked through indivisible machine operations (such as test-and-set).</p> <p>In monoprocessor systems, spin locks are implemented by disabling preemption or interrupts.</p> <p>Note that while a thread is holding a spin lock, a masked spin lock, or has disabled preemption, it is no longer allowed to block. In this case, the list of system calls available to the thread is restricted to the following:</p> <pre>semV threadSemPost eventPost msgPut svSpinLockGet svSpinLockRel svMaskedLockGet svMaskedLockRel svMaskAll svUnmaskAll svUnmask</pre>

```
svPreemptable  
svIntrLevel  
svTimeoutSet  
svTimeoutCancel  
svPagesAllocate
```

The system not force the thread to restrict itself to the permitted calls. If the thread issues non-permitted calls, the system will hang.

The *svMaskedLockInit* and *svSpinLockInit* functions initialize a spin lock in the free state. The *info* argument allows the user to enforce lock hierarchy and/or to gather lock use/contention statistics in the future.

In multiprocessor systems, *svMaskedLockGet* atomically disables interrupts on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with interrupts enabled to reduce interrupt latency. This function can be called from within an interrupt handler.

The *svMaskedLockTry* function attempts to acquire a spin lock, but returns immediately without spinning if there is none available. This function can be called from an interrupt handler.

The *svMaskedLockRel* function releases a spin lock and re-enables interrupts on the current processor. This function can be called from within an interrupt handler.

The *svMaskedLockGet*, *svMaskedLockTry* and *svMaskedLockRel* functions only relate to the interrupt mask of the current processor, and do not affect interrupt processing on other processors. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.

In multiprocessor systems, *svSpinLockGet* atomically disables preemption on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with preemption enabled to reduce scheduling latency. On a monoprocessor, *svSpinLockGet* disables preemption on the current processor. This function cannot be called from within an interrupt handler.

The *svSpinLockTry* try attempts to acquire a spin lock, returning immediately without spinning if there is none available. This function cannot be called from within an interrupt handler.

The *svSpinLockRel* function releases a spin lock. This function ca not be called from within an interrupt handler.

If the current thread is preemptable, *svPreemptable* returns a non-zero value, otherwise it returns zero.

If called from an interrupt level, *svIntrLevel* returns a non-zero value, otherwise it returns zero.

**RETURN VALUE**

If the lock is successfully acquired, *svSpinLockTry* returns the previous state of the lock (K\_AVAIL). The *svMaskedLockTry* function returns K\_AVAIL. The *svPreemptable* and *svIntrLevel* calls return a value which depends on the current execution context. Other calls return 0.

**ERRORS**

No error codes are returned. In DEBUG mode, the kernel debugger is invoked if the lock level hierarchy is violated, or if a spin lock cannot be acquired after a very large number of iterations, indicating a probable deadlock.

**RESTRICTIONS**

Hierarchy level checks are not performed in this version.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**NAME** svMaskedLockInit, svMaskedLockGet, svMaskedLockTry, svMaskedLockRel, svSpinLockInit, svSpinLockGet, svSpinLockTry, svSpinLockRel – Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock

**SYNOPSIS**

```
#include <sync/chSpinLock.h>
void svMaskedLockInit(KnMaskedSpinLock * lockaddr, KnMaskedSpinLockInfo * info);

void svMaskedLockGet(KnMaskedSpinLock * lockaddr);

int svMaskedLockTry(KnMaskedSpinLock * lockaddr);

void svMaskedLockRel(KnMaskedSpinLock * lockaddr);

void svSpinLockInit(KnSpinLock * lockaddr, KnSpinLockInfo * info);

void svSpinLockGet(KnSpinLock * lockaddr);

int svSpinLockTry(KnSpinLock * lockaddr);

void svSpinLockRel(KnSpinLock * lockaddr);

int svPreemptable(void);

int svIntrLevel(void);
```

**DESCRIPTION**

Spin locks are mutual exclusion objects, used to protect very short critical sections. These critical sections must not include kernel calls and the data they manipulate must be locked in memory to avoid page faults.

In multiprocessor systems, threads loop actively when resources are not available. A spin lock is an object in memory, typically modified and checked through indivisible machine operations (such as test-and-set).

In monoprocessor systems, spin locks are implemented by disabling preemption or interrupts.

Note that while a thread is holding a spin lock, a masked spin lock, or has disabled preemption, it is no longer allowed to block. In this case, the list of system calls available to the thread is restricted to the following:

```
semV
threadSemPost
eventPost
msgPut
svSpinLockGet
svSpinLockRel
svMaskedLockGet
svMaskedLockRel
svMaskAll
svUnmaskAll
svUnmask
```

```
svPreemptable  
svIntrLevel  
svTimeoutSet  
svTimeoutCancel  
svPagesAllocate
```

The system not force the thread to restrict itself to the permitted calls. If the thread issues non-permitted calls, the system will hang.

The *svMaskedLockInit* and *svSpinLockInit* functions initialize a spin lock in the free state. The *info* argument allows the user to enforce lock hierarchy and/or to gather lock use/contention statistics in the future.

In multiprocessor systems, *svMaskedLockGet* atomically disables interrupts on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with interrupts enabled to reduce interrupt latency. This function can be called from within an interrupt handler.

The *svMaskedLockTry* function attempts to acquire a spin lock, but returns immediately without spinning if there is none available. This function can be called from an interrupt handler.

The *svMaskedLockRel* function releases a spin lock and re-enables interrupts on the current processor. This function can be called from within an interrupt handler.

The *svMaskedLockGet*, *svMaskedLockTry* and *svMaskedLockRel* functions only relate to the interrupt mask of the current processor, and do not affect interrupt processing on other processors. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.

In multiprocessor systems, *svSpinLockGet* atomically disables preemption on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with preemption enabled to reduce scheduling latency. On a monoprocessor, *svSpinLockGet* disables preemption on the current processor. This function cannot be called from within an interrupt handler.

The *svSpinLockTry* try attempts to acquire a spin lock, returning immediately without spinning if there is none available. This function cannot be called from within an interrupt handler.

The *svSpinLockRel* function releases a spin lock. This function ca not be called from within an interrupt handler.

If the current thread is preemptable, *svPreemptable* returns a non-zero value, otherwise it returns zero.

If called from an interrupt level, *svIntrLevel* returns a non-zero value, otherwise it returns zero.

**RETURN VALUE**

If the lock is successfully acquired, *svSpinLockTry* returns the previous state of the lock (K\_AVAIL). The *svMaskedLockTry* function returns K\_AVAIL. The *svPreemptable* and *svIntrLevel* calls return a value which depends on the current execution context. Other calls return 0.

**ERRORS**

No error codes are returned. In DEBUG mode, the kernel debugger is invoked if the lock level hierarchy is violated, or if a spin lock cannot be acquired after a very large number of iterations, indicating a probable deadlock.

**RESTRICTIONS**

Hierarchy level checks are not performed in this version.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	svMaskedLockInit, svMaskedLockGet, svMaskedLockTry, svMaskedLockRel, svSpinLockInit, svSpinLockGet, svSpinLockTry, svSpinLockRel – Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSpinLock.h&gt; void <b>svMaskedLockInit</b>(KnMaskedSpinLock * <i>lockaddr</i>, KnMaskedSpinLockInfo * <i>info</i>);  void <b>svMaskedLockGet</b>(KnMaskedSpinLock * <i>lockaddr</i>);  int <b>svMaskedLockTry</b>(KnMaskedSpinLock * <i>lockaddr</i>);  void <b>svMaskedLockRel</b>(KnMaskedSpinLock * <i>lockaddr</i>);  void <b>svSpinLockInit</b>(KnSpinLock * <i>lockaddr</i>, KnSpinLockInfo * <i>info</i>);  void <b>svSpinLockGet</b>(KnSpinLock * <i>lockaddr</i>);  int <b>svSpinLockTry</b>(KnSpinLock * <i>lockaddr</i>);  void <b>svSpinLockRel</b>(KnSpinLock * <i>lockaddr</i>);  int <b>svPreemptable</b>(void);  int <b>svIntrLevel</b>(void);</pre>
<b>DESCRIPTION</b>	<p>Spin locks are mutual exclusion objects, used to protect very short critical sections. These critical sections must not include kernel calls and the data they manipulate must be locked in memory to avoid page faults.</p> <p>In multiprocessor systems, threads loop actively when resources are not available. A spin lock is an object in memory, typically modified and checked through indivisible machine operations (such as test-and-set).</p> <p>In monoprocessor systems, spin locks are implemented by disabling preemption or interrupts.</p> <p>Note that while a thread is holding a spin lock, a masked spin lock, or has disabled preemption, it is no longer allowed to block. In this case, the list of system calls available to the thread is restricted to the following:</p> <pre>semV threadSemPost eventPost msgPut svSpinLockGet svSpinLockRel svMaskedLockGet svMaskedLockRel svMaskAll svUnmaskAll svUnmask</pre>

```
svPreemptable  
svIntrLevel  
svTimeoutSet  
svTimeoutCancel  
svPagesAllocate
```

The system not force the thread to restrict itself to the permitted calls. If the thread issues non-permitted calls, the system will hang.

The *svMaskedLockInit* and *svSpinLockInit* functions initialize a spin lock in the free state. The *info* argument allows the user to enforce lock hierarchy and/or to gather lock use/contention statistics in the future.

In multiprocessor systems, *svMaskedLockGet* atomically disables interrupts on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with interrupts enabled to reduce interrupt latency. This function can be called from within an interrupt handler.

The *svMaskedLockTry* function attempts to acquire a spin lock, but returns immediately without spinning if there is none available. This function can be called from an interrupt handler.

The *svMaskedLockRel* function releases a spin lock and re-enables interrupts on the current processor. This function can be called from within an interrupt handler.

The *svMaskedLockGet*, *svMaskedLockTry* and *svMaskedLockRel* functions only relate to the interrupt mask of the current processor, and do not affect interrupt processing on other processors. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.

In multiprocessor systems, *svSpinLockGet* atomically disables preemption on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with preemption enabled to reduce scheduling latency. On a monoprocessor, *svSpinLockGet* disables preemption on the current processor. This function cannot be called from within an interrupt handler.

The *svSpinLockTry* try attempts to acquire a spin lock, returning immediately without spinning if there is none available. This function cannot be called from within an interrupt handler.

The *svSpinLockRel* function releases a spin lock. This function ca not be called from within an interrupt handler.

If the current thread is preemptable, *svPreemptable* returns a non-zero value, otherwise it returns zero.

	If called from an interrupt level, <i>svIntrLevel</i> returns a non-zero value, otherwise it returns zero.
<b>RETURN VALUE</b>	If the lock is successfully acquired, <i>svSpinLockTry</i> returns the previous state of the lock (K_AVAIL). The <i>svMaskedLockTry</i> function returns K_AVAIL. The <i>svPreemptable</i> and <i>svIntrLevel</i> calls return a value which depends on the current execution context. Other calls return 0.
<b>ERRORS</b>	No error codes are returned. In DEBUG mode, the kernel debugger is invoked if the lock level hierarchy is violated, or if a spin lock cannot be acquired after a very large number of iterations, indicating a probable deadlock.
<b>RESTRICTIONS</b>	Hierarchy level checks are not performed in this version.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**NAME** svMaskedLockInit, svMaskedLockGet, svMaskedLockTry, svMaskedLockRel, svSpinLockInit, svSpinLockGet, svSpinLockTry, svSpinLockRel – Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock

**SYNOPSIS**

```
#include <sync/chSpinLock.h>
void svMaskedLockInit(KnMaskedSpinLock * lockaddr, KnMaskedSpinLockInfo * info);

void svMaskedLockGet(KnMaskedSpinLock * lockaddr);

int svMaskedLockTry(KnMaskedSpinLock * lockaddr);

void svMaskedLockRel(KnMaskedSpinLock * lockaddr);

void svSpinLockInit(KnSpinLock * lockaddr, KnSpinLockInfo * info);

void svSpinLockGet(KnSpinLock * lockaddr);

int svSpinLockTry(KnSpinLock * lockaddr);

void svSpinLockRel(KnSpinLock * lockaddr);

int svPreemptable(void);

int svIntrLevel(void);
```

**DESCRIPTION**

Spin locks are mutual exclusion objects, used to protect very short critical sections. These critical sections must not include kernel calls and the data they manipulate must be locked in memory to avoid page faults.

In multiprocessor systems, threads loop actively when resources are not available. A spin lock is an object in memory, typically modified and checked through indivisible machine operations (such as test-and-set).

In monoprocessor systems, spin locks are implemented by disabling preemption or interrupts.

Note that while a thread is holding a spin lock, a masked spin lock, or has disabled preemption, it is no longer allowed to block. In this case, the list of system calls available to the thread is restricted to the following:

```
semV
threadSemPost
eventPost
msgPut
svSpinLockGet
svSpinLockRel
svMaskedLockGet
svMaskedLockRel
svMaskAll
svUnmaskAll
svUnmask
```

```
svPreemptable  
svIntrLevel  
svTimeoutSet  
svTimeoutCancel  
svPagesAllocate
```

The system not force the thread to restrict itself to the permitted calls. If the thread issues non-permitted calls, the system will hang.

The *svMaskedLockInit* and *svSpinLockInit* functions initialize a spin lock in the free state. The *info* argument allows the user to enforce lock hierarchy and/or to gather lock use/contention statistics in the future.

In multiprocessor systems, *svMaskedLockGet* atomically disables interrupts on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with interrupts enabled to reduce interrupt latency. This function can be called from within an interrupt handler.

The *svMaskedLockTry* function attempts to acquire a spin lock, but returns immediately without spinning if there is none available. This function can be called from an interrupt handler.

The *svMaskedLockRel* function releases a spin lock and re-enables interrupts on the current processor. This function can be called from within an interrupt handler.

The *svMaskedLockGet*, *svMaskedLockTry* and *svMaskedLockRel* functions only relate to the interrupt mask of the current processor, and do not affect interrupt processing on other processors. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.

In multiprocessor systems, *svSpinLockGet* atomically disables preemption on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with preemption enabled to reduce scheduling latency. On a monoprocessor, *svSpinLockGet* disables preemption on the current processor. This function cannot be called from within an interrupt handler.

The *svSpinLockTry* try attempts to acquire a spin lock, returning immediately without spinning if there is none available. This function cannot be called from within an interrupt handler.

The *svSpinLockRel* function releases a spin lock. This function ca not be called from within an interrupt handler.

If the current thread is preemptable, *svPreemptable* returns a non-zero value, otherwise it returns zero.

If called from an interrupt level, *svIntrLevel* returns a non-zero value, otherwise it returns zero.

**RETURN VALUE**

If the lock is successfully acquired, *svSpinLockTry* returns the previous state of the lock (K\_AVAIL). The *svMaskedLockTry* function returns K\_AVAIL. The *svPreemptable* and *svIntrLevel* calls return a value which depends on the current execution context. Other calls return 0.

**ERRORS**

No error codes are returned. In DEBUG mode, the kernel debugger is invoked if the lock level hierarchy is violated, or if a spin lock cannot be acquired after a very large number of iterations, indicating a probable deadlock.

**RESTRICTIONS**

Hierarchy level checks are not performed in this version.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	svMemRead, svMemWrite – Copy from supervisor caller space; Copy to supervisor caller space				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int svMemRead(void * src, void * dest, int count);  int svMemWrite(void * src, void * dest, int count);</pre>				
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>These calls are used by programs to access supervisor address space memory safely.</p> <p>The <i>svMemRead</i> function copies <i>count</i> bytes from the address <i>src</i> in the supervisor address space, to the address <i>dest</i> in the supervisor address space. If access to the area defined by <i>src</i> and <i>count</i> in the supervisor address space produces a memory access fault, <i>svMemRead</i> will return an error code (K_EFAULT).</p> <p>The <i>svMemWrite</i> call copies <i>count</i> bytes from the address <i>src</i> in the supervisor address space, to the address <i>dest</i> in the supervisor address space. If access to the area defined by <i>dest</i> and <i>count</i> in the supervisor address space produces a memory access fault, <i>svMemWrite</i> will return an error code (K_EFAULT).</p>				
<b>RETURN VALUE</b>	If successful, <i>svMemRead</i> and <i>svMemWrite</i> return 0. Otherwise, they return K_EFAULT.				
<b>ERRORS</b>	[K_EFAULT]                      Access to the provided supervisor data produced a memory access fault.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				

<b>NAME</b>	svMemRead, svMemWrite – Copy from supervisor caller space; Copy to supervisor caller space				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int svMemRead(void * src, void * dest, int count);  int svMemWrite(void * src, void * dest, int count);</pre>				
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>These calls are used by programs to access supervisor address space memory safely.</p> <p>The <i>svMemRead</i> function copies <i>count</i> bytes from the address <i>src</i> in the supervisor address space, to the address <i>dest</i> in the supervisor address space. If access to the area defined by <i>src</i> and <i>count</i> in the supervisor address space produces a memory access fault, <i>svMemRead</i> will return an error code (K_EFAULT).</p> <p>The <i>svMemWrite</i> call copies <i>count</i> bytes from the address <i>src</i> in the supervisor address space, to the address <i>dest</i> in the supervisor address space. If access to the area defined by <i>dest</i> and <i>count</i> in the supervisor address space produces a memory access fault, <i>svMemWrite</i> will return an error code (K_EFAULT).</p>				
<b>RETURN VALUE</b>	If successful, <i>svMemRead</i> and <i>svMemWrite</i> return 0. Otherwise, they return K_EFAULT.				
<b>ERRORS</b>	[K_EFAULT]                      Access to the provided supervisor data produced a memory access fault.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				

<b>NAME</b>	svMsgHandler, svMsgHdlReply – Connect/disconnect a message handler; Prepare a reply to a handled message
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int svMsgHandler(KnCap * act, int portli, KnMsgHdl handler, unsigned int stacksize, void * cookie);  int svMsgHdlReply(KnMsgDesc * msg);</pre>
<b>FEATURES</b>	IPC
<b>DESCRIPTION</b>	<p>The <i>svMsgHandler</i> function connects a handler to the arrival of messages on the port whose local identifier is <i>portli</i> in the actor whose capability is given by <i>act</i>. If <i>act</i> is <i>K_MYACTOR</i>, the current actor is used.</p> <p>If <i>handler</i> is NULL, <i>svMsgHandler</i> disconnects the handler connected previously.</p> <p>The message handler will be called each time a message is posted to the port. If messages are queued behind the port at the time the handler is connected, the handler will be invoked for these messages.</p> <p>The <i>cookie</i> argument is user-defined and is passed as the last argument to the handler. This allows the application to identify the handler easily when multiple ports are using the same handler.</p> <p>When a message handler has been connected for a given port, <i>ipcReceive</i> (2K) calls on that port are refused (<i>ipcReceive</i> (2K) returns the <i>K_EINVAL</i> error code).</p> <p>A port to which a message handler is connected may not be migrated (see <i>portMigrate</i> (2K)).</p> <p>The stack size required by the handler execution is defined by <i>stacksize</i>. In the current version, this argument is ignored. The system only insures that when the handler is called, the system stack of the thread is a constant (defined at kernel generation time) number of free bytes.</p> <p>The handler is a pointer to a function whose arguments are the following:</p> <pre>int handler (request, src, dest, port, reply, cookie)     KnMsgDesc*   request ;     KnUniqueId*  src ;     KnIpcDest*   dest ;     int          port ;     KnMsgDesc*   reply ;     void*        cookie ;</pre> <p>The <i>request</i> message descriptor describes the message to be handled (the <i>KnMsgDesc</i> structure is described in <i>ipcSend</i> (2K)). The <i>seqNum</i> member of this structure represents the count of messages that have been received behind the port up to and including the message being handled (see also <i>ipcReceive</i> (2K)). If the <i>K_USERBODY</i> or <i>K_USERANNEX</i> flag is present in the <i>flags</i> field of</p>

this descriptor, the corresponding data must be obtained using *svCopyIn* (2K). Otherwise, the data may be read directly from the system address space.

The *src* pointer points to the UI describing the port from which the message has been sent.

The *dest* pointer points to the destination to which the message has been sent.

The *port* pointer is the local identifier of the port to which the handler is attached.

The location for a reply to the message may already have been allocated by the system or the sender. In this case, *reply* is a non NULL pointer to a message descriptor. Otherwise, *reply* is NULL. Note that the K\_USERBODY and K\_USERANNEX flags may be present in the *reply* message descriptor. In this case, the handler should use *svCopyOut* (2K) to update the reply message annex and/or body.

To have the system update the handler's reply data, or if *reply* is NULL, the handler can invoke the *svMsgHdlReply* function. This function updates the reply data according to the *msg* message description. This function does not cause the destinator of the reply to receive it: this operation is delayed until the handler exits. The value passed to *svMsgHandler* is defined by *cookie* .

The handler is executed by a system thread. This thread is a restricted form of a thread, a *handler thread*. The handler thread is temporarily attached to the actor which owns the port to which the message handler has been attached. Any reference to the "current actor" (K\_MYACTOR) within the context of the handler thread execution is interpreted as a reference to the actor which owns the port to which the handler is connected.

The handler thread possesses a local identifier within its actor, it may be obtained using *threadSelf* (2K). However, this thread is a restricted type of thread, its local identifier may not be used for external management of the thread ( *threadAbort* (2K), *threadSuspend* (2K), *threadResume* (2K), *threadContext* (2K), *threadTimes* (2K)). Applying these services to a handler thread results in the K\_EINVAL error code being returned (invalid thread)..

The only way to delete the handler thread is to return from the handler.

When entering the handler, the handler thread has no current message. This means that *ipcSave* (2K), *ipcReply* (2K) and *ipcSysInfo* (2K) return the K\_EINVAL error code. However, the handler thread can perform *ipcReceive* (2K) and get a current message. If the handler returns while the handler thread has a current message, this message is canceled.

When the handler exits, the handler return value defines the format of the reply to the sender's request, as follows:

If *svMsgHdlReply* was not called during handler execution, and the handler returns a positive or null value, the message body and annex are updated by the handler (using the *reply* argument). In this case, the handler returns the message body size. If the handler return value is a negative error code, the caller does not receive a reply; however, if the message was sent using *ipcCall* (2K), the negative error code returned by the handler is returned to the caller.

If *svMsgHdlReply* was called during handler execution, the handler return code is ignored and the caller receives zero if the call was successful or a negative error code (see below).

Both *svMsgHandler* and *svMsgHdlReply* are restricted to supervisor threads. The handler code must be located in the system address space.

**RETURN VALUE**

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

- [K\_EINVAL] ( *svMsgHandler* ) *actorcap* is an inconsistent actor capability, or *portli* is not a valid port identifier within the actor whose capability is given by *actorcap*. ( *svMsgHdlReply* ) the current thread is not a handler thread.
- [K\_EUNKNOWN] ( *svMsgHandler* only) *actorcap* does not specify a reachable actor.
- [K\_EFAULT] Some of the data provided are outside the current actor's address space.

**RESTRICTIONS**

Currently, if a handler processes an asynchronous message ( *ipcSend* ) or a message sent remotely, the invocation of *svMsgHdlReply* sends the reply to the caller, instead of preparing it and delaying it until the handler returns .

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*ipcCall*(2K) , *ipcSend*(2K) , *threadAbort*(2K) , *svCopyIn*(2K) , *svCopyOut*(2K)

<b>NAME</b>	svMsgHandler, svMsgHdlReply – Connect/disconnect a message handler; Prepare a reply to a handled message
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chIpc.h&gt; int svMsgHandler(KnCap * act, int portli, KnMsgHdl handler, unsigned int stacksize, void * cookie);  int svMsgHdlReply(KnMsgDesc * msg);</pre>
<b>FEATURES</b>	IPC
<b>DESCRIPTION</b>	<p>The <i>svMsgHandler</i> function connects a handler to the arrival of messages on the port whose local identifier is <i>portli</i> in the actor whose capability is given by <i>act</i>. If <i>act</i> is <i>K_MYACTOR</i>, the current actor is used.</p> <p>If <i>handler</i> is <i>NULL</i>, <i>svMsgHandler</i> disconnects the handler connected previously..</p> <p>The message handler will be called each time a message is posted to the port. If messages are queued behind the port at the time the handler is connected, the handler will be invoked for these messages.</p> <p>The <i>cookie</i> argument is user-defined and is passed as the last argument to the handler. This allows the application to identify the handler easily when multiple ports are using the same handler.</p> <p>When a message handler has been connected for a given port, <i>ipcReceive</i> (2K) calls on that port are refused ( <i>ipcReceive</i> (2K) returns the <i>K_EINVAL</i> error code).</p> <p>A port to which a message handler is connected may not be migrated (see <i>portMigrate</i> (2K)).</p> <p>The stack size required by the handler execution is defined by <i>stacksize</i> . In the current version, this argument is ignored. The system only insures that when the handler is called, the system stack of the thread is a constant (defined at kernel generation time) number of free bytes.</p> <p>The handler is a pointer to a function whose arguments are the following:</p> <pre>int handler (request, src, dest, port, reply, cookie)     KnMsgDesc*   request ;     KnUniqueId*  src ;     KnIpcDest*   dest ;     int          port ;     KnMsgDesc*   reply ;     void*        cookie ;</pre> <p>The <i>request</i> message descriptor describes the message to be handled (the <i>KnMsgDesc</i> structure is described in <i>ipcSend</i> (2K)). The <i>seqNum</i> member of this structure represents the count of messages that have been received behind the port up to and including the message being handled (see also <i>ipcReceive</i> (2K)). If the <i>K_USERBODY</i> or <i>K_USERANNEX</i> flag is present in the <i>flags</i> field of</p>

this descriptor, the corresponding data must be obtained using *svCopyIn* (2K). Otherwise, the data may be read directly from the system address space.

The *src* pointer points to the UI describing the port from which the message has been sent.

The *dest* pointer points to the destination to which the message has been sent.

The *port* pointer is the local identifier of the port to which the handler is attached.

The location for a reply to the message may already have been allocated by the system or the sender. In this case, *reply* is a non NULL pointer to a message descriptor. Otherwise, *reply* is NULL. Note that the K\_USERBODY and K\_USERANNEX flags may be present in the *reply* message descriptor. In this case, the handler should use *svCopyOut* (2K) to update the reply message annex and/or body.

To have the system update the handler's reply data, or if *reply* is NULL, the handler can invoke the *svMsgHdlReply* function. This function updates the reply data according to the *msg* message description. This function does not cause the destinator of the reply to receive it: this operation is delayed until the handler exits. The value passed to *svMsgHandler* is defined by *cookie*.

The handler is executed by a system thread. This thread is a restricted form of a thread, a *handler thread*. The handler thread is temporarily attached to the actor which owns the port to which the message handler has been attached. Any reference to the "current actor" (K\_MYACTOR) within the context of the handler thread execution is interpreted as a reference to the actor which owns the port to which the handler is connected.

The handler thread possesses a local identifier within its actor, it may be obtained using *threadSelf* (2K). However, this thread is a restricted type of thread, its local identifier may not be used for external management of the thread ( *threadAbort* (2K), *threadSuspend* (2K), *threadResume* (2K), *threadContext* (2K), *threadTimes* (2K)). Applying these services to a handler thread results in the K\_EINVAL error code being returned (invalid thread)..

The only way to delete the handler thread is to return from the handler.

When entering the handler, the handler thread has no current message. This means that *ipcSave* (2K), *ipcReply* (2K) and *ipcSysInfo* (2K) return the K\_EINVAL error code. However, the handler thread can perform *ipcReceive* (2K) and get a current message. If the handler returns while the handler thread has a current message, this message is canceled.

When the handler exits, the handler return value defines the format of the reply to the sender's request, as follows:

If *svMsgHdlReply* was not called during handler execution, and the handler returns a positive or null value, the message body and annex are updated by the handler (using the *reply* argument). In this case, the handler returns the message body size. If the handler return value is a negative error code, the caller does not receive a reply; however, if the message was sent using *ipcCall* (2K), the negative error code returned by the handler is returned to the caller.

If *svMsgHdlReply* was called during handler execution, the handler return code is ignored and the caller receives zero if the call was successful or a negative error code (see below).

Both *svMsgHandler* and *svMsgHdlReply* are restricted to supervisor threads. The handler code must be located in the system address space.

**RETURN VALUE**

Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

- [K\_EINVAL] ( *svMsgHandler* ) *actorcap* is an inconsistent actor capability, or *portli* is not a valid port identifier within the actor whose capability is given by *actorcap*. ( *svMsgHdlReply* ) the current thread is not a handler thread.
- [K\_EUNKNOWN] ( *svMsgHandler* only ) *actorcap* does not specify a reachable actor.
- [K\_EFAULT] Some of the data provided are outside the current actor's address space.

**RESTRICTIONS**

Currently, if a handler processes an asynchronous message ( *ipcSend* ) or a message sent remotely, the invocation of *svMsgHdlReply* sends the reply to the caller, instead of preparing it and delaying it until the handler returns .

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*ipcCall*(2K) , *ipcSend*(2K) , *threadAbort*(2K) , *svCopyIn*(2K) , *svCopyOut*(2K)

<b>NAME</b>	svPagesAllocate, svPagesFree – supervisor address space memory allocator
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int svPagesAllocate(VmAddr * address, VmSize size, VmFlags flags, KnCap * actcap); int svPagesFree(VmAddr address, VmSize size, KnCap * actcap);</pre>
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL
<b>DESCRIPTION</b>	<p>The <i>svPagesAllocate</i> function allocates an amount of (usually discontinuous) physical memory and a range of virtual addresses within the supervisor address space. It then maps the physical memory to this address range.</p> <p>The <i>size</i> argument specifies the size (in bytes) of the memory to be allocated. There are no alignment requirements on the <i>size</i> value, but <i>svPagesAllocate</i> always aligns the size of allocated memory to the next page boundary.</p> <p>The starting address of the allocated virtual memory range is returned in the variable pointed to by the <i>address</i> argument.</p> <p>The <i>flags</i> argument is a combination of the following flags:</p> <p><b>K_NOWAITFORMEMORY</b> If the K_NOWAITFORMEMORY flag is set, <i>svPagesAllocate</i> fails when there is not enough physical memory. In this case it returns the K_ENOMEM error; otherwise, <i>svPagesAllocate</i> waits for memory to become available as the result of memory swap out activity. If the swap outs is not available <i>svPagesAllocate</i> will never wait.</p> <p><b>K_REDALLOC</b> This flag is for internal operating system usage only.</p> <p>The <i>actcap</i> argument specifies an actor with which the allocated memory will be associated. It is required for various functions, such as actor dumps. Note that allocated memory is not automatically freed when the corresponding actor is deleted using <i>actorDelete</i> (2K). If the actor association is not needed, K_SVACTOR can be specified. In this case the memory will be anonymous (not linked to any actor).</p> <p>The <i>svPagesFree</i> call releases a virtual address range and the physical memory mapped to it. The <i>address</i> argument specifies the range's starting address and the <i>size</i> argument specifies its size in bytes.</p> <p>The address range must be a sub-range of an address range previously allocated by a <i>svPagesAllocate</i> call. The range start address must be page aligned but there</p>

are no alignment requirements on the range size. Nevertheless *svPagesFree* always aligns the size of deallocated memory to the next page boundary.

The *svPagesAllocate* and *svPagesFree* calls may be called by an interrupt handler. In this case *svPagesAllocate* must be called with `K_NOWAITFORMEMORY`.

The *actcap* argument specifies the actor from which the pages must be detached. It must match the actor that was specified when the pages were allocated.

**RETURN VALUES**

If successful `K_OK` is returned, otherwise a negative error code is returned.

**ERRORS**

[ <code>K_EFAULT</code> ]	The <i>address</i> argument points to the outside of the caller's address space.
[ <code>K_EINVAL</code> ]	<i>svPagesAllocate</i> was called by an interrupt handler and <code>K_NOWAITFORMEMORY</code> flag was not specified.
[ <code>K_EROUND</code> ]	The <i>address</i> argument of <i>svPagesFree</i> is not page aligned.
[ <code>K_ESIZE</code> ]	The <i>size</i> argument is equal to 0.
[ <code>K_ENOMEM</code> ]	The system is out of memory.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`vmSetPar(2K)`

<b>NAME</b>	svPagesAllocate, svPagesFree – supervisor address space memory allocator
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int svPagesAllocate(VmAddr * address, VmSize size, VmFlags flags, KnCap * actcap); int svPagesFree(VmAddr address, VmSize size, KnCap * actcap);</pre>
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL
<b>DESCRIPTION</b>	<p>The <i>svPagesAllocate</i> function allocates an amount of (usually discontinuous) physical memory and a range of virtual addresses within the supervisor address space. It then maps the physical memory to this address range.</p> <p>The <i>size</i> argument specifies the size (in bytes) of the memory to be allocated. There are no alignment requirements on the <i>size</i> value, but <i>svPagesAllocate</i> always aligns the size of allocated memory to the next page boundary.</p> <p>The starting address of the allocated virtual memory range is returned in the variable pointed to by the <i>address</i> argument.</p> <p>The <i>flags</i> argument is a combination of the following flags:</p> <p><b>K_NOWAITFORMEMORY</b> If the K_NOWAITFORMEMORY flag is set, <i>svPagesAllocate</i> fails when there is not enough physical memory. In this case it returns the K_ENOMEM error; otherwise, <i>svPagesAllocate</i> waits for memory to become available as the result of memory swap out activity. If the swap outs is not available <i>svPagesAllocate</i> will never wait.</p> <p><b>K_REDALLOC</b> This flag is for internal operating system usage only.</p> <p>The <i>actcap</i> argument specifies an actor with which the allocated memory will be associated. It is required for various functions, such as actor dumps. Note that allocated memory is not automatically freed when the corresponding actor is deleted using <i>actorDelete</i> (2K). If the actor association is not needed, K_SVACTOR can be specified. In this case the memory will be anonymous (not linked to any actor).</p> <p>The <i>svPagesFree</i> call releases a virtual address range and the physical memory mapped to it. The <i>address</i> argument specifies the range's starting address and the <i>size</i> argument specifies its size in bytes.</p> <p>The address range must be a sub-range of an address range previously allocated by a <i>svPagesAllocate</i> call. The range start address must be page aligned but there</p>

are no alignment requirements on the range size. Nevertheless *svPagesFree* always aligns the size of deallocated memory to the next page boundary.

The *svPagesAllocate* and *svPagesFree* calls may be called by an interrupt handler. In this case *svPagesAllocate* must be called with `K_NOWAITFORMEMORY`.

The *actcap* argument specifies the actor from which the pages must be detached. It must match the actor that was specified when the pages were allocated.

**RETURN VALUES**

If successful `K_OK` is returned, otherwise a negative error code is returned.

**ERRORS**

[ <code>K_EFAULT</code> ]	The <i>address</i> argument points to the outside of the caller's address space.
[ <code>K_EINVAL</code> ]	<i>svPagesAllocate</i> was called by an interrupt handler and <code>K_NOWAITFORMEMORY</code> flag was not specified.
[ <code>K_EROUND</code> ]	The <i>address</i> argument of <i>svPagesFree</i> is not page aligned.
[ <code>K_ESIZE</code> ]	The <i>size</i> argument is equal to 0.
[ <code>K_ENOMEM</code> ]	The system is out of memory.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`vmSetPar(2K)`

<b>NAME</b>	svMaskedLockInit, svMaskedLockGet, svMaskedLockTry, svMaskedLockRel, svSpinLockInit, svSpinLockGet, svSpinLockTry, svSpinLockRel – Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSpinLock.h&gt; void <b>svMaskedLockInit</b>(KnMaskedSpinLock * <i>lockaddr</i>, KnMaskedSpinLockInfo * <i>info</i>);  void <b>svMaskedLockGet</b>(KnMaskedSpinLock * <i>lockaddr</i>);  int <b>svMaskedLockTry</b>(KnMaskedSpinLock * <i>lockaddr</i>);  void <b>svMaskedLockRel</b>(KnMaskedSpinLock * <i>lockaddr</i>);  void <b>svSpinLockInit</b>(KnSpinLock * <i>lockaddr</i>, KnSpinLockInfo * <i>info</i>);  void <b>svSpinLockGet</b>(KnSpinLock * <i>lockaddr</i>);  int <b>svSpinLockTry</b>(KnSpinLock * <i>lockaddr</i>);  void <b>svSpinLockRel</b>(KnSpinLock * <i>lockaddr</i>);  int <b>svPreemptable</b>(void);  int <b>svIntrLevel</b>(void);</pre>
<b>DESCRIPTION</b>	<p>Spin locks are mutual exclusion objects, used to protect very short critical sections. These critical sections must not include kernel calls and the data they manipulate must be locked in memory to avoid page faults.</p> <p>In multiprocessor systems, threads loop actively when resources are not available. A spin lock is an object in memory, typically modified and checked through indivisible machine operations (such as test-and-set).</p> <p>In monoprocessor systems, spin locks are implemented by disabling preemption or interrupts.</p> <p>Note that while a thread is holding a spin lock, a masked spin lock, or has disabled preemption, it is no longer allowed to block. In this case, the list of system calls available to the thread is restricted to the following:</p> <pre>semV threadSemPost eventPost msgPut svSpinLockGet svSpinLockRel svMaskedLockGet svMaskedLockRel svMaskAll svUnmaskAll svUnmask</pre>

```
svPreemptable  
svIntrLevel  
svTimeoutSet  
svTimeoutCancel  
svPagesAllocate
```

The system not force the thread to restrict itself to the permitted calls. If the thread issues non-permitted calls, the system will hang.

The *svMaskedLockInit* and *svSpinLockInit* functions initialize a spin lock in the free state. The *info* argument allows the user to enforce lock hierarchy and/or to gather lock use/contention statistics in the future.

In multiprocessor systems, *svMaskedLockGet* atomically disables interrupts on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with interrupts enabled to reduce interrupt latency. This function can be called from within an interrupt handler.

The *svMaskedLockTry* function attempts to acquire a spin lock, but returns immediately without spinning if there is none available. This function can be called from an interrupt handler.

The *svMaskedLockRel* function releases a spin lock and re-enables interrupts on the current processor. This function can be called from within an interrupt handler.

The *svMaskedLockGet*, *svMaskedLockTry* and *svMaskedLockRel* functions only relate to the interrupt mask of the current processor, and do not affect interrupt processing on other processors. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.

In multiprocessor systems, *svSpinLockGet* atomically disables preemption on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with preemption enabled to reduce scheduling latency. On a monoprocessor, *svSpinLockGet* disables preemption on the current processor. This function cannot be called from within an interrupt handler.

The *svSpinLockTry* try attempts to acquire a spin lock, returning immediately without spinning if there is none available. This function cannot be called from within an interrupt handler.

The *svSpinLockRel* function releases a spin lock. This function ca not be called from within an interrupt handler.

If the current thread is preemptable, *svPreemptable* returns a non-zero value, otherwise it returns zero.

If called from an interrupt level, *svIntrLevel* returns a non-zero value, otherwise it returns zero.

**RETURN VALUE**

If the lock is successfully acquired, *svSpinLockTry* returns the previous state of the lock (K\_AVAIL). The *svMaskedLockTry* function returns K\_AVAIL. The *svPreemptable* and *svIntrLevel* calls return a value which depends on the current execution context. Other calls return 0.

**ERRORS**

No error codes are returned. In DEBUG mode, the kernel debugger is invoked if the lock level hierarchy is violated, or if a spin lock cannot be acquired after a very large number of iterations, indicating a probable deadlock.

**RESTRICTIONS**

Hierarchy level checks are not performed in this version.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**NAME** svMaskedLockInit, svMaskedLockGet, svMaskedLockTry, svMaskedLockRel, svSpinLockInit, svSpinLockGet, svSpinLockTry, svSpinLockRel – Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock

**SYNOPSIS**

```
#include <sync/chSpinLock.h>
void svMaskedLockInit(KnMaskedSpinLock * lockaddr, KnMaskedSpinLockInfo * info);

void svMaskedLockGet(KnMaskedSpinLock * lockaddr);

int svMaskedLockTry(KnMaskedSpinLock * lockaddr);

void svMaskedLockRel(KnMaskedSpinLock * lockaddr);

void svSpinLockInit(KnSpinLock * lockaddr, KnSpinLockInfo * info);

void svSpinLockGet(KnSpinLock * lockaddr);

int svSpinLockTry(KnSpinLock * lockaddr);

void svSpinLockRel(KnSpinLock * lockaddr);

int svPreemptable(void);

int svIntrLevel(void);
```

**DESCRIPTION**

Spin locks are mutual exclusion objects, used to protect very short critical sections. These critical sections must not include kernel calls and the data they manipulate must be locked in memory to avoid page faults.

In multiprocessor systems, threads loop actively when resources are not available. A spin lock is an object in memory, typically modified and checked through indivisible machine operations (such as test-and-set).

In monoprocessor systems, spin locks are implemented by disabling preemption or interrupts.

Note that while a thread is holding a spin lock, a masked spin lock, or has disabled preemption, it is no longer allowed to block. In this case, the list of system calls available to the thread is restricted to the following:

```
semV
threadSemPost
eventPost
msgPut
svSpinLockGet
svSpinLockRel
svMaskedLockGet
svMaskedLockRel
svMaskAll
svUnmaskAll
svUnmask
```

```
svPreemptable  
svIntrLevel  
svTimeoutSet  
svTimeoutCancel  
svPagesAllocate
```

The system not force the thread to restrict itself to the permitted calls. If the thread issues non-permitted calls, the system will hang.

The *svMaskedLockInit* and *svSpinLockInit* functions initialize a spin lock in the free state. The *info* argument allows the user to enforce lock hierarchy and/or to gather lock use/contention statistics in the future.

In multiprocessor systems, *svMaskedLockGet* atomically disables interrupts on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with interrupts enabled to reduce interrupt latency. This function can be called from within an interrupt handler.

The *svMaskedLockTry* function attempts to acquire a spin lock, but returns immediately without spinning if there is none available. This function can be called from an interrupt handler.

The *svMaskedLockRel* function releases a spin lock and re-enables interrupts on the current processor. This function can be called from within an interrupt handler.

The *svMaskedLockGet*, *svMaskedLockTry* and *svMaskedLockRel* functions only relate to the interrupt mask of the current processor, and do not affect interrupt processing on other processors. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.

In multiprocessor systems, *svSpinLockGet* atomically disables preemption on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with preemption enabled to reduce scheduling latency. On a monoprocessor, *svSpinLockGet* disables preemption on the current processor. This function cannot be called from within an interrupt handler.

The *svSpinLockTry* try attempts to acquire a spin lock, returning immediately without spinning if there is none available. This function cannot be called from within an interrupt handler.

The *svSpinLockRel* function releases a spin lock. This function ca not be called from within an interrupt handler.

If the current thread is preemptable, *svPreemptable* returns a non-zero value, otherwise it returns zero.

If called from an interrupt level, *svIntrLevel* returns a non-zero value, otherwise it returns zero.

**RETURN VALUE**

If the lock is successfully acquired, *svSpinLockTry* returns the previous state of the lock (K\_AVAIL). The *svMaskedLockTry* function returns K\_AVAIL. The *svPreemptable* and *svIntrLevel* calls return a value which depends on the current execution context. Other calls return 0.

**ERRORS**

No error codes are returned. In DEBUG mode, the kernel debugger is invoked if the lock level hierarchy is violated, or if a spin lock cannot be acquired after a very large number of iterations, indicating a probable deadlock.

**RESTRICTIONS**

Hierarchy level checks are not performed in this version.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	svMaskedLockInit, svMaskedLockGet, svMaskedLockTry, svMaskedLockRel, svSpinLockInit, svSpinLockGet, svSpinLockTry, svSpinLockRel – Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSpinLock.h&gt; void <b>svMaskedLockInit</b>(KnMaskedSpinLock * <i>lockaddr</i>, KnMaskedSpinLockInfo * <i>info</i>);  void <b>svMaskedLockGet</b>(KnMaskedSpinLock * <i>lockaddr</i>);  int <b>svMaskedLockTry</b>(KnMaskedSpinLock * <i>lockaddr</i>);  void <b>svMaskedLockRel</b>(KnMaskedSpinLock * <i>lockaddr</i>);  void <b>svSpinLockInit</b>(KnSpinLock * <i>lockaddr</i>, KnSpinLockInfo * <i>info</i>);  void <b>svSpinLockGet</b>(KnSpinLock * <i>lockaddr</i>);  int <b>svSpinLockTry</b>(KnSpinLock * <i>lockaddr</i>);  void <b>svSpinLockRel</b>(KnSpinLock * <i>lockaddr</i>);  int <b>svPreemptable</b>(void);  int <b>svIntrLevel</b>(void);</pre>
<b>DESCRIPTION</b>	<p>Spin locks are mutual exclusion objects, used to protect very short critical sections. These critical sections must not include kernel calls and the data they manipulate must be locked in memory to avoid page faults.</p> <p>In multiprocessor systems, threads loop actively when resources are not available. A spin lock is an object in memory, typically modified and checked through indivisible machine operations (such as test-and-set).</p> <p>In monoprocessor systems, spin locks are implemented by disabling preemption or interrupts.</p> <p>Note that while a thread is holding a spin lock, a masked spin lock, or has disabled preemption, it is no longer allowed to block. In this case, the list of system calls available to the thread is restricted to the following:</p> <pre>semV threadSemPost eventPost msgPut svSpinLockGet svSpinLockRel svMaskedLockGet svMaskedLockRel svMaskAll svUnmaskAll svUnmask</pre>

```
svPreemptable  
svIntrLevel  
svTimeoutSet  
svTimeoutCancel  
svPagesAllocate
```

The system not force the thread to restrict itself to the permitted calls. If the thread issues non-permitted calls, the system will hang.

The *svMaskedLockInit* and *svSpinLockInit* functions initialize a spin lock in the free state. The *info* argument allows the user to enforce lock hierarchy and/or to gather lock use/contention statistics in the future.

In multiprocessor systems, *svMaskedLockGet* atomically disables interrupts on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with interrupts enabled to reduce interrupt latency. This function can be called from within an interrupt handler.

The *svMaskedLockTry* function attempts to acquire a spin lock, but returns immediately without spinning if there is none available. This function can be called from an interrupt handler.

The *svMaskedLockRel* function releases a spin lock and re-enables interrupts on the current processor. This function can be called from within an interrupt handler.

The *svMaskedLockGet*, *svMaskedLockTry* and *svMaskedLockRel* functions only relate to the interrupt mask of the current processor, and do not affect interrupt processing on other processors. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.

In multiprocessor systems, *svSpinLockGet* atomically disables preemption on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with preemption enabled to reduce scheduling latency. On a monoprocessor, *svSpinLockGet* disables preemption on the current processor. This function cannot be called from within an interrupt handler.

The *svSpinLockTry* try attempts to acquire a spin lock, returning immediately without spinning if there is none available. This function cannot be called from within an interrupt handler.

The *svSpinLockRel* function releases a spin lock. This function ca not be called from within an interrupt handler.

If the current thread is preemptable, *svPreemptable* returns a non-zero value, otherwise it returns zero.

If called from an interrupt level, *svIntrLevel* returns a non-zero value, otherwise it returns zero.

**RETURN VALUE**

If the lock is successfully acquired, *svSpinLockTry* returns the previous state of the lock (K\_AVAIL). The *svMaskedLockTry* function returns K\_AVAIL. The *svPreemptable* and *svIntrLevel* calls return a value which depends on the current execution context. Other calls return 0.

**ERRORS**

No error codes are returned. In DEBUG mode, the kernel debugger is invoked if the lock level hierarchy is violated, or if a spin lock cannot be acquired after a very large number of iterations, indicating a probable deadlock.

**RESTRICTIONS**

Hierarchy level checks are not performed in this version.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**NAME** svMaskedLockInit, svMaskedLockGet, svMaskedLockTry, svMaskedLockRel, svSpinLockInit, svSpinLockGet, svSpinLockTry, svSpinLockRel – Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock

**SYNOPSIS**

```
#include <sync/chSpinLock.h>
void svMaskedLockInit(KnMaskedSpinLock * lockaddr, KnMaskedSpinLockInfo * info);

void svMaskedLockGet(KnMaskedSpinLock * lockaddr);

int svMaskedLockTry(KnMaskedSpinLock * lockaddr);

void svMaskedLockRel(KnMaskedSpinLock * lockaddr);

void svSpinLockInit(KnSpinLock * lockaddr, KnSpinLockInfo * info);

void svSpinLockGet(KnSpinLock * lockaddr);

int svSpinLockTry(KnSpinLock * lockaddr);

void svSpinLockRel(KnSpinLock * lockaddr);

int svPreemptable(void);

int svIntrLevel(void);
```

**DESCRIPTION**

Spin locks are mutual exclusion objects, used to protect very short critical sections. These critical sections must not include kernel calls and the data they manipulate must be locked in memory to avoid page faults.

In multiprocessor systems, threads loop actively when resources are not available. A spin lock is an object in memory, typically modified and checked through indivisible machine operations (such as test-and-set).

In monoprocessor systems, spin locks are implemented by disabling preemption or interrupts.

Note that while a thread is holding a spin lock, a masked spin lock, or has disabled preemption, it is no longer allowed to block. In this case, the list of system calls available to the thread is restricted to the following:

```
semV
threadSemPost
eventPost
msgPut
svSpinLockGet
svSpinLockRel
svMaskedLockGet
svMaskedLockRel
svMaskAll
svUnmaskAll
svUnmask
```

```
svPreemptable  
svIntrLevel  
svTimeoutSet  
svTimeoutCancel  
svPagesAllocate
```

The system not force the thread to restrict itself to the permitted calls. If the thread issues non-permitted calls, the system will hang.

The *svMaskedLockInit* and *svSpinLockInit* functions initialize a spin lock in the free state. The *info* argument allows the user to enforce lock hierarchy and/or to gather lock use/contention statistics in the future.

In multiprocessor systems, *svMaskedLockGet* atomically disables interrupts on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with interrupts enabled to reduce interrupt latency. This function can be called from within an interrupt handler.

The *svMaskedLockTry* function attempts to acquire a spin lock, but returns immediately without spinning if there is none available. This function can be called from an interrupt handler.

The *svMaskedLockRel* function releases a spin lock and re-enables interrupts on the current processor. This function can be called from within an interrupt handler.

The *svMaskedLockGet*, *svMaskedLockTry* and *svMaskedLockRel* functions only relate to the interrupt mask of the current processor, and do not affect interrupt processing on other processors. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.

In multiprocessor systems, *svSpinLockGet* atomically disables preemption on the current processor and acquires a spin lock. This routine spins actively on the lock value until it becomes free. Note that spinning is performed with preemption enabled to reduce scheduling latency. On a monoprocessor, *svSpinLockGet* disables preemption on the current processor. This function cannot be called from within an interrupt handler.

The *svSpinLockTry* try attempts to acquire a spin lock, returning immediately without spinning if there is none available. This function cannot be called from within an interrupt handler.

The *svSpinLockRel* function releases a spin lock. This function ca not be called from within an interrupt handler.

If the current thread is preemptable, *svPreemptable* returns a non-zero value, otherwise it returns zero.

If called from an interrupt level, *svIntrLevel* returns a non-zero value, otherwise it returns zero.

**RETURN VALUE**

If the lock is successfully acquired, *svSpinLockTry* returns the previous state of the lock (K\_AVAIL). The *svMaskedLockTry* function returns K\_AVAIL. The *svPreemptable* and *svIntrLevel* calls return a value which depends on the current execution context. Other calls return 0.

**ERRORS**

No error codes are returned. In DEBUG mode, the kernel debugger is invoked if the lock level hierarchy is violated, or if a spin lock cannot be acquired after a very large number of iterations, indicating a probable deadlock.

**RESTRICTIONS**

Hierarchy level checks are not performed in this version.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**NAME** svSysCtx – get system context table address

**SYNOPSIS** `#include <exec/chContext.h>`  
`KnChorusContext *svSysCtx(void);`

**DESCRIPTION** The *svSysCtx* function returns the address of the system context table which structure is defined in *exec/chContext.h*.

**RETURN VALUE** Address of the system context table.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	svSysPanic – trigger the invocation of the panic handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chPanic.h&gt; int svSysPanic(unsigned panicNo /* panic number */, void *args /* panic argument */, unsigned argsSize /* panic argument size */, char *mess /* panic message */, char *fileName /* file name */, int lineNumber /* line number */);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svSysPanic</i> system call allows Chorus applications to trigger the execution of the kernel panic handling sequence. When <i>svSysPanic</i> is invoked by an application, the kernel executes the panic handling strategy, just as if the panic had been detected by the kernel itself. If a panic handler has been connected by the BSP the kernel invokes this handler, passing the <i>svSysPanic</i> arguments unchanged. If no panic handler has been connected, the kernel performs default panic handling (usually by entering the kernel debugger).</p> <p>The <i>panicNo</i> field identifies the panic. It should normally take a value not already assigned to Chorus components, except in the case of generic panics such as debug assertions, for which a single generic number has been assigned.</p> <p>The <i>panicArgs</i> member is a pointer to a panic-specific data structure. The <i>panicArgsSize</i> member indicates the size of the panic-specific structure.</p> <p>This kernel service is only available to supervisor threads.</p>
<b>RETURN VALUE</b>	The <i>svSysPanic</i> system call never returns.
<b>RESTRICTIONS</b>	This call is restricted to <i>SUPERVISOR</i> threads.
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	svSysTimeout, svSysTimeoutSet, svSysTimeoutCancel, svTimeoutGetRes – Request a timeout; Cancel a timeout; Get timeout resolution
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTimeout.h&gt; int svSysTimeoutSet(KnTimeout * timeout, KnTimeVal * waitLimit, int flag, KnLapDesc * lapdesc);  int svSysTimeoutCancel(KnTimeout * timeout);  int svTimeoutGetRes(KnTimeVal * resolution);</pre>
<b>FEATURES</b>	TIMEOUT
<b>DESCRIPTION</b>	<p>The <code>svSysTimeoutSet()</code> system call declares <i>lapdesc</i> as the descriptor of a raw lap (see <code>svLapCreate(2K)</code>), which must be invoked when the interval of time specified by <i>*waitLimit</i> (see <code>intro(2K)</code>) has elapsed.</p> <p>The <i>waitLimit</i> pointer argument refers to a <code>KnTimeVal</code> structure whose members are defined in <code>sysTime(2K)</code>. If the <code>KnTimeVal</code> structure indicated by <i>waitLimit</i> is not equal to the <i>resolution</i>, then the time value composed by the <code>tmSec</code> and <code>tmNSec</code> fields is rounded. The fields are rounded down to the nearest multiple of the values in <i>resolution</i> unless the time is zero, in which case they are rounded up to the values in <i>resolution</i>.</p> <p>The lap handler of the <i>lapdesc</i> descriptor is a special kind of interrupt handler which is executed solely in SUPERVISOR execution mode, its code and accessed data must be within the locked-in-memory regions of a SUPERVISOR actor. The set of kernel calls that can be used in this timeout handler is limited, as for any interrupt handler.</p> <p>The argument passed to the lap handler is the <i>timeout</i> parameter specified by <code>svSysTimeoutSet()</code>. Fields within the <code>KnTimeout</code> structure are initialized and modified solely within the nucleus and are inaccessible to the application.</p> <p>The <i>flag</i> argument is reserved for future use and should therefore be set to 0 to ensure compatibility with future releases of the system.</p> <p>The <code>svSysTimeoutCancel()</code> system call attempts to cancel a timeout request. It takes as an argument the address of the <code>KnTimeout</code> object used in the call to <code>svSysTimeoutSet()</code>. If the timeout request is still pending, it is immediately canceled and a value is returned showing that the timer had not yet expired. Otherwise, the timeout interval had expired and the timeout handler was called, in which case nothing occurs and the value returned indicates this.</p> <p>The <code>svTimeoutGetRes()</code> system call obtains the resolution of the <i>*waitLimit</i> argument of <code>svSysTimeoutSet()</code>. The time value returned in <i>resolution</i> represents the smallest possible difference between two distinct <i>*waitLimit</i> values.</p>

The `svSysTimeoutSet()`, `svTimeoutGetRes()` and `svSysTimeoutCancel()` system calls are restricted to SUPERVISOR threads.

The `svSysTimeoutSet()` system call must not be called using a `KnTimeout` which has an existing timeout pending (use `svSysTimeoutCancel()` to cancel the request first).

**RETURN VALUES**

Upon completion, `svSysTimeoutSet()` returns `K_OK` if the timeout was successfully entered. Otherwise, a negative error code is returned specifying the reason it was not successful. Upon completion, `svSysTimeoutCancel()` returns a boolean value set to 1 if the request was still pending, or 0 if the timeout handler was called.

**ERRORS**

[`K_EINVAL`] The *\*waitLimit* value specified is invalid, or *lapdesc* is not a valid raw lap descriptor.

The value indicated by *\*waitLimit* is invalid if either the `tmSec` or `tmNSec` fields of the structure are negative, or if the structure is poorly constructed (for example, if the number of nanoseconds in `tmNSec` makes `tmNSec` greater than one second).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svIntrConnect(2K)`, `svLapCreate(2K)`, `sysTime(2K)`

<b>NAME</b>	svSysTimeout, svSysTimeoutSet, svSysTimeoutCancel, svTimeoutGetRes – Request a timeout; Cancel a timeout; Get timeout resolution
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTimeout.h&gt; int svSysTimeoutSet(KnTimeout * timeout, KnTimeVal * waitLimit, int flag, KnLapDesc * lapdesc);  int svSysTimeoutCancel(KnTimeout * timeout);  int svTimeoutGetRes(KnTimeVal * resolution);</pre>
<b>FEATURES</b>	TIMEOUT
<b>DESCRIPTION</b>	<p>The <code>svSysTimeoutSet()</code> system call declares <i>lapdesc</i> as the descriptor of a raw lap (see <code>svLapCreate(2K)</code>), which must be invoked when the interval of time specified by <i>*waitLimit</i> (see <code>intro(2K)</code>) has elapsed.</p> <p>The <i>waitLimit</i> pointer argument refers to a <code>KnTimeVal</code> structure whose members are defined in <code>sysTime(2K)</code>. If the <code>KnTimeVal</code> structure indicated by <i>waitLimit</i> is not equal to the <i>resolution</i>, then the time value composed by the <code>tmSec</code> and <code>tmNSec</code> fields is rounded. The fields are rounded down to the nearest multiple of the values in <i>resolution</i> unless the time is zero, in which case they are rounded up to the values in <i>resolution</i>.</p> <p>The lap handler of the <i>lapdesc</i> descriptor is a special kind of interrupt handler which is executed solely in SUPERVISOR execution mode, its code and accessed data must be within the locked-in-memory regions of a SUPERVISOR actor. The set of kernel calls that can be used in this timeout handler is limited, as for any interrupt handler.</p> <p>The argument passed to the lap handler is the <i>timeout</i> parameter specified by <code>svSysTimeoutSet()</code>. Fields within the <code>KnTimeout</code> structure are initialized and modified solely within the nucleus and are inaccessible to the application.</p> <p>The <i>flag</i> argument is reserved for future use and should therefore be set to 0 to ensure compatibility with future releases of the system.</p> <p>The <code>svSysTimeoutCancel()</code> system call attempts to cancel a timeout request. It takes as an argument the address of the <code>KnTimeout</code> object used in the call to <code>svSysTimeoutSet()</code>. If the timeout request is still pending, it is immediately canceled and a value is returned showing that the timer had not yet expired. Otherwise, the timeout interval had expired and the timeout handler was called, in which case nothing occurs and the value returned indicates this.</p> <p>The <code>svTimeoutGetRes()</code> system call obtains the resolution of the <i>*waitLimit</i> argument of <code>svSysTimeoutSet()</code>. The time value returned in <i>resolution</i> represents the smallest possible difference between two distinct <i>*waitLimit</i> values.</p>

The `svSysTimeoutSet()`, `svTimeoutGetRes()` and `svSysTimeoutCancel()` system calls are restricted to SUPERVISOR threads.

The `svSysTimeoutSet()` system call must not be called using a `KnTimeout` which has an existing timeout pending (use `svSysTimeoutCancel()` to cancel the request first).

**RETURN VALUES**

Upon completion, `svSysTimeoutSet()` returns `K_OK` if the timeout was successfully entered. Otherwise, a negative error code is returned specifying the reason it was not successful. Upon completion, `svSysTimeoutCancel()` returns a boolean value set to 1 if the request was still pending, or 0 if the timeout handler was called.

**ERRORS**

[`K_EINVAL`] The *\*waitLimit* value specified is invalid, or *lapdesc* is not a valid raw lap descriptor.

The value indicated by *\*waitLimit* is invalid if either the `tmSec` or `tmNSec` fields of the structure are negative, or if the structure is poorly constructed (for example, if the number of nanoseconds in `tmNSec` makes `tmNSec` greater than one second).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svIntrConnect(2K)`, `svLapCreate(2K)`, `sysTime(2K)`

<b>NAME</b>	svSysTimeout, svSysTimeoutSet, svSysTimeoutCancel, svTimeoutGetRes – Request a timeout; Cancel a timeout; Get timeout resolution
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTimeout.h&gt; int svSysTimeoutSet(KnTimeout * timeout, KnTimeVal * waitLimit, int flag, KnLapDesc * lapdesc);  int svSysTimeoutCancel(KnTimeout * timeout);  int svTimeoutGetRes(KnTimeVal * resolution);</pre>
<b>FEATURES</b>	TIMEOUT
<b>DESCRIPTION</b>	<p>The <code>svSysTimeoutSet()</code> system call declares <i>lapdesc</i> as the descriptor of a raw lap (see <code>svLapCreate(2K)</code>), which must be invoked when the interval of time specified by <i>*waitLimit</i> (see <code>intro(2K)</code>) has elapsed.</p> <p>The <i>waitLimit</i> pointer argument refers to a <code>KnTimeVal</code> structure whose members are defined in <code>sysTime(2K)</code>. If the <code>KnTimeVal</code> structure indicated by <i>waitLimit</i> is not equal to the <i>resolution</i>, then the time value composed by the <code>tmSec</code> and <code>tmNSec</code> fields is rounded. The fields are rounded down to the nearest multiple of the values in <i>resolution</i> unless the time is zero, in which case they are rounded up to the values in <i>resolution</i>.</p> <p>The lap handler of the <i>lapdesc</i> descriptor is a special kind of interrupt handler which is executed solely in SUPERVISOR execution mode, its code and accessed data must be within the locked-in-memory regions of a SUPERVISOR actor. The set of kernel calls that can be used in this timeout handler is limited, as for any interrupt handler.</p> <p>The argument passed to the lap handler is the <i>timeout</i> parameter specified by <code>svSysTimeoutSet()</code>. Fields within the <code>KnTimeout</code> structure are initialized and modified solely within the nucleus and are inaccessible to the application.</p> <p>The <i>flag</i> argument is reserved for future use and should therefore be set to 0 to ensure compatibility with future releases of the system.</p> <p>The <code>svSysTimeoutCancel()</code> system call attempts to cancel a timeout request. It takes as an argument the address of the <code>KnTimeout</code> object used in the call to <code>svSysTimeoutSet()</code>. If the timeout request is still pending, it is immediately canceled and a value is returned showing that the timer had not yet expired. Otherwise, the timeout interval had expired and the timeout handler was called, in which case nothing occurs and the value returned indicates this.</p> <p>The <code>svTimeoutGetRes()</code> system call obtains the resolution of the <i>*waitLimit</i> argument of <code>svSysTimeoutSet()</code>. The time value returned in <i>resolution</i> represents the smallest possible difference between two distinct <i>*waitLimit</i> values.</p>

The `svSysTimeoutSet()`, `svTimeoutGetRes()` and `svSysTimeoutCancel()` system calls are restricted to SUPERVISOR threads.

The `svSysTimeoutSet()` system call must not be called using a `KnTimeout` which has an existing timeout pending (use `svSysTimeoutCancel()` to cancel the request first).

**RETURN VALUES**

Upon completion, `svSysTimeoutSet()` returns `K_OK` if the timeout was successfully entered. Otherwise, a negative error code is returned specifying the reason it was not successful. Upon completion, `svSysTimeoutCancel()` returns a boolean value set to 1 if the request was still pending, or 0 if the timeout handler was called.

**ERRORS**

[K\_EINVAL] The *\*waitLimit* value specified is invalid, or *lapdesc* is not a valid raw lap descriptor.

The value indicated by *\*waitLimit* is invalid if either the `tmSec` or `tmNSec` fields of the structure are negative, or if the structure is poorly constructed (for example, if the number of nanoseconds in `tmNSec` makes `tmNSec` greater than one second).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svIntrConnect(2K)`, `svLapCreate(2K)`, `sysTime(2K)`

<b>NAME</b>	svSysTrapHandler, svSysTrapHandlerConnect, svSysTrapHandlerDisconnect, svSysTrapHandlerGetConnected – Connect a trap handler; Disconnect a trap handler; Get a trap handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTrap.h&gt; int svSysTrapHandlerConnect(unsigned int trapnumber, KnLapDesc * newhandler);  int svSysTrapHandlerDisconnect(unsigned int trapnumber, KnLapDesc * curhandler);  int svSysTrapHandlerGetConnected(unsigned int trapnumber, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svSysTrapHandler</i> calls are used to manage system trap handlers. These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>The <i>trapnumber</i> parameter identifies the trap number to which the handler must be connected.</p> <p>Trap handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>When a thread executes a hardware trap instruction, the system invokes (see <i>lapInvoke(2K)</i>) the trap handler connected to the trap number specified by the trap instruction.</p> <p>The argument of the lap trap handler is a pointer to a <i>KnSysTrapDesc</i> data structure. This structure has the following fields:</p> <pre>KnThreadCtx *threadCtx ; unsigned int trapNumber ;</pre> <p>The <i>threadCtx</i> field gives access to the register context of the thread saved immediatly after it performed the hardware trap instruction. This context is processor—dependent and may be modified by the trap handler.</p> <p>The <i>trapNumber</i> field is a processor—dependent trap number.</p> <p>The <i>svSysTrapHandlerConnect</i> system call duplicates the lap descriptor pointed to by <i>newhandler</i> into the lap descriptor associated with <i>trapNumber</i> .</p> <p>The <i>svSysTrapHandlerDisconnect</i> system call clears the lap descriptor associated with <i>trapNumber</i> . If <i>curhandler</i> is not equal to <i>K_CONNECTED_LAP</i>, it must point to a lap descriptor identical to the lap descriptor currently installed.</p> <p>The <i>svSysTrapHandlerGetConnected</i> system call copies the lap descriptor associated with <i>trapNumber</i> at the location pointed to by <i>curhandler</i> ..</p>
<b>RETURN VALUE</b>	Upon successful completion, the calls return <i>K_OK</i> , otherwise a negative error code is returned.

**ERRORS**

[K\_EBUSY]

*svSysTrapHandlerConnect* is called and there is already a valid lap descriptor installed for the trap *trapnumber* .

[K\_EINVAL]

*trapnumber* is invalid. The *curhandler* lap descriptor specified to *svSysTrapHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svLapCreate(2K)` , `lapInvoke(2K)`

<b>NAME</b>	svSysTrapHandler, svSysTrapHandlerConnect, svSysTrapHandlerDisconnect, svSysTrapHandlerGetConnected – Connect a trap handler; Disconnect a trap handler; Get a trap handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTrap.h&gt; int svSysTrapHandlerConnect(unsigned int trapnumber, KnLapDesc * newhandler);  int svSysTrapHandlerDisconnect(unsigned int trapnumber, KnLapDesc * curhandler);  int svSysTrapHandlerGetConnected(unsigned int trapnumber, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svSysTrapHandler</i> calls are used to manage system trap handlers. These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>The <i>trapnumber</i> parameter identifies the trap number to which the handler must be connected.</p> <p>Trap handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>When a thread executes a hardware trap instruction, the system invokes (see <i>lapInvoke(2K)</i>) the trap handler connected to the trap number specified by the trap instruction.</p> <p>The argument of the lap trap handler is a pointer to a <i>KnSysTrapDesc</i> data structure. This structure has the following fields:</p> <pre>KnThreadCtx *threadCtx ; unsigned int trapNumber ;</pre> <p>The <i>threadCtx</i> field gives access to the register context of the thread saved immediatly after it performed the hardware trap instruction. This context is processor—dependent and may be modified by the trap handler.</p> <p>The <i>trapNumber</i> field is a processor—dependent trap number.</p> <p>The <i>svSysTrapHandlerConnect</i> system call duplicates the lap descriptor pointed to by <i>newhandler</i> into the lap descriptor associated with <i>trapNumber</i> .</p> <p>The <i>svSysTrapHandlerDisconnect</i> system call clears the lap descriptor associated with <i>trapNumber</i> . If <i>curhandler</i> is not equal to <i>K_CONNECTED_LAP</i>, it must point to a lap descriptor identical to the lap descriptor currently installed.</p> <p>The <i>svSysTrapHandlerGetConnected</i> system call copies the lap descriptor associated with <i>trapNumber</i> at the location pointed to by <i>curhandler</i> ..</p>
<b>RETURN VALUE</b>	Upon successful completion, the calls return <i>K_OK</i> , otherwise a negative error code is returned.

**ERRORS**

[K\_EBUSY]

*svSysTrapHandlerConnect* is called and there is already a valid lap descriptor installed for the trap *trapnumber* .

[K\_EINVAL]

*trapnumber* is invalid. The *curhandler* lap descriptor specified to *svSysTrapHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svLapCreate(2K)` , `lapInvoke(2K)`

<b>NAME</b>	svSysTrapHandler, svSysTrapHandlerConnect, svSysTrapHandlerDisconnect, svSysTrapHandlerGetConnected – Connect a trap handler; Disconnect a trap handler; Get a trap handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTrap.h&gt; int svSysTrapHandlerConnect(unsigned int trapnumber, KnLapDesc * newhandler);  int svSysTrapHandlerDisconnect(unsigned int trapnumber, KnLapDesc * curhandler);  int svSysTrapHandlerGetConnected(unsigned int trapnumber, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svSysTrapHandler</i> calls are used to manage system trap handlers. These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>The <i>trapnumber</i> parameter identifies the trap number to which the handler must be connected.</p> <p>Trap handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>When a thread executes a hardware trap instruction, the system invokes (see <i>lapInvoke(2K)</i>) the trap handler connected to the trap number specified by the trap instruction.</p> <p>The argument of the lap trap handler is a pointer to a <i>KnSysTrapDesc</i> data structure. This structure has the following fields:</p> <pre>KnThreadCtx *threadCtx ; unsigned int trapNumber ;</pre> <p>The <i>threadCtx</i> field gives access to the register context of the thread saved immediatly after it performed the hardware trap instruction. This context is processor—dependent and may be modified by the trap handler.</p> <p>The <i>trapNumber</i> field is a processor—dependent trap number.</p> <p>The <i>svSysTrapHandlerConnect</i> system call duplicates the lap descriptor pointed to by <i>newhandler</i> into the lap descriptor associated with <i>trapNumber</i> .</p> <p>The <i>svSysTrapHandlerDisconnect</i> system call clears the lap descriptor associated with <i>trapNumber</i> . If <i>curhandler</i> is not equal to <i>K_CONNECTED_LAP</i>, it must point to a lap descriptor identical to the lap descriptor currently installed.</p> <p>The <i>svSysTrapHandlerGetConnected</i> system call copies the lap descriptor associated with <i>trapNumber</i> at the location pointed to by <i>curhandler</i> ..</p>
<b>RETURN VALUE</b>	Upon successful completion, the calls return <i>K_OK</i> , otherwise a negative error code is returned.

**ERRORS**

[K\_EBUSY]

*svSysTrapHandlerConnect* is called and there is already a valid lap descriptor installed for the trap *trapnumber* .

[K\_EINVAL]

*trapnumber* is invalid. The *curhandler* lap descriptor specified to *svSysTrapHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svLapCreate(2K)` , `lapInvoke(2K)`

<b>NAME</b>	svSysTrapHandler, svSysTrapHandlerConnect, svSysTrapHandlerDisconnect, svSysTrapHandlerGetConnected – Connect a trap handler; Disconnect a trap handler; Get a trap handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTrap.h&gt; int svSysTrapHandlerConnect(unsigned int trapnumber, KnLapDesc * newhandler);  int svSysTrapHandlerDisconnect(unsigned int trapnumber, KnLapDesc * curhandler);  int svSysTrapHandlerGetConnected(unsigned int trapnumber, KnLapDesc * curhandler);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svSysTrapHandler</i> calls are used to manage system trap handlers. These calls are restricted to <i>SUPERVISOR</i> threads.</p> <p>The <i>trapnumber</i> parameter identifies the trap number to which the handler must be connected.</p> <p>Trap handlers are specified in the form of lap descriptors (see <i>svLapCreate(2K)</i>).</p> <p>When a thread executes a hardware trap instruction, the system invokes (see <i>lapInvoke(2K)</i>) the trap handler connected to the trap number specified by the trap instruction.</p> <p>The argument of the lap trap handler is a pointer to a <i>KnSysTrapDesc</i> data structure. This structure has the following fields:</p> <pre>KnThreadCtx *threadCtx ; unsigned int trapNumber ;</pre> <p>The <i>threadCtx</i> field gives access to the register context of the thread saved immediatly after it performed the hardware trap instruction. This context is processor—dependent and may be modified by the trap handler.</p> <p>The <i>trapNumber</i> field is a processor—dependent trap number.</p> <p>The <i>svSysTrapHandlerConnect</i> system call duplicates the lap descriptor pointed to by <i>newhandler</i> into the lap descriptor associated with <i>trapNumber</i> .</p> <p>The <i>svSysTrapHandlerDisconnect</i> system call clears the lap descriptor associated with <i>trapNumber</i> . If <i>curhandler</i> is not equal to <i>K_CONNECTED_LAP</i>, it must point to a lap descriptor identical to the lap descriptor currently installed.</p> <p>The <i>svSysTrapHandlerGetConnected</i> system call copies the lap descriptor associated with <i>trapNumber</i> at the location pointed to by <i>curhandler</i> ..</p>
<b>RETURN VALUE</b>	Upon successful completion, the calls return <i>K_OK</i> , otherwise a negative error code is returned.

**ERRORS**

[K\_EBUSY]

*svSysTrapHandlerConnect* is called and there is already a valid lap descriptor installed for the trap *trapnumber* .

[K\_EINVAL]

*trapnumber* is invalid. The *curhandler* lap descriptor specified to *svSysTrapHandlerDisconnect* is not equal to K\_CONNECTED\_LAP and does not match the lap descriptor currently installed.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svLapCreate(2K)` , `lapInvoke(2K)`

<b>NAME</b>	svThreadVirtualTimeout, svThreadVirtualTimeoutSet, svThreadVirtualTimeoutCancel – Set a thread's virtual timeout; Cancel a thread's virtual timeout
<b>SYNOPSIS</b>	<pre>#include &lt;vtimer/chVtimer.h&gt; int svThreadVirtualTimeoutSet(KnCap * actor, KnThreadLid threadLi, KnVirtTimeout * vtimeout, KnTimeVal * cputime, int flag, KnLapDesc * lapdesc);  int svThreadVirtualTimeoutCancel(KnVirtTimeout * vtimeout);</pre>
<b>FEATURES</b>	VTIMER
<b>DESCRIPTION</b>	<p>The <i>svThreadVirtualTimeoutSet</i> system call sets a timeout on the execution time of the thread whose local identifier is <i>threadLi</i> within the actor specified by <i>actorcap</i>.</p> <p>Once the designated thread has consumed <i>cputime</i> of additional execution time, the lap handler designated by <i>lapdesc</i> is invoked with <i>vtimeout</i> as its argument (see <i>lapInvoke(2K)</i>).</p> <p>If <i>flag</i> is set to <code>K_VTIME_INTERNAL</code>, only execution time in the threads' owning actor is counted toward the timeout.</p> <p>If <i>flag</i> is set to <code>K_VTIME_TOTAL</code>, all execution time is counted, regardless of whether or not the thread is executing in a cross-actor invocation.</p> <p>The <i>lapDesc</i> argument is a lap descriptor previously created by <i>svLapCreate</i>.</p> <p>The <i>KnVirtTimeout</i> object is opaque; it must be pre-allocated by the caller, but its fields are inaccessible to the caller. Virtual timeouts are always relative.</p> <p>Virtual timeout handlers are not invoked at interrupt level. The execution environment is similar to that of an abort handler. The handler is executed by the affected thread, and has no restrictions on use of system calls. A thread can be diverted to execute a handler only when executing in its owning actor; during any cross-actor invocation (trap, ipc, or lap call), virtual timeout handler execution is masked.</p> <p>When a thread is deleted, all the timeouts posted for this thread are implicitly cancelled.</p> <p>The <i>svThreadVirtualTimeoutCancel</i> system call attempts to cancel a virtual timeout request. The <i>vtimeout</i> argument contains the address of the <i>KnVirtTimeout</i> object used in the corresponding call to <i>svThreadVirtualTimeoutSet</i>. If the timeout request is still pending, it is cancelled and the call returns <code>K_EOK</code>. If the virtual timeout interval had passed and the handler invocation has been posted, the call takes no action and returns <code>K_ETIMEOUT</code>. In the latter case, no information is available as to whether the handler execution had actually begun, nor whether the handler is still executing. A given timeout must not be</p>

cancelled a second time. This implies as well, that the timeouts of a thread which has been previously deleted, must not be cancelled.

A timeout cannot be reinstalled. To reinstall a timeout request, the timeout must be previously cancelled, otherwise an unpredictable behaviour may occur).

**RETURN VALUES**

Upon successful completion, *svThreadVirtualTimeoutSet* and *svThreadVirtualTimeoutCancel* return K\_OK. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EINVAL] The actor capability is not valid.  
 [K\_EUNKNOWN] *actor* is unreachable.  
 [K\_ETIMEOUT] (*svThreadVirtualTimeoutCancel* only) the virtual interval had already expired, and the handler had been queued for execution.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke(2K)*, *svActorVirtualTimeoutSet(2K)*, *svLapCreate(2K)*, *threadTimes(2K)*, *virtualTimeGetRes(2K)*

<b>NAME</b>	svThreadVirtualTimeout, svThreadVirtualTimeoutSet, svThreadVirtualTimeoutCancel – Set a thread’s virtual timeout; Cancel a thread’s virtual timeout
<b>SYNOPSIS</b>	<pre>#include &lt;vtimer/chVtimer.h&gt; int svThreadVirtualTimeoutSet(KnCap * actor, KnThreadLid threadLi, KnVirtTimeout * vtimeout, KnTimeVal * cputime, int flag, KnLapDesc * lapdesc);  int svThreadVirtualTimeoutCancel(KnVirtTimeout * vtimeout);</pre>
<b>FEATURES</b>	VTIMER
<b>DESCRIPTION</b>	<p>The <i>svThreadVirtualTimeoutSet</i> system call sets a timeout on the execution time of the thread whose local identifier is <i>threadLi</i> within the actor specified by <i>actorcap</i> .</p> <p>Once the designated thread has consumed <i>cputime</i> of additional execution time, the lap handler designated by <i>lapdesc</i> is invoked with <i>vtimeout</i> as its argument (see <i>lapInvoke(2K)</i> ).</p> <p>If <i>flag</i> is set to <code>K_VTIME_INTERNAL</code>, only execution time in the threads’ owning actor is counted toward the timeout.</p> <p>If <i>flag</i> is set to <code>K_VTIME_TOTAL</code>, all execution time is counted, regardless of whether or not the thread is executing in a cross-actor invocation.</p> <p>The <i>lapDesc</i> argument is a lap descriptor previously created by <i>svLapCreate</i> .</p> <p>The <i>KnVirtTimeout</i> object is opaque; it must be pre-allocated by the caller, but its fields are inaccessible to the caller. Virtual timeouts are always relative.</p> <p>Virtual timeout handlers are not invoked at interrupt level. The execution environment is similar to that of an abort handler. The handler is executed by the affected thread, and has no restrictions on use of system calls. A thread can be diverted to execute a handler only when executing in its owning actor; during any cross-actor invocation (trap, ipc, or lap call), virtual timeout handler execution is masked.</p> <p>When a thread is deleted, all the timeouts posted for this thread are implicitly cancelled.</p> <p>The <i>svThreadVirtualTimeoutCancel</i> system call attempts to cancel a virtual timeout request. The <i>vtimeout</i> argument contains the address of the <i>KnVirtTimeout</i> object used in the corresponding call to <i>svThreadVirtualTimeoutSet</i> . If the timeout request is still pending, it is cancelled and the call returns <code>K_EOK</code>. If the virtual timeout interval had passed and the handler invocation has been posted, the call takes no action and returns <code>K_ETIMEOUT</code>. In the latter case, no information is available as to whether the handler execution had actually begun, nor whether the handler is still executing. A given timeout must not be</p>

cancelled a second time. This implies as well, that the timeouts of a thread which has been previously deleted, must not be cancelled.

A timeout cannot be reinstalled. To reinstall a timeout request, the timeout must be previously cancelled, otherwise an unpredictable behaviour may occur).

**RETURN VALUES**

Upon successful completion, *svThreadVirtualTimeoutSet* and *svThreadVirtualTimeoutCancel* return K\_OK. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EINVAL]                      The actor capability is not valid.  
 [K\_EUNKNOWN]                  *actor* is unreachable.  
 [K\_ETIMEOUT]                  (*svThreadVirtualTimeoutCancel* only) the virtual interval had already expired, and the handler had been queued for execution.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke(2K)* , *svActorVirtualTimeoutSet(2K)* , *svLapCreate(2K)* , *threadTimes(2K)* , *virtualTimeGetRes(2K)*

<b>NAME</b>	svThreadVirtualTimeout, svThreadVirtualTimeoutSet, svThreadVirtualTimeoutCancel – Set a thread’s virtual timeout; Cancel a thread’s virtual timeout
<b>SYNOPSIS</b>	<pre>#include &lt;vtimer/chVtimer.h&gt; int svThreadVirtualTimeoutSet(KnCap * actor, KnThreadLid threadLi, KnVirtTimeout * vtimeout, KnTimeVal * cputime, int flag, KnLapDesc * lapdesc);  int svThreadVirtualTimeoutCancel(KnVirtTimeout * vtimeout);</pre>
<b>FEATURES</b>	VTIMER
<b>DESCRIPTION</b>	<p>The <i>svThreadVirtualTimeoutSet</i> system call sets a timeout on the execution time of the thread whose local identifier is <i>threadLi</i> within the actor specified by <i>actorcap</i>.</p> <p>Once the designated thread has consumed <i>cputime</i> of additional execution time, the lap handler designated by <i>lapdesc</i> is invoked with <i>vtimeout</i> as its argument (see <i>lapInvoke(2K)</i>).</p> <p>If <i>flag</i> is set to <code>K_VTIME_INTERNAL</code>, only execution time in the threads’ owning actor is counted toward the timeout.</p> <p>If <i>flag</i> is set to <code>K_VTIME_TOTAL</code>, all execution time is counted, regardless of whether or not the thread is executing in a cross-actor invocation.</p> <p>The <i>lapDesc</i> argument is a lap descriptor previously created by <i>svLapCreate</i>.</p> <p>The <i>KnVirtTimeout</i> object is opaque; it must be pre-allocated by the caller, but its fields are inaccessible to the caller. Virtual timeouts are always relative.</p> <p>Virtual timeout handlers are not invoked at interrupt level. The execution environment is similar to that of an abort handler. The handler is executed by the affected thread, and has no restrictions on use of system calls. A thread can be diverted to execute a handler only when executing in its owning actor; during any cross-actor invocation (trap, ipc, or lap call), virtual timeout handler execution is masked.</p> <p>When a thread is deleted, all the timeouts posted for this thread are implicitly cancelled.</p> <p>The <i>svThreadVirtualTimeoutCancel</i> system call attempts to cancel a virtual timeout request. The <i>vtimeout</i> argument contains the address of the <i>KnVirtTimeout</i> object used in the corresponding call to <i>svThreadVirtualTimeoutSet</i>. If the timeout request is still pending, it is cancelled and the call returns <code>K_EOK</code>. If the virtual timeout interval had passed and the handler invocation has been posted, the call takes no action and returns <code>K_ETIMEOUT</code>. In the latter case, no information is available as to whether the handler execution had actually begun, nor whether the handler is still executing. A given timeout must not be</p>

cancelled a second time. This implies as well, that the timeouts of a thread which has been previously deleted, must not be cancelled.

A timeout cannot be reinstalled. To reinstall a timeout request, the timeout must be previously cancelled, otherwise an unpredictable behaviour may occur).

**RETURN VALUES**

Upon successful completion, *svThreadVirtualTimeoutSet* and *svThreadVirtualTimeoutCancel* return K\_OK. Otherwise, a negative error code is returned.

**ERRORS**

[K\_EINVAL]                   The actor capability is not valid.  
 [K\_EUNKNOWN]                *actor* is unreachable.  
 [K\_ETIMEOUT]                ( *svThreadVirtualTimeoutCancel* only) the virtual interval had already expired, and the handler had been queued for execution.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*lapInvoke(2K)* , *svActorVirtualTimeoutSet(2K)* , *svLapCreate(2K)* , *threadTimes(2K)* , *virtualTimeGetRes(2K)*

<b>NAME</b>	svSysTimeout, svSysTimeoutSet, svSysTimeoutCancel, svTimeoutGetRes – Request a timeout; Cancel a timeout; Get timeout resolution
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTimeout.h&gt; int svSysTimeoutSet(KnTimeout * timeout, KnTimeVal * waitLimit, int flag, KnLapDesc * lapdesc);  int svSysTimeoutCancel(KnTimeout * timeout);  int svTimeoutGetRes(KnTimeVal * resolution);</pre>
<b>FEATURES</b>	TIMEOUT
<b>DESCRIPTION</b>	<p>The <code>svSysTimeoutSet()</code> system call declares <i>lapdesc</i> as the descriptor of a raw lap (see <code>svLapCreate(2K)</code>), which must be invoked when the interval of time specified by <i>*waitLimit</i> (see <code>intro(2K)</code>) has elapsed.</p> <p>The <i>waitLimit</i> pointer argument refers to a <code>KnTimeVal</code> structure whose members are defined in <code>sysTime(2K)</code>. If the <code>KnTimeVal</code> structure indicated by <i>waitLimit</i> is not equal to the <i>resolution</i>, then the time value composed by the <code>tmSec</code> and <code>tmNSec</code> fields is rounded. The fields are rounded down to the nearest multiple of the values in <i>resolution</i> unless the time is zero, in which case they are rounded up to the values in <i>resolution</i>.</p> <p>The lap handler of the <i>lapdesc</i> descriptor is a special kind of interrupt handler which is executed solely in SUPERVISOR execution mode, its code and accessed data must be within the locked-in-memory regions of a SUPERVISOR actor. The set of kernel calls that can be used in this timeout handler is limited, as for any interrupt handler.</p> <p>The argument passed to the lap handler is the <i>timeout</i> parameter specified by <code>svSysTimeoutSet()</code>. Fields within the <code>KnTimeout</code> structure are initialized and modified solely within the nucleus and are inaccessible to the application.</p> <p>The <i>flag</i> argument is reserved for future use and should therefore be set to 0 to ensure compatibility with future releases of the system.</p> <p>The <code>svSysTimeoutCancel()</code> system call attempts to cancel a timeout request. It takes as an argument the address of the <code>KnTimeout</code> object used in the call to <code>svSysTimeoutSet()</code>. If the timeout request is still pending, it is immediately canceled and a value is returned showing that the timer had not yet expired. Otherwise, the timeout interval had expired and the timeout handler was called, in which case nothing occurs and the value returned indicates this.</p> <p>The <code>svTimeoutGetRes()</code> system call obtains the resolution of the <i>*waitLimit</i> argument of <code>svSysTimeoutSet()</code>. The time value returned in <i>resolution</i> represents the smallest possible difference between two distinct <i>*waitLimit</i> values.</p>

The `svSysTimeoutSet()`, `svTimeoutGetRes()` and `svSysTimeoutCancel()` system calls are restricted to SUPERVISOR threads.

The `svSysTimeoutSet()` system call must not be called using a `KnTimeout` which has an existing timeout pending (use `svSysTimeoutCancel()` to cancel the request first).

**RETURN VALUES**

Upon completion, `svSysTimeoutSet()` returns `K_OK` if the timeout was successfully entered. Otherwise, a negative error code is returned specifying the reason it was not successful. Upon completion, `svSysTimeoutCancel()` returns a boolean value set to 1 if the request was still pending, or 0 if the timeout handler was called.

**ERRORS**

[K\_EINVAL] The *\*waitLimit* value specified is invalid, or *lapdesc* is not a valid raw lap descriptor.

The value indicated by *\*waitLimit* is invalid if either the `tmSec` or `tmNSec` fields of the structure are negative, or if the structure is poorly constructed (for example, if the number of nanoseconds in `tmNSec` makes `tmNSec` greater than one second).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svIntrConnect(2K)`, `svLapCreate(2K)`, `sysTime(2K)`

<b>NAME</b>	svTrapConnect, svTrapDisconnect – Connect a trap handler; Disconnect a trap handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTrap.h&gt; int svTrapConnect(unsigned int trapno, KnTrapHdl routine);  int svTrapDisconnect(unsigned int trapno);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svTrapConnect</i> system call connects a trap handler to the hardware trap specified by <i>trapno</i>. The <i>trapno</i> value is machine-dependent. The <i>routine</i> pointer points to the handler. The handler is a function which takes two arguments:</p> <pre>int handler (ctx, itno)      KnThreadCtx* ctx ;     unsigned int itno ;</pre> <p>After being connected, a handler will be activated upon invocation of the trap <i>trapno</i> by a thread. The current thread register context is saved. The handler takes as arguments a pointer to the saved context ( <i>ctx</i> ) and the trap number ( <i>trapno</i> ).</p> <p>Only one handler may be connected to a given trap number at a time. The <i>svTrapDisconnect</i> system call disconnects the handler <i>routine</i> previously associated with the <i>trapno</i> trap number. These calls are restricted to <i>SUPERVISOR</i> threads. Trap handlers are executed in <i>SUPERVISOR</i> execution mode. The handler's code and data accessed must be parts of locked-in-memory regions of a <i>SUPERVISOR</i> actor.</p> <p>When a thread executes a trap handler (via a hardware trap invocation), its execution actor is changed to match the actor which declared the handler. It is still attached to the actor within which it was created, although the executed code is part of another actor's address space (a <i>SUPERVISOR</i> actor).</p> <p>The execution actor of a thread determines the actor which must be used for the invocation of exception handlers (see <i>svExcHandler</i> (2K)) if the thread encounters an exception. It also determines when abort handlers are invoked (see <i>svAbortHandler</i> (2K)).</p> <p>The execution actor does not determine the interpretation of actor references, the thread's home actor is always used. Consequently, care must be taken when invoking <i>actorSelf</i>(2K) or using the <i>K_MYACTOR</i> value in actor—related calls (as in <i>threadCreate</i> (2K) for example) within a trap handler; they refer to the thread's home actor. The execution actor must be explicitly named by its capability, if needed (for example if the handler is to manage this actor's address space).</p> <p>When the trap handler returns, the previous thread's execution actor is restored.</p>
<b>ERRORS</b>	<p>[K_EBUSY]                      There is already a valid descriptor installed for the trap number.</p>

[K\_EINVAL] Incorrect trap number.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svExcHandler(2K)`

<b>NAME</b>	svTrapConnect, svTrapDisConnect – Connect a trap handler; Disconnect a trap handler
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTrap.h&gt; int svTrapConnect(unsigned int trapno, KnTrapHdl routine);  int svTrapDisConnect(unsigned int trapno);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>svTrapConnect</i> system call connects a trap handler to the hardware trap specified by <i>trapno</i>. The <i>trapno</i> value is machine-dependent. The <i>routine</i> pointer points to the handler. The handler is a function which takes two arguments:</p> <pre>int handler (ctx, itno)      KnThreadCtx* ctx ;     unsigned int itno ;</pre> <p>After being connected, a handler will be activated upon invocation of the trap <i>trapno</i> by a thread. The current thread register context is saved. The handler takes as arguments a pointer to the saved context (<i>ctx</i>) and the trap number (<i>trapno</i>).</p> <p>Only one handler may be connected to a given trap number at a time. The <i>svTrapDisConnect</i> system call disconnects the handler <i>routine</i> previously associated with the <i>trapno</i> trap number. These calls are restricted to <i>SUPERVISOR</i> threads. Trap handlers are executed in <i>SUPERVISOR</i> execution mode. The handler's code and data accessed must be parts of locked-in-memory regions of a <i>SUPERVISOR</i> actor.</p> <p>When a thread executes a trap handler (via a hardware trap invocation), its execution actor is changed to match the actor which declared the handler. It is still attached to the actor within which it was created, although the executed code is part of another actor's address space (a <i>SUPERVISOR</i> actor).</p> <p>The execution actor of a thread determines the actor which must be used for the invocation of exception handlers (see <i>svExcHandler</i> (2K)) if the thread encounters an exception. It also determines when abort handlers are invoked (see <i>svAbortHandler</i> (2K)).</p> <p>The execution actor does not determine the interpretation of actor references, the thread's home actor is always used. Consequently, care must be taken when invoking <i>actorSelf</i>(2K) or using the <i>K_MYACTOR</i> value in actor—related calls (as in <i>threadCreate</i> (2K) for example) within a trap handler; they refer to the thread's home actor. The execution actor must be explicitly named by its capability, if needed (for example if the handler is to manage this actor's address space).</p> <p>When the trap handler returns, the previous thread's execution actor is restored.</p>
<b>ERRORS</b>	<p>[K_EBUSY]                      There is already a valid descriptor installed for the trap number.</p>

[K\_EINVAL] Incorrect trap number.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`svExceptionHandler(2K)`

<b>NAME</b>	svMaskAll, svUnmaskAll, svUnmask – Disable interrupt processing; Enable interrupt processing; Reenable interrupt processing				
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSpinLock.h&gt; int svMaskAll(void);  void svUnmaskAll(void);  void svUnmask(int oldval);</pre>				
<b>DESCRIPTION</b>	<p>The <i>svMaskAll</i> function disables interrupts on the current processor by setting the processor interrupt mask to its maximum value. There is no effect on other processors. It returns the previous interrupt mask.</p> <p>The <i>svUnmaskAll</i> function enables interrupts on the current processor by setting the processor interrupt mask to its minimum value. There is no effect on other processors.</p> <p>The <i>svUnmask</i> function resets the interrupt mask to a state represented by <i>oldval</i>. The <i>oldval</i> value must have been returned by a previous invocation of <i>svMaskAll</i>. Pairs of <i>svMaskAll</i> and <i>svUnmask</i> may be nested. The interrupt mask is part of a thread's execution context, it is saved as part of the thread's status when the thread is suspended. If a thread performs a blocking call after calling <i>svMaskAll</i>, interrupt processing may be enabled by another scheduled thread. It will be disabled again when the thread is restarted. Such situations should be avoided.</p> <p>The <i>svMaskAll</i>, <i>svUnmaskAll</i> and <i>svUnmask</i> functions only relate to the interrupt mask from processor standpoint. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.</p> <p>These calls are restricted to <i>SUPERVISOR</i> threads. Their use must be limited to a very small number of instructions. They should not be used to ensure mutual exclusion. The routines <i>svMaskedLockGet</i> and <i>svMaskedLockRel</i> should be used for this purpose.</p>				
<b>RETURN VALUE</b>	<i>svMaskAll</i> returns the current interrupt masking status.				
<b>ERRORS</b>	No error messages are returned .				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>svMaskedLockInit(2K)</code>				

<b>NAME</b>	svMaskAll, svUnmaskAll, svUnmask – Disable interrupt processing; Enable interrupt processing; Reenable interrupt processing				
<b>SYNOPSIS</b>	<pre>#include &lt;sync/chSpinLock.h&gt; int svMaskAll(void);  void svUnmaskAll(void);  void svUnmask(int oldval);</pre>				
<b>DESCRIPTION</b>	<p>The <i>svMaskAll</i> function disables interrupts on the current processor by setting the processor interrupt mask to its maximum value. There is no effect on other processors. It returns the previous interrupt mask.</p> <p>The <i>svUnmaskAll</i> function enables interrupts on the current processor by setting the processor interrupt mask to its minimum value. There is no effect on other processors.</p> <p>The <i>svUnmask</i> function resets the interrupt mask to a state represented by <i>oldval</i>. The <i>oldval</i> value must have been returned by a previous invocation of <i>svMaskAll</i>. Pairs of <i>svMaskAll</i> and <i>svUnmask</i> may be nested. The interrupt mask is part of a thread's execution context, it is saved as part of the thread's status when the thread is suspended. If a thread performs a blocking call after calling <i>svMaskAll</i>, interrupt processing may be enabled by another scheduled thread. It will be disabled again when the thread is restarted. Such situations should be avoided.</p> <p>The <i>svMaskAll</i>, <i>svUnmaskAll</i> and <i>svUnmask</i> functions only relate to the interrupt mask from processor standpoint. If the machine is equipped with an external interrupt controller, these calls have no effect on the controller status.</p> <p>These calls are restricted to <i>SUPERVISOR</i> threads. Their use must be limited to a very small number of instructions. They should not be used to ensure mutual exclusion. The routines <i>svMaskedLockGet</i> and <i>svMaskedLockRel</i> should be used for this purpose.</p>				
<b>RETURN VALUE</b>	<i>svMaskAll</i> returns the current interrupt masking status.				
<b>ERRORS</b>	No error messages are returned .				
<b>ATTRIBUTES</b>	See <a href="#">attributes(5)</a> 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;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<a href="#">svMaskedLockInit(2K)</a>				

<b>NAME</b>	svVirtualTimeoutSet, svVirtualTimeoutCancel – Set a virtual timeout; Cancel a virtual timeout
<b>SYNOPSIS</b>	<pre>#include &lt;vtimer/chVtimer.h&gt; int svVirtualTimeoutSet(KnCap * actor, KnThreadLid threadLi, KnVirtTimeout * vtimeout, KnVirtToHdl handler, KnTimeVal * cputime, int flag);  int svVirtualTimeoutCancel(KnVirtTimeout * vtimeout);</pre>
<b>FEATURES</b>	VTIMER
<b>DESCRIPTION</b>	<p>The <i>svVirtualTimeoutSet</i> system call sets a timeout on the execution time of the thread whose local identifier is <i>threadLi</i> within the actor specified by <i>actorcap</i>.</p> <p>If <i>threadLi</i> is set to <code>K_ALLACTORTHREADS</code>, the timeout applies to the total execution time of all threads in the actor designated by <i>actorcap</i>. Once the designated thread or threads have consumed <i>cputime</i> of additional execution time, the handler designated by <i>handler</i> is entered with <i>vtimeout</i> as its sole argument. (Virtual timeouts are always relative.)</p> <p>If <i>flag</i> is set to <code>K_VTIME_INTERNAL</code>, only execution time in the threads' home actor is counted toward the timeout.</p> <p>If <i>flag</i> is set to <code>K_VTIME_TOTAL</code>, all execution time is counted, regardless of whether or not the thread is executing in a cross-actor invocation.</p> <p>In the <code>K_ALLACTORTHREADS</code> case, thread execution time is charged to the thread's owning actor only, not to the execution actor.</p> <p>The <i>KnVirtTimeout</i> object is opaque; it must be pre-allocated by the caller, but its fields are inaccessible to the caller.</p> <p>Virtual timeout handlers are not invoked at interrupt level. The execution environment is similar to that of an abort handler. The handler is executed by the affected thread (or the next affected thread to run, if an actor-level timeout), and has no restrictions on use of system calls. A thread can be diverted to execute a handler only when executing in its owning actor; during any cross-actor invocation (trap, ipc, or lap call), handler execution is masked.</p> <p>The <i>svVirtualTimeoutCancel</i> system call attempts to cancel a virtual timeout request. The <i>vtimeout</i> argument contains the address of the <i>KnVirtTimeout</i> object used in the corresponding call to <i>svVirtTimeoutSet</i>. If the timeout request is still pending, it is cancelled and the call returns <code>K_EOK</code>. If the virtual timeout interval expired and the handler invocation has been posted, the call takes no action and returns <code>K_ETIMEOUT</code>. In the latter case, no information is available as to whether the handler execution has actually begun, nor whether the handler is still executing.</p>

A timeout cannot be reinstalled. To reinstall a timeout request, the timeout must be previously cancelled, otherwise an unpredictable behaviour may occur.

**RETURN VALUES**

Upon successful completion, *svVirtualTimeoutSet* and *svTimeoutCancel* return `K_OK`. Otherwise, a negative error code is returned.

**ERRORS**

[`K_EINVAL`] *threadli* does not designate a valid thread, or the time value referenced by *cputime* is invalid, or *flag* contains an invalid value, or actor capability is not valid.

[`K_ENOMEM`] The system is out of resources.

[`K_EUNKNOWN`] *actor* is unreachable.

[`K_ETIMEOUT`] (*svVirtualTimeoutCancel* only) the virtual interval had already expired, and the handler had been queued for execution.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`threadTimes(2K)`, `virtualTimeGetRes(2K)`

<b>NAME</b>	svVirtualTimeoutSet, svVirtualTimeoutCancel – Set a virtual timeout; Cancel a virtual timeout
<b>SYNOPSIS</b>	<pre>#include &lt;vtimer/chVtimer.h&gt; int svVirtualTimeoutSet(KnCap * actor, KnThreadLid threadLi, KnVirtTimeout * vtimeout, KnVirtToHdl handler, KnTimeVal * cputime, int flag);  int svVirtualTimeoutCancel(KnVirtTimeout * vtimeout);</pre>
<b>FEATURES</b>	VTIMER
<b>DESCRIPTION</b>	<p>The <i>svVirtualTimeoutSet</i> system call sets a timeout on the execution time of the thread whose local identifier is <i>threadLi</i> within the actor specified by <i>actorcap</i>.</p> <p>If <i>threadLi</i> is set to <i>K_ALLACTORTHREADS</i>, the timeout applies to the total execution time of all threads in the actor designated by <i>actorcap</i>. Once the designated thread or threads have consumed <i>cputime</i> of additional execution time, the handler designated by <i>handler</i> is entered with <i>vtimeout</i> as its sole argument. (Virtual timeouts are always relative.)</p> <p>If <i>flag</i> is set to <i>K_VTIME_INTERNAL</i>, only execution time in the threads' home actor is counted toward the timeout.</p> <p>If <i>flag</i> is set to <i>K_VTIME_TOTAL</i>, all execution time is counted, regardless of whether or not the thread is executing in a cross-actor invocation.</p> <p>In the <i>K_ALLACTORTHREADS</i> case, thread execution time is charged to the thread's owning actor only, not to the execution actor.</p> <p>The <i>KnVirtTimeout</i> object is opaque; it must be pre-allocated by the caller, but its fields are inaccessible to the caller.</p> <p>Virtual timeout handlers are not invoked at interrupt level. The execution environment is similar to that of an abort handler. The handler is executed by the affected thread (or the next affected thread to run, if an actor-level timeout), and has no restrictions on use of system calls. A thread can be diverted to execute a handler only when executing in its owning actor; during any cross-actor invocation (trap, ipc, or lap call), handler execution is masked.</p> <p>The <i>svVirtualTimeoutCancel</i> system call attempts to cancel a virtual timeout request. The <i>vtimeout</i> argument contains the address of the <i>KnVirtTimeout</i> object used in the corresponding call to <i>svVirtTimeoutSet</i>. If the timeout request is still pending, it is cancelled and the call returns <i>K_EOK</i>. If the virtual timeout interval expired and the handler invocation has been posted, the call takes no action and returns <i>K_ETIMEOUT</i>. In the latter case, no information is available as to whether the handler execution has actually begun, nor whether the handler is still executing.</p>

A timeout cannot be reinstalled. To reinstall a timeout request, the timeout must be previously cancelled, otherwise an unpredictable behaviour may occur.

**RETURN VALUES**

Upon successful completion, *svVirtualTimeoutSet* and *svTimeoutCancel* return `K_OK`. Otherwise, a negative error code is returned.

**ERRORS**

[`K_EINVAL`] *threadli* does not designate a valid thread, or the time value referenced by *cputime* is invalid, or *flag* contains an invalid value, or actor capability is not valid.

[`K_ENOMEM`] The system is out of resources.

[`K_EUNKNOWN`] *actor* is unreachable.

[`K_ETIMEOUT`] (*svVirtualTimeoutCancel* only) the virtual interval had already expired, and the handler had been queued for execution.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`threadTimes(2K)`, `virtualTimeGetRes(2K)`

<b>NAME</b>	sysBench – kernel benchmark utility
<b>SYNOPSIS</b>	<pre>#include &lt;timer/chBench.h&gt; int sysBench(int option, int *uTime);</pre>
<b>FEATURES</b>	PERF
<b>DESCRIPTION</b>	<p>The <i>sysBench</i> system call returns a microsecond time value suitable for benchmarking. The <i>option</i> argument specifies whether the benchmark time value is for the beginning or end point of the benchmark interval. The argument <i>uTime</i> is the address of an integer which <i>sysBench</i> uses to store a time value in microseconds.</p> <p>It should be noted that the time value returned by <i>sysBench</i> does not necessarily correspond to any actual time, it could simply be an instantaneous value from a clock chip. On some platforms, time values can "wrap-around" after a certain amount of time, causing a subsequent bench point to return a lower time value, making the benchmark interval invalid. For example, on x86 architectures the maximum interval measurable using <i>sysBench</i> is 5 milliseconds.</p> <p>Three options are currently supported:</p> <p><b>K_BSTART</b>      This option starts a benchmark interval by enabling benchmarking, disabling interrupts, and returns an instantaneous microsecond time value in <i>uTime</i>.</p> <p><b>K_BPOINT</b>      This option returns an instantaneous microsecond time value in <i>uTime</i>. This option should be called between a <b>K_BSTART</b> and a <b>K_BSTOP</b> call.</p> <p><b>K_BSTOP</b>        This disables benchmarking, re-enables interrupts appropriately, and returns another instantaneous microsecond time value in <i>uTime</i>.</p>
<b>RETURN VALUE</b>	Upon successful completion <b>K_OK</b> is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	<b>[K_EINVAL]</b> An invalid option was passed as an argument, or the benchmark commands were not issued in the correct order (for example, <b>K_BSTART</b> was not issued before <b>B_STOP</b> or two <b>K_BSTART</b> commands were issued one after another).
<b>EXAMPLES</b>	<p>This small code fragment illustrates a typical usage of <i>sysBench</i>.</p> <pre>int startTime, stopTime, intervalTime; if ((res = sysBench(K_BSTART, &amp;startTime)) != K_OK) {     /* A benchmark was pending */     return res; }</pre>

```
/* ... code to be measured ... */
if ((res = sysBench(K_BSTOP, &stopTime)) != K_OK) {
/* K_BSTOP was already issued */
return res;
}
intervalTime = stopTime - startTime;
```

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**NAME** sysGetConf – Get Chorus module configuration value

**SYNOPSIS** `#include <exec/chModules.h>`  
`int sysGetConf(const char *modName, int var, int *pvalue);`

**FEATURES** CORE

**DESCRIPTION** The *sysGetConf* system call copies the value from the Chorus module corresponding to the module name *modName*, as defined in *chModules.h* into the integer pointed to by *pvalue*. The *modName* pointer must point to a null terminated string containing the name of the module whose configuration value is required. The *pvalue* pointer points to an integer which, on successful completion, will contain the value corresponding to the *var* configuration parameter of *modName*.

The *var* value may be either *K\_GETCONF\_VERSION* to get the version of a specified module, or a value defined in a specific module header (for example, *K\_GETCONF\_ACTOR\_MAX* in *exec/chExec.h* for the maximum number of actors on a site, or *K\_GETCONF\_PORT\_MAX* in *ipc/chIpc.h* for the maximum number of ports on a site).

**RETURN VALUE** Upon successful completion, if the module specified is present, *K\_OK* is returned. Otherwise, a negative error code is returned.

**ERRORS** `[K_ENOTAVAILABLE]` *modName*; the module is not present in the current Chorus run—time configuration. This is an easy way of testing for the presence of a module at run—time.

`[K_EFAULT]` Some of the data provided are outside the current actor’s address space.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	sysGetEnv – Get a value from the Chorus configuration environment				
<b>SYNOPSIS</b>	#include <env/chEnv.h> int <b>sysGetEnv</b> (const char *envName, char *envValue, unsigned int *size);				
<b>FEATURES</b>	ENV				
<b>DESCRIPTION</b>	<p>The <i>sysGetEnv</i> system call copies the value from the Chorus configuration environment defined by the variable <i>envName</i>, into the buffer pointed to by <i>envValue</i>. The <i>envName</i> configuration variable must point to a null terminated string containing the name of the configuration variable whose value is required. The <i>envValue</i> pointer points to a buffer which, on successful completion, will contain the value corresponding to <i>envName</i>.</p> <p>On input, <i>size</i> points to the size in bytes of the buffer <i>envName</i>. The <i>sysGetEnv</i> call uses this value to determine whether there is sufficient space to copy the value of the variable.</p> <p>On output, <i>size</i> points to the number of bytes in the value, including the terminating null character.</p>				
<b>RETURN VALUE</b>	Upon successful completion K_OK is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EFAIL]                    <i>envName</i> the configuration variable does not exist in the current Chorus environment.</p> <p>[K_ENOMEM]                 For both user and supervisor access, the value pointed to by <i>size</i> was insufficient to store the entire value of <i>envName</i> in <i>envValue</i>. In this case, <i>size</i> points to the number of bytes necessary to contain the entire value. For user mode access, the string size of either <i>envName</i> or <i>envValue</i> exceeded the internal size limits imposed by the user system call implementation. The default values of these limits are 32 bytes for <i>envName</i>, and 256 bytes for <i>envValue</i>.</p> <p>[K_EFAULT]                 Some of the data provided are outside the current actor's address space.</p>				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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				
<b>SEE ALSO</b>	<i>sysSetEnv(2K)</i> , <i>sysUnsetEnv(2K)</i>				

<b>NAME</b>	sysLog – log a message in the kernel’s cyclical buffer				
<b>SYNOPSIS</b>	<pre>#include &lt;log/chLog.h&gt; void sysLog(char *format, ... /* arg */);</pre>				
<b>FEATURES</b>	LOG				
<b>DESCRIPTION</b>	<p>The <code>sysLog()</code> call logs a message in the kernel’s cyclical buffer. The syntax of <code>sysLog()</code> is similar to that of <code>printf()</code> with the restriction that the only conversion specifications supported are “%s”, “%d”, and “%c”.</p> <p>The formatted string is truncated to <code>SYSLOG_MAX_LINE</code> characters, as defined in <code>&lt;log/chLog.h&gt;</code>.</p>				
<b>ATTRIBUTES</b>	<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;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>NOTES</b>	<p>Important actors related to system administration, for example <code>pppstart.r</code>, <code>slattach.r</code> and <code>chat.r</code>, use <code>sysLog()</code>.</p> <p>The system administrator can view lines in the system log using <code>cs(1CC)</code> and <code>chls(1CC)</code>, as shown in the following examples:</p> <pre>host% chls -ll nlines host% rsh target arun /bin/cs -ll nlines</pre> <p>where <code>nlines</code> is the number of lines to display counting from the end of the system log, and <code>target</code> is the target system hostname.</p>				
<b>SEE ALSO</b>	<code>cs(1CC)</code> , <code>chls(1CC)</code>				

<b>NAME</b>	sysRead, sysWrite, sysPoll – Read characters from the system console; Write characters to the system console; Poll characters from the system console
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chIo.h&gt; int sysRead(void * buf, unsigned int nchar);  int sysWrite(void * buf, unsigned int nchar);  int sysPoll(void * buf);</pre>
<b>DESCRIPTION</b>	<p>The <i>sysRead</i> system call reads characters from the system console into the buffer pointed to by <i>buf</i>, until either <i>nchar</i> -1 characters are read, or a newline character is read and not transferred to <i>buf</i>. The characters will be echoed on the system console and the string is terminated with a null character. If <i>nchar</i> is equal to 1, one character is read, without echo.</p> <p>The <i>sysWrite</i> system call attempts to write <i>nchar</i> characters to the system console from the buffer pointed to by <i>buf</i>.</p> <p>The <i>sysPoll</i> system call attempts to read one character from the system console into the buffer pointed to by <i>buf</i>. The character will be echoed on the system console. The <i>sysPoll</i> system call is non-blocking.</p>
<b>RETURN VALUE</b>	If successful, <i>sysRead</i> and <i>sysPoll</i> return the number of characters actually read, and <i>sysWrite</i> returns the number of characters actually written. Otherwise, a negative error code is returned.
<b>ERRORS</b>	[K_EFAULT]                      Some of the arguments provided are outside the caller's address space.
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**NAME** sysRead, sysWrite, sysPoll – Read characters from the system console; Write characters to the system console; Poll characters from the system console

**SYNOPSIS**

```
#include <exec/chIo.h>
int sysRead(void * buf, unsigned int nchar);

int sysWrite(void * buf, unsigned int nchar);

int sysPoll(void * buf);
```

**DESCRIPTION**

The *sysRead* system call reads characters from the system console into the buffer pointed to by *buf*, until either *nchar* - 1 characters are read, or a newline character is read and not transferred to *buf*. The characters will be echoed on the system console and the string is terminated with a null character. If *nchar* is equal to 1, one character is read, without echo.

The *sysWrite* system call attempts to write *nchar* characters to the system console from the buffer pointed to by *buf*.

The *sysPoll* system call attempts to read one character from the system console into the buffer pointed to by *buf*. The character will be echoed on the system console. The *sysPoll* system call is non-blocking.

**RETURN VALUE** If successful, *sysRead* and *sysPoll* return the number of characters actually read, and *sysWrite* returns the number of characters actually written. Otherwise, a negative error code is returned.

**ERRORS** [K\_EFAULT] Some of the arguments provided are outside the caller's address space.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	sysReboot – request a reboot of the local site						
<b>SYNOPSIS</b>	<pre>#include &lt;restart/chRestart.h&gt; int sysReboot(KnRebootDesc*rebootdesc);</pre>						
<b>FEATURES</b>	CORE						
<b>DESCRIPTION</b>	<p>The <i>sysReboot</i> system call requests a reboot of the local site.</p> <p>The caller thread must be a supervisor thread (see <i>threadCreate(2K)</i>) or must belong to a system actor (see <i>actorCreate(2K)</i>).</p> <p>The <i>rebootdesc</i> pointer must be either 0 or a pointer to a <i>KnRebootDesc</i> structure:</p> <pre>KnRebootMode mode ;</pre> <p>The <i>mode</i> field specifies the reboot request and can be one of the following values:</p> <table border="0"> <tr> <td style="vertical-align: top;">K_REBOOT_COLD</td> <td>stops all kernel drivers and initiates a BSP-dependent cold reboot procedure. Note that <i>sysReboot</i> doesn't perform any shutdown operation on the POSIX personality.</td> </tr> <tr> <td style="vertical-align: top;">K_REBOOT_HOT</td> <td>is not available for applications and must be used by the operating system only</td> </tr> <tr> <td style="vertical-align: top;">K_REBOOT_NEW</td> <td>is not available for applications and must be used by the operating system only</td> </tr> </table> <p><i>sysReboot</i> called with the argument 0 performs the above described K_REBOOT_COLD reboot request.</p>	K_REBOOT_COLD	stops all kernel drivers and initiates a BSP-dependent cold reboot procedure. Note that <i>sysReboot</i> doesn't perform any shutdown operation on the POSIX personality.	K_REBOOT_HOT	is not available for applications and must be used by the operating system only	K_REBOOT_NEW	is not available for applications and must be used by the operating system only
K_REBOOT_COLD	stops all kernel drivers and initiates a BSP-dependent cold reboot procedure. Note that <i>sysReboot</i> doesn't perform any shutdown operation on the POSIX personality.						
K_REBOOT_HOT	is not available for applications and must be used by the operating system only						
K_REBOOT_NEW	is not available for applications and must be used by the operating system only						
<b>RETURN VALUE</b>	If successful, <i>sysReboot</i> does not return, otherwise a negative error code is returned.						
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EPRIV]</td> <td>The caller thread is not a SUPERVISOR thread or does not belong to a SYSTEM actor.</td> </tr> <tr> <td style="vertical-align: top;">[K_ENOTIMP]</td> <td>The BSP doesn't implement cold reboot procedure.</td> </tr> </table>	[K_EPRIV]	The caller thread is not a SUPERVISOR thread or does not belong to a SYSTEM actor.	[K_ENOTIMP]	The BSP doesn't implement cold reboot procedure.		
[K_EPRIV]	The caller thread is not a SUPERVISOR thread or does not belong to a SYSTEM actor.						
[K_ENOTIMP]	The BSP doesn't implement cold reboot procedure.						
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
<b>SEE ALSO</b>	<i>sysShutdown(2K)</i>						

<b>NAME</b>	sysSetEnv – Set a value in the ChorusOS configuration environment				
<b>SYNOPSIS</b>	<pre>#include &lt;env/chEnv.h&gt; int <b>sysSetEnv</b>(const char *envName, const char *envValue);</pre>				
<b>DESCRIPTION</b>	<p>The <code>sysSetEnv( )</code> system call inserts an environment variable and its value into the the ChorusOS configuration environment. The <code>envName</code> pointer points to a null terminated string containing the environment variable to be set. The <code>envValue</code> pointer points to a null-terminated string containing the corresponding value of the environment variable.</p> <p>Any existing value for <code>envName</code> in the configuration environment is removed and replaced by <code>envValue</code>.</p>				
<b>RETURN VALUES</b>	Upon successful completion <code>K_OK</code> is returned and the environment variable name and value are inserted into the configuration environment. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EINVAL]     <code>envName</code> was zero length, or an invalide name for a ChorusOS configuration environment variable.</p> <p>[K_ENOMEM]     For both user and supervisor access, there was insufficient space in the ChorusOS configuration environment to store <code>envName</code> and <code>envValue</code>. For user mode access, the string size of either <code>envName</code> or <code>envValue</code> exceeded the internal size limits imposed by the user system call implementation. The default values of these limits are 32 bytes for <code>envName</code>, and 256 bytes for <code>envValue</code>.</p> <p>[K_EFAULT]     Some of the data provided are outside the current actor's address space.</p>				
<b>ATTRIBUTES</b>	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;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>sysGetEnv(2K)</code> , <code>sysUnsetEnv(2K)</code>				

<b>NAME</b>	sysShutdown – shut down or restart the system				
<b>SYNOPSIS</b>	int <b>sysShutdown</b> (int <i>argc</i> , char ** <i>argv</i> );				
<b>DESCRIPTION</b>	<b>sysShutdown</b> ( ) is executed by trusted user actors or supervisor actors to shut down the OS. It is a high-level interface intended for use in multi-command actors such as <b>C_INIT</b> (1M).				
<b>PARAMETERS</b>	<p><b>sysShutdown</b>( ) parameters are similar to the options for <b>shutdown</b>(1M). If the first parameter is <b>-i</b>, then the subsequent digit indicates the system state and actions to perform as follows:</p> <table border="0"> <tr> <td style="padding-right: 20px;">0</td> <td>Shut down the entire system safely and reboot.</td> </tr> <tr> <td>1</td> <td>Perform a site restart. A site restart means that the entire system is shut down safely and then restarted from the system image, without accessing stable storage, such that direct hot restartable actors are automatically restarted from persistent memory. This state requires the <b>HOT_RESTART</b> feature.</td> </tr> </table> <p>See the manual entry for <b>restart</b>(1M) for a more detailed description of site restart.</p>	0	Shut down the entire system safely and reboot.	1	Perform a site restart. A site restart means that the entire system is shut down safely and then restarted from the system image, without accessing stable storage, such that direct hot restartable actors are automatically restarted from persistent memory. This state requires the <b>HOT_RESTART</b> feature.
0	Shut down the entire system safely and reboot.				
1	Perform a site restart. A site restart means that the entire system is shut down safely and then restarted from the system image, without accessing stable storage, such that direct hot restartable actors are automatically restarted from persistent memory. This state requires the <b>HOT_RESTART</b> feature.				
<b>RETURN VALUES</b>	Upon successful completion, 0 is returned and the system is no longer available. Otherwise, a negative error code is returned if a nucleus error is encountered, or a positive error code is returned if an OS error is encountered.				
<b>ERRORS</b>	<p>Nucleus errors include any errors reported by <b>threadCreate</b>(2K) or <b>threadScheduler</b>(2K).</p> <table border="0"> <tr> <td style="padding-right: 20px;">[EFAULT]</td> <td>Actor capability of the invoker cannot be obtained.</td> </tr> <tr> <td>[EINVAL]</td> <td>Invalid argument.</td> </tr> </table>	[EFAULT]	Actor capability of the invoker cannot be obtained.	[EINVAL]	Invalid argument.
[EFAULT]	Actor capability of the invoker cannot be obtained.				
[EINVAL]	Invalid argument.				
<b>ATTRIBUTES</b>	See <b>attributes</b> (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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<b>shutdown</b> (1M), <b>restart</b> (1M)				

**NAME** sysTime, sysTimeGetRes – get system time; get system time resolution

**SYNOPSIS**

```
#include <exec/chTime.h>
int sysTime(KnTimeVal * time);

int sysTimeGetRes(KnTimeVal * resolution);
```

**FEATURES** SYSTIME

**DESCRIPTION** The kernel maintains its time using a *KnTimeVal* structure, the members of which are the following:

```
long tmSec ; /* seconds */
long tmNSec ; /* nanoseconds */
```

The *tmSec* field represents the number of seconds, and the *tmNSec* field represents the number of additional nanoseconds.

This time is set to 0 when the kernel is initialized, and incremented by the timer period at each timer interrupt.

*sysTime* allows the user to get the value of the kernel time. *sysTime* returns the kernel time in the *KnTimeVal* structure pointed to by *time*.

*sysTimeGetRes* obtains the resolution of the *sysTime* system call. The time value returned in *resolution* represents the smallest possible difference between two distinct values of the system time.

**RETURN VALUE** Upon successful completion K\_OK is returned. Otherwise, a negative error code is returned.

**ERRORS** [K\_EFAULT] Some of the data provided are outside the current actor's address space.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *univTime(2K)*

<b>NAME</b>	sysTime, sysTimeGetRes – get system time; get system time resolution				
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chTime.h&gt; int sysTime(KnTimeVal * time);  int sysTimeGetRes(KnTimeVal * resolution);</pre>				
<b>FEATURES</b>	SYSTIME				
<b>DESCRIPTION</b>	<p>The kernel maintains its time using a <i>KnTimeVal</i> structure, the members of which are the following:</p> <pre>long tmSec ; /* seconds */ long tmNSec ; /* nanoseconds */</pre> <p>The <i>tmSec</i> field represents the number of seconds, and the <i>tmNSec</i> field represents the number of additional nanoseconds.</p> <p>This time is set to 0 when the kernel is initialized, and incremented by the timer period at each timer interrupt.</p> <p><i>sysTime</i> allows the user to get the value of the kernel time. <i>sysTime</i> returns the kernel time in the <i>KnTimeVal</i> structure pointed to by <i>time</i>.</p> <p><i>sysTimeGetRes</i> obtains the resolution of the <i>sysTime</i> system call. The time value returned in <i>resolution</i> represents the smallest possible difference between two distinct values of the system time.</p>				
<b>RETURN VALUE</b>	Upon successful completion K_OK is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	[K_EFAULT]                      Some of the data provided are outside the current actor's address space.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	univTime(2K)				

<b>NAME</b>	sysTimer, sysTimerStartPeriodic, sysTimerStartFreerun, sysTimerReadCounter, sysTimerGetCounterPeriod, sysTimerGetCounterFrequency, sysTimerStop – system timer management
<b>SYNOPSIS</b>	<pre>#include &lt;perf/chBench.h&gt; KnError sysTimerStartPeriodic(KnTimeVal * period, TimerTickHandler handler, void * cookie);  KnError sysTimerStartFreerun(KnTimeVal * minperiod);  unsigned long sysTimerReadCounter(void);  unsigned long sysTimerGetCounterPeriod(void);  unsigned long sysTimerGetCounterFrequency(void);  void sysTimerStop(void);</pre>
<b>FEATURES</b>	PERF
<b>DESCRIPTION</b>	<p>The system timer management services are provided by the microkernel to share the system timer with the application. They are intended mainly for benchmarking and profiling.</p> <p>Typically, bench programs require free running timers and profiling handlers require to be invoked periodically.</p> <p>sysTimerStartPeriodic must be issued to share the system timer between the system and the application. This starts the timer in periodic mode.</p> <p>The sysTimerStartPeriodic is restricted to SUPERVISOR actors.</p> <p>In case of failure, an error code is returned as described below:</p> <p>K_EBUSY           The system timer is already in use by another application.</p> <p>K_EINVAL           The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.</p> <p>K_EFAULT           <i>KnTimeVal *period</i> points to an illegal address.</p> <p>K_EFAIL            The timer can't be initialized with the required period.</p> <p>The <i>handler</i> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level. The handler is invoked passing <i>cookie</i> as the only argument.</p>

`sysTimerStartFreerun` disconnects the system periodic timer. This starts the timer in freerun mode. The *minperiod* specifies the minimum period required for the free-run timer. The implementation is free to program the timer with any period higher than the required one.

In case of failure, an error code is returned as described below:

`K_EBUSY`            The system timer is already shared by another application.  
`K_EFAIL`            The timer can't be initialized with the required interval.

The `sysTimerReadCounter` routine returns the current value of the timer counter. The value can be used to measure the elapsed time between two calls to the `read_counter` routine. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterPeriod` returns the difference between the maximum and minimum values of the timer counter. The routine may be called when the timer operates in any mode. The `get_counter_period` routine may be called from interrupt level. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterFrequency` returns the frequency of the timer counter in Hz. The routine may be called when the timer is operating in any mode. The `get_counter_frequency` routine may be called from interrupt level.

Invocation of the above three calls without a previous call to `sysTimerStartPeriodic` or `sysTimerStartFreerun` returns an undocumented value.

`sysTimerStop` returns the timer to the system. The timer is reprogrammed to the system frequency.

#### ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

#### SEE ALSO

`timer(9DDI)`

<b>NAME</b>	sysTimer, sysTimerStartPeriodic, sysTimerStartFreerun, sysTimerReadCounter, sysTimerGetCounterPeriod, sysTimerGetCounterFrequency, sysTimerStop – system timer management								
<b>SYNOPSIS</b>	<pre>#include &lt;perf/chBench.h&gt; KnError sysTimerStartPeriodic(KnTimeVal * period, TimerTickHandler handler, void * cookie);  KnError sysTimerStartFreerun(KnTimeVal * minperiod);  unsigned long sysTimerReadCounter(void);  unsigned long sysTimerGetCounterPeriod(void);  unsigned long sysTimerGetCounterFrequency(void);  void sysTimerStop(void);</pre>								
<b>FEATURES</b>	PERF								
<b>DESCRIPTION</b>	<p>The system timer management services are provided by the microkernel to share the system timer with the application. They are intended mainly for benchmarking and profiling.</p> <p>Typically, bench programs require free running timers and profiling handlers require to be invoked periodically.</p> <p>sysTimerStartPeriodic must be issued to share the system timer between the system and the application. This starts the timer in periodic mode.</p> <p>The sysTimerStartPeriodic is restricted to SUPERVISOR actors.</p> <p>In case of failure, an error code is returned as described below:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">K_EBUSY</td> <td>The system timer is already in use by another application.</td> </tr> <tr> <td style="padding-right: 20px;">K_EINVAL</td> <td>The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAULT</td> <td><i>KnTimeVal *period</i> points to an illegal address.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAIL</td> <td>The timer can't be initialized with the required period.</td> </tr> </table> <p>The <i>handler</i> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level. The handler is invoked passing <i>cookie</i> as the only argument.</p>	K_EBUSY	The system timer is already in use by another application.	K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.	K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.	K_EFAIL	The timer can't be initialized with the required period.
K_EBUSY	The system timer is already in use by another application.								
K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.								
K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.								
K_EFAIL	The timer can't be initialized with the required period.								

`sysTimerStartFreerun` disconnects the system periodic timer. This starts the timer in freerun mode. The *minperiod* specifies the minimum period required for the free-run timer. The implementation is free to program the timer with any period higher than the required one.

In case of failure, an error code is returned as described below:

- `K_EBUSY`            The system timer is already shared by another application.
- `K_EFAIL`            The timer can't be initialized with the required interval.

The `sysTimerReadCounter` routine returns the current value of the timer counter. The value can be used to measure the elapsed time between two calls to the `read_counter` routine. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterPeriod` returns the difference between the maximum and minimum values of the timer counter. The routine may be called when the timer operates in any mode. The `get_counter_period` routine may be called from interrupt level. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterFrequency` returns the frequency of the timer counter in Hz. The routine may be called when the timer is operating in any mode. The `get_counter_frequency` routine may be called from interrupt level.

Invocation of the above three calls without a previous call to `sysTimerStartPeriodic` or `sysTimerStartFreerun` returns an undocumented value.

`sysTimerStop` returns the timer to the system. The timer is reprogrammed to the system frequency.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`timer(9DDI)`

<b>NAME</b>	sysTimer, sysTimerStartPeriodic, sysTimerStartFreerun, sysTimerReadCounter, sysTimerGetCounterPeriod, sysTimerGetCounterFrequency, sysTimerStop – system timer management
<b>SYNOPSIS</b>	<pre>#include &lt;perf/chBench.h&gt; KnError sysTimerStartPeriodic(KnTimeVal * period, TimerTickHandler handler, void * cookie);  KnError sysTimerStartFreerun(KnTimeVal * minperiod);  unsigned long sysTimerReadCounter(void);  unsigned long sysTimerGetCounterPeriod(void);  unsigned long sysTimerGetCounterFrequency(void);  void sysTimerStop(void);</pre>
<b>FEATURES</b>	PERF
<b>DESCRIPTION</b>	<p>The system timer management services are provided by the microkernel to share the system timer with the application. They are intended mainly for benchmarking and profiling.</p> <p>Typically, bench programs require free running timers and profiling handlers require to be invoked periodically.</p> <p>sysTimerStartPeriodic must be issued to share the system timer between the system and the application. This starts the timer in periodic mode.</p> <p>The sysTimerStartPeriodic is restricted to SUPERVISOR actors.</p> <p>In case of failure, an error code is returned as described below:</p> <p>K_EBUSY           The system timer is already in use by another application.</p> <p>K_EINVAL           The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.</p> <p>K_EFAULT           <i>KnTimeVal *period</i> points to an illegal address.</p> <p>K_EFAIL            The timer can't be initialized with the required period.</p> <p>The <i>handler</i> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level. The handler is invoked passing <i>cookie</i> as the only argument.</p>

sysTimerStartFreerun disconnects the system periodic timer. This starts the timer in freerun mode. The *minperiod* specifies the minimum period required for the free-run timer. The implementation is free to program the timer with any period higher than the required one.

In case of failure, an error code is returned as described below:

- K\_EBUSY            The system timer is already shared by another application.
- K\_EFAIL            The timer can't be initialized with the required interval.

The sysTimerReadCounter routine returns the current value of the timer counter. The value can be used to measure the elapsed time between two calls to the read\_counter routine. Invocation of this call without a previous call to share the system timer returns an undocumented value.

sysTimerGetCounterPeriod returns the difference between the maximum and minimum values of the timer counter. The routine may be called when the timer operates in any mode. The get\_counter\_period routine may be called from interrupt level. Invocation of this call without a previous call to share the system timer returns an undocumented value.

sysTimerGetCounterFrequency returns the frequency of the timer counter in Hz. The routine may be called when the timer is operating in any mode. The get\_counter\_frequency routine may be called from interrupt level.

Invocation of the above three calls without a previous call to sysTimerStartPeriodic or sysTimerStartFreerun returns an undocumented value.

sysTimerStop returns the timer to the system. The timer is reprogrammed to the system frequency.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

timer(9DDI)

<b>NAME</b>	sysTimer, sysTimerStartPeriodic, sysTimerStartFreerun, sysTimerReadCounter, sysTimerGetCounterPeriod, sysTimerGetCounterFrequency, sysTimerStop – system timer management								
<b>SYNOPSIS</b>	<pre>#include &lt;perf/chBench.h&gt; KnError sysTimerStartPeriodic(KnTimeVal * period, TimerTickHandler handler, void * cookie);  KnError sysTimerStartFreerun(KnTimeVal * minperiod);  unsigned long sysTimerReadCounter(void);  unsigned long sysTimerGetCounterPeriod(void);  unsigned long sysTimerGetCounterFrequency(void);  void sysTimerStop(void);</pre>								
<b>FEATURES</b>	PERF								
<b>DESCRIPTION</b>	<p>The system timer management services are provided by the microkernel to share the system timer with the application. They are intended mainly for benchmarking and profiling.</p> <p>Typically, bench programs require free running timers and profiling handlers require to be invoked periodically.</p> <p>sysTimerStartPeriodic must be issued to share the system timer between the system and the application. This starts the timer in periodic mode.</p> <p>The sysTimerStartPeriodic is restricted to SUPERVISOR actors.</p> <p>In case of failure, an error code is returned as described below:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">K_EBUSY</td> <td>The system timer is already in use by another application.</td> </tr> <tr> <td style="padding-right: 20px;">K_EINVAL</td> <td>The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAULT</td> <td><i>KnTimeVal *period</i> points to an illegal address.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAIL</td> <td>The timer can't be initialized with the required period.</td> </tr> </table> <p>The <i>handler</i> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level. The handler is invoked passing <i>cookie</i> as the only argument.</p>	K_EBUSY	The system timer is already in use by another application.	K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.	K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.	K_EFAIL	The timer can't be initialized with the required period.
K_EBUSY	The system timer is already in use by another application.								
K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.								
K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.								
K_EFAIL	The timer can't be initialized with the required period.								

`sysTimerStartFreerun` disconnects the system periodic timer. This starts the timer in freerun mode. The *minperiod* specifies the minimum period required for the free-run timer. The implementation is free to program the timer with any period higher than the required one.

In case of failure, an error code is returned as described below:

- `K_EBUSY`            The system timer is already shared by another application.
- `K_EFAIL`            The timer can't be initialized with the required interval.

The `sysTimerReadCounter` routine returns the current value of the timer counter. The value can be used to measure the elapsed time between two calls to the `read_counter` routine. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterPeriod` returns the difference between the maximum and minimum values of the timer counter. The routine may be called when the timer operates in any mode. The `get_counter_period` routine may be called from interrupt level. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterFrequency` returns the frequency of the timer counter in Hz. The routine may be called when the timer is operating in any mode. The `get_counter_frequency` routine may be called from interrupt level.

Invocation of the above three calls without a previous call to `sysTimerStartPeriodic` or `sysTimerStartFreerun` returns an undocumented value.

`sysTimerStop` returns the timer to the system. The timer is reprogrammed to the system frequency.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`timer(9DDI)`

<b>NAME</b>	sysTimer, sysTimerStartPeriodic, sysTimerStartFreerun, sysTimerReadCounter, sysTimerGetCounterPeriod, sysTimerGetCounterFrequency, sysTimerStop – system timer management								
<b>SYNOPSIS</b>	<pre>#include &lt;perf/chBench.h&gt; KnError sysTimerStartPeriodic(KnTimeVal * period, TimerTickHandler handler, void * cookie);  KnError sysTimerStartFreerun(KnTimeVal * minperiod);  unsigned long sysTimerReadCounter(void);  unsigned long sysTimerGetCounterPeriod(void);  unsigned long sysTimerGetCounterFrequency(void);  void sysTimerStop(void);</pre>								
<b>FEATURES</b>	PERF								
<b>DESCRIPTION</b>	<p>The system timer management services are provided by the microkernel to share the system timer with the application. They are intended mainly for benchmarking and profiling.</p> <p>Typically, bench programs require free running timers and profiling handlers require to be invoked periodically.</p> <p>sysTimerStartPeriodic must be issued to share the system timer between the system and the application. This starts the timer in periodic mode.</p> <p>The sysTimerStartPeriodic is restricted to SUPERVISOR actors.</p> <p>In case of failure, an error code is returned as described below:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">K_EBUSY</td> <td>The system timer is already in use by another application.</td> </tr> <tr> <td style="padding-right: 20px;">K_EINVAL</td> <td>The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAULT</td> <td><i>KnTimeVal *period</i> points to an illegal address.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAIL</td> <td>The timer can't be initialized with the required period.</td> </tr> </table> <p>The <i>handler</i> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level. The handler is invoked passing <i>cookie</i> as the only argument.</p>	K_EBUSY	The system timer is already in use by another application.	K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.	K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.	K_EFAIL	The timer can't be initialized with the required period.
K_EBUSY	The system timer is already in use by another application.								
K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.								
K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.								
K_EFAIL	The timer can't be initialized with the required period.								

sysTimerStartFreerun disconnects the system periodic timer. This starts the timer in freerun mode. The *minperiod* specifies the minimum period required for the free-run timer. The implementation is free to program the timer with any period higher than the required one.

In case of failure, an error code is returned as described below:

- K\_EBUSY            The system timer is already shared by another application.
- K\_EFAIL            The timer can't be initialized with the required interval.

The sysTimerReadCounter routine returns the current value of the timer counter. The value can be used to measure the elapsed time between two calls to the read\_counter routine. Invocation of this call without a previous call to share the system timer returns an undocumented value.

sysTimerGetCounterPeriod returns the difference between the maximum and minimum values of the timer counter. The routine may be called when the timer operates in any mode. The get\_counter\_period routine may be called from interrupt level. Invocation of this call without a previous call to share the system timer returns an undocumented value.

sysTimerGetCounterFrequency returns the frequency of the timer counter in Hz. The routine may be called when the timer is operating in any mode. The get\_counter\_frequency routine may be called from interrupt level.

Invocation of the above three calls without a previous call to sysTimerStartPeriodic or sysTimerStartFreerun returns an undocumented value.

sysTimerStop returns the timer to the system. The timer is reprogrammed to the system frequency.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

timer(9DDI)

<b>NAME</b>	sysTimer, sysTimerStartPeriodic, sysTimerStartFreerun, sysTimerReadCounter, sysTimerGetCounterPeriod, sysTimerGetCounterFrequency, sysTimerStop – system timer management								
<b>SYNOPSIS</b>	<pre>#include &lt;perf/chBench.h&gt; KnError sysTimerStartPeriodic(KnTimeVal * period, TimerTickHandler handler, void * cookie);  KnError sysTimerStartFreerun(KnTimeVal * minperiod);  unsigned long sysTimerReadCounter(void);  unsigned long sysTimerGetCounterPeriod(void);  unsigned long sysTimerGetCounterFrequency(void);  void sysTimerStop(void);</pre>								
<b>FEATURES</b>	PERF								
<b>DESCRIPTION</b>	<p>The system timer management services are provided by the microkernel to share the system timer with the application. They are intended mainly for benchmarking and profiling.</p> <p>Typically, bench programs require free running timers and profiling handlers require to be invoked periodically.</p> <p>sysTimerStartPeriodic must be issued to share the system timer between the system and the application. This starts the timer in periodic mode.</p> <p>The sysTimerStartPeriodic is restricted to SUPERVISOR actors.</p> <p>In case of failure, an error code is returned as described below:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">K_EBUSY</td> <td>The system timer is already in use by another application.</td> </tr> <tr> <td style="padding-right: 20px;">K_EINVAL</td> <td>The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAULT</td> <td><i>KnTimeVal *period</i> points to an illegal address.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAIL</td> <td>The timer can't be initialized with the required period.</td> </tr> </table> <p>The <i>handler</i> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level. The handler is invoked passing <i>cookie</i> as the only argument.</p>	K_EBUSY	The system timer is already in use by another application.	K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.	K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.	K_EFAIL	The timer can't be initialized with the required period.
K_EBUSY	The system timer is already in use by another application.								
K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.								
K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.								
K_EFAIL	The timer can't be initialized with the required period.								

`sysTimerStartFreerun` disconnects the system periodic timer. This starts the timer in freerun mode. The *minperiod* specifies the minimum period required for the free-run timer. The implementation is free to program the timer with any period higher than the required one.

In case of failure, an error code is returned as described below:

- `K_EBUSY`           The system timer is already shared by another application.
- `K_EFAIL`           The timer can't be initialized with the required interval.

The `sysTimerReadCounter` routine returns the current value of the timer counter. The value can be used to measure the elapsed time between two calls to the `read_counter` routine. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterPeriod` returns the difference between the maximum and minimum values of the timer counter. The routine may be called when the timer operates in any mode. The `get_counter_period` routine may be called from interrupt level. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterFrequency` returns the frequency of the timer counter in Hz. The routine may be called when the timer is operating in any mode. The `get_counter_frequency` routine may be called from interrupt level.

Invocation of the above three calls without a previous call to `sysTimerStartPeriodic` or `sysTimerStartFreerun` returns an undocumented value.

`sysTimerStop` returns the timer to the system. The timer is reprogrammed to the system frequency.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`timer(9DDI)`

<b>NAME</b>	sysTimer, sysTimerStartPeriodic, sysTimerStartFreerun, sysTimerReadCounter, sysTimerGetCounterPeriod, sysTimerGetCounterFrequency, sysTimerStop – system timer management								
<b>SYNOPSIS</b>	<pre>#include &lt;perf/chBench.h&gt; KnError sysTimerStartPeriodic(KnTimeVal * period, TimerTickHandler handler, void * cookie);  KnError sysTimerStartFreerun(KnTimeVal * minperiod);  unsigned long sysTimerReadCounter(void);  unsigned long sysTimerGetCounterPeriod(void);  unsigned long sysTimerGetCounterFrequency(void);  void sysTimerStop(void);</pre>								
<b>FEATURES</b>	PERF								
<b>DESCRIPTION</b>	<p>The system timer management services are provided by the microkernel to share the system timer with the application. They are intended mainly for benchmarking and profiling.</p> <p>Typically, bench programs require free running timers and profiling handlers require to be invoked periodically.</p> <p>sysTimerStartPeriodic must be issued to share the system timer between the system and the application. This starts the timer in periodic mode.</p> <p>The sysTimerStartPeriodic is restricted to SUPERVISOR actors.</p> <p>In case of failure, an error code is returned as described below:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">K_EBUSY</td> <td>The system timer is already in use by another application.</td> </tr> <tr> <td style="padding-right: 20px;">K_EINVAL</td> <td>The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAULT</td> <td><i>KnTimeVal *period</i> points to an illegal address.</td> </tr> <tr> <td style="padding-right: 20px;">K_EFAIL</td> <td>The timer can't be initialized with the required period.</td> </tr> </table> <p>The <i>handler</i> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level. The handler is invoked passing <i>cookie</i> as the only argument.</p>	K_EBUSY	The system timer is already in use by another application.	K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.	K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.	K_EFAIL	The timer can't be initialized with the required period.
K_EBUSY	The system timer is already in use by another application.								
K_EINVAL	The given period <i>KnTimeVal *period</i> specified an illegal nanosecond value: lesser than 0 or greater than 1,000,000,000, or a non zero seconds value.								
K_EFAULT	<i>KnTimeVal *period</i> points to an illegal address.								
K_EFAIL	The timer can't be initialized with the required period.								

`sysTimerStartFreerun` disconnects the system periodic timer. This starts the timer in freerun mode. The *minperiod* specifies the minimum period required for the free-run timer. The implementation is free to program the timer with any period higher than the required one.

In case of failure, an error code is returned as described below:

- `K_EBUSY`            The system timer is already shared by another application.
- `K_EFAIL`            The timer can't be initialized with the required interval.

The `sysTimerReadCounter` routine returns the current value of the timer counter. The value can be used to measure the elapsed time between two calls to the `read_counter` routine. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterPeriod` returns the difference between the maximum and minimum values of the timer counter. The routine may be called when the timer operates in any mode. The `get_counter_period` routine may be called from interrupt level. Invocation of this call without a previous call to share the system timer returns an undocumented value.

`sysTimerGetCounterFrequency` returns the frequency of the timer counter in Hz. The routine may be called when the timer is operating in any mode. The `get_counter_frequency` routine may be called from interrupt level.

Invocation of the above three calls without a previous call to `sysTimerStartPeriodic` or `sysTimerStartFreerun` returns an undocumented value.

`sysTimerStop` returns the timer to the system. The timer is reprogrammed to the system frequency.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`timer(9DDI)`

<b>NAME</b>	sysUnsetEnv – delete a value from the ChorusOS configuration environment				
<b>SYNOPSIS</b>	<pre>#include &lt;env/chEnv.h&gt; int sysUnsetEnv(const char *envName);</pre>				
<b>DESCRIPTION</b>	The <code>sysUnsetEnv( )</code> system call deletes the environment variable pointed to by <code>envName</code> . This is a null-terminated string.				
<b>RETURN VALUES</b>	Upon successful completion <code>K_OK</code> is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EFAIL]        <i>envName</i> was not an existing configuration variable in the current ChorusOS environment.</p> <p>[K_EFAULT]      Some of the data provided are outside the current actor's address space.</p> <p>[K_ENOMEM]     For user mode access, the string size of <i>envName</i> exceeds the internal size limit imposed by the user system call implementation. The default value of this limit is 32 bytes.</p>				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>sysSetEnv(2K)</code> , <code>sysGetEnv(2K)</code>				

<b>NAME</b>	sysRead, sysWrite, sysPoll – Read characters from the system console; Write characters to the system console; Poll characters from the system console
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chIo.h&gt; int sysRead(void * buf, unsigned int nchar);  int sysWrite(void * buf, unsigned int nchar);  int sysPoll(void * buf);</pre>
<b>DESCRIPTION</b>	<p>The <i>sysRead</i> system call reads characters from the system console into the buffer pointed to by <i>buf</i>, until either <i>nchar</i> -1 characters are read, or a newline character is read and not transferred to <i>buf</i>. The characters will be echoed on the system console and the string is terminated with a null character. If <i>nchar</i> is equal to 1, one character is read, without echo.</p> <p>The <i>sysWrite</i> system call attempts to write <i>nchar</i> characters to the system console from the buffer pointed to by <i>buf</i>.</p> <p>The <i>sysPoll</i> system call attempts to read one character from the system console into the buffer pointed to by <i>buf</i>. The character will be echoed on the system console. The <i>sysPoll</i> system call is non-blocking.</p>
<b>RETURN VALUE</b>	If successful, <i>sysRead</i> and <i>sysPoll</i> return the number of characters actually read, and <i>sysWrite</i> returns the number of characters actually written. Otherwise, a negative error code is returned.
<b>ERRORS</b>	[K_EFAULT]                      Some of the arguments provided are outside the caller's address space.
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	threadAbort, threadAborted – Abort a thread; Check whether the current thread has been aborted
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadAborted(void);  int threadAbort(KnCap * actorcap, int threadli);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>threadAbort</i> system call aborts the thread whose local identifier is <i>threadli</i> in the actor whose capability is given by <i>actorcap</i> .</p> <p>If <i>actorcap</i> is K_MYACTOR, the aborted thread will be in the current actor.</p> <p>Aborting a thread has different effects depending on the thread's state:</p> <ul style="list-style-type: none"> <li>■ If the thread is blocked by an ABORTABLE call, ( <i>threadDelay</i> (2K)), the thread exits its blocked state and the blocking kernel call returns a specific error code (K_EABORT); the abortion has been processed by the thread.</li> <li>■ Otherwise (the thread is running or blocked in a NON ABORTABLE call), the thread enters the ABORTED state; the abortion still has to be processed. The abortion processing can take three distinct forms: <ul style="list-style-type: none"> <li>■ When a running thread in the ABORTED state issues an ABORTABLE blocking call ( <i>semP</i> (2K)), it does not enter the blocked state; the blocking call returns the specific error code K_EABORT immediately and the thread quits the ABORTED state (the abort has been processed by the thread).</li> <li>■ A running thread can invoke the <i>threadAborted</i> system call at any time, in order to test whether it is in the ABORTED state. If the thread is in the ABORTED state <i>threadAborted</i> resets this state, and returns 1 (the abort has been processed by the thread). Subsequent calls to <i>threadAborted</i> will return 0 (if the thread is not aborted again).</li> <li>■ Finally, SUPERVISOR actors may also be informed via a (asynchronous) handler call, when threads enter the ABORTED state, using <i>svAbortHandler</i> (2K). In this case, when the thread in the ABORTED state returns to execute within its home actor execution environment (see <i>svTrapConnect</i> (2K)), the actor's abort handler is invoked (if any), and the thread quits the ABORTED state (the abort has been processed by the actor). Consequently, calls to <i>threadAborted</i> from the thread will return 0 (if the thread is not aborted again).</li> </ul> </li> </ul>

**RETURN VALUES**

Upon successful completion, *threadAbort* returns 0. Otherwise, a negative error code is returned.

The *threadAborted* call returns 1 if the current thread is in the *ABORTED* state, 0 if not.

**ERRORS**

[K\_EINVAL] *actorcap* is an inconsistent actor capability, or *threadli* is not a valid thread identifier in the specified actor.

[K\_EUNKNOWN] *actorcap* does not specify a reachable actor.

[K\_EFAULT] Some of the data provided are outside the current actor's address space.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*semP(2K)*, *svAbortHandler(2K)*, *svTrapConnect(2K)*, *threadDelay(2K)*

<b>NAME</b>	threadAbort, threadAborted – Abort a thread; Check whether the current thread has been aborted
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadAborted(void);  int threadAbort(KnCap * actorcap, int threadli);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>threadAbort</i> system call aborts the thread whose local identifier is <i>threadli</i> in the actor whose capability is given by <i>actorcap</i> .</p> <p>If <i>actorcap</i> is K_MYACTOR, the aborted thread will be in the current actor.</p> <p>Aborting a thread has different effects depending on the thread's state:</p> <ul style="list-style-type: none"> <li>■ If the thread is blocked by an ABORTABLE call, ( <i>threadDelay</i> (2K)), the thread exits its blocked state and the blocking kernel call returns a specific error code (K_EABORT); the abortion has been processed by the thread.</li> <li>■ Otherwise (the thread is running or blocked in a NON ABORTABLE call), the thread enters the ABORTED state; the abortion still has to be processed. The abortion processing can take three distinct forms: <ul style="list-style-type: none"> <li>■ When a running thread in the ABORTED state issues an ABORTABLE blocking call ( <i>semP</i> (2K)), it does not enter the blocked state; the blocking call returns the specific error code K_EABORT immediately and the thread quits the ABORTED state (the abort has been processed by the thread).</li> <li>■ A running thread can invoke the <i>threadAborted</i> system call at any time, in order to test whether it is in the ABORTED state. If the thread is in the ABORTED state <i>threadAborted</i> resets this state, and returns 1 (the abort has been processed by the thread). Subsequent calls to <i>threadAborted</i> will return 0 (if the thread is not aborted again).</li> <li>■ Finally, SUPERVISOR actors may also be informed via a (asynchronous) handler call, when threads enter the ABORTED state, using <i>svAbortHandler</i> (2K). In this case, when the thread in the ABORTED state returns to execute within its home actor execution environment (see <i>svTrapConnect</i> (2K)), the actor's abort handler is invoked (if any), and the thread quits the ABORTED state (the abort has been processed by the actor). Consequently, calls to <i>threadAborted</i> from the thread will return 0 (if the thread is not aborted again).</li> </ul> </li> </ul>

**RETURN VALUES**

Upon successful completion, *threadAbort* returns 0. Otherwise, a negative error code is returned.

The *threadAborted* call returns 1 if the current thread is in the *ABORTED* state, 0 if not.

**ERRORS**

[K\_EINVAL] *actorcap* is an inconsistent actor capability, or *threadli* is not a valid thread identifier in the specified actor.

[K\_EUNKNOWN] *actorcap* does not specify a reachable actor.

[K\_EFAULT] Some of the data provided are outside the current actor's address space.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*semP(2K)*, *svAbortHandler(2K)*, *svTrapConnect(2K)*, *threadDelay(2K)*

**NAME** threadActivate – make a thread active

**SYNOPSIS** #include <exec/chExec.h>  
int threadActivate(KnCap \*actorcap, KnThreadLid threadli);

**FEATURES** CORE

**DESCRIPTION** The *threadActivate* system call activates thread *threadli* in the actor *actorcap*. If the thread was created in *K\_INACTIVE*, *threadActivate* will make it eligible to execute.

**RETURN VALUE** Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EFAULT]	Some of the data provided are outside the current actor's address space.

**RESTRICTIONS** The target actor and the current actor must be located on the same site.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** threadCreate(2K)

<b>NAME</b>	threadBind – bind a thread to a processor
<b>SYNOPSIS</b>	<pre>#include &lt;sched/chSched.h&gt; int threadBind(KnCap *actorcap, KnThreadLid threadli, KnCpuId newcpunb, KnCpuId *oldcpunb, KnBindId bindid);</pre>
<b>FEATURES</b>	SCHED_FIFO, SCHED_CLASS
<b>DESCRIPTION</b>	<p>On a Shared Memory Processor based Architecture, a thread is bound by default to the processor on which <i>threadCreate</i> was performed. The thread binding can be changed using <i>threadBind</i>.</p> <p>The <i>threadBind</i> system call binds the thread whose local identifier is <i>threadli</i>, in the actor whose capability is given by <i>actorcap</i>, to the processor designated by <i>newcpunb</i>.</p> <p>If the thread is currently running on a different processor, it is cancelled and scheduled for execution on the processor specified.</p> <p>If <i>actorcap</i> is K_MYACTOR, the thread is a thread of the current actor. In this case, if <i>threadli</i> is K_MYSELF, the current thread is used.</p> <p>If <i>newcpunb</i> is K_SAMECPU, the thread keeps its current binding and the function returns the current binding if <i>oldcpunb</i> is not a null pointer.</p> <p>The <i>bindid</i> parameter defines the type of binding applied to <i>threadli</i>. In the current release, the K_NONEXCLUSIVE binding type is mandatory.</p>
<b>RETURN VALUE</b>	Upon successful completion K_OK is returned, otherwise a negative error code is returned.
<b>ERRORS</b>	<p>[K_EINVAL] <i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor, or <i>newcpunb</i> is not a valid processor number, or <i>bindid</i> is not a valid binding type.</p> <p>[K_EUNKNOWN] <i>actorcap</i> does not specify a reachable actor.</p> <p>[K_EFAULT] Some of the data provided are outside the current actor's address space.</p> <p>[K_ENOTIMP] For a given thread, <i>threadBind</i> is not performed from the CPU of the current (old) binding or <i>threadBind</i> is called on a site running a kernel configured for uniprocessor architecture only.</p> <p>[K_EBUSY] <i>bindid</i> is not a valid binding type.</p>
<b>RESTRICTIONS</b>	The target thread and the current thread must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`threadCreate(2K)`

<b>NAME</b>	threadContext – get and/or set the context of a thread
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadContext(KnCap *actorcap, KnThreadLid threadli, unsigned int type, void *oldcontext, void *newcontext);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>threadContext</i> system call gets and/or sets the execution context of the thread whose local identifier is <i>threadli</i>, in the actor whose capability is given by <i>actorcap</i>. If <i>actorcap</i> is K_MYACTOR, the thread is a thread of the current actor. In this case, if <i>threadli</i> is K_MYSELF, the current thread is used.</p> <p>The <i>type</i> parameter selects the component of the execution context which is affected. If <i>type</i> is K_CURRCTX, the thread's <i>hardware</i> context is affected. In this case, <i>oldcontext</i> and <i>newcontext</i> must be pointers to <i>KnThreadCtx</i> structures defined in the <i>exec/chThCtx.f.h</i> header file. The fields of this structure are machine-dependent. A thread's <i>hardware</i> context is defined as the set of general—purpose machine register values (including stack pointer, program counter, and so on) put on the stack when an initial (non-nested) trap, exception, or preemption occurred. Any subsequent nested events (in the case of a trap or exception) have no effect on the thread's hardware context as accessed by K_CURRCTX.</p> <p>If <i>type</i> is K_SUPPRIV, the thread's supervisor software register context is affected (see <i>threadLoadR(2K)</i>). In this case, <i>oldcontext</i> and <i>newcontext</i> must be pointers to an array of <code>void*</code> whose size is K_NBPRIVDATA.</p> <p>If <i>type</i> is K_USERPRIV, the thread's user software register context is affected (see <i>threadLoadR(2K)</i>). In this case, <i>oldcontext</i> and <i>newcontext</i> must be pointers to an array of <code>void*</code> whose size is K_NBPRIVDATA. This type is valid only for threads of user privilege.</p> <p>If <i>type</i> is K_SOFTCTX, the thread's software register context is affected (see <i>threadLoadR(2K)</i>). In this case, <i>oldcontext</i> and <i>newcontext</i> must be pointers to <i>KnThreadSoftCtx</i> structures, whose members are the following:</p> <pre>long userReg ; /* user privilege register */ long supReg ; /* supervisor privilege register */</pre> <p>Note that the caller thread must be a supervisor thread (see <i>threadCreate(2K)</i>) or must belong to a system actor (see <i>actorCreate(2K)</i>) in order to modify the <i>supReg</i> value.</p> <p>In either case, the thread's current execution context is copied into the structure pointed to by <i>oldcontext</i> (if not NULL) in the client address space. The thread's new context will be taken from the structure pointed to by <i>newcontext</i> (if not NULL).</p>

**RETURN VALUE** Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

<b>ERRORS</b>	[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.
	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
	[K_EFAULT]	Some of the data provided is out of the current actor's address space.
	[K_EPRIV]	The caller thread is not allowed to perform the operation.
	[K_ENOTAVAILABLE]	The hardware execution context of the target thread is not available.

**RESTRICTIONS** The target actor and the current actor must be located on the same site.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `actorCreate(2K)`, `threadCreate(2K)`, `threadLoadR(2K)`

<b>NAME</b>	threadCreate – create a thread
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadCreate(KnCap *actorcap, KnThreadLid *threadli, KnThreadStatus status, void *schedparam, void *startinfo);</pre>
<b>DESCRIPTION</b>	<p><i>threadCreate</i> creates a new thread in the actor the capability of which is given by <i>actorcap</i>. If <i>actorcap</i> is <code>K_MYACTOR</code>, the thread is created within the current actor.</p> <p><i>threadCreate</i> returns the local identifier of the new thread in <i>threadli</i>.</p> <p><i>status</i> is the thread's initial status: if <i>status</i> is <code>K_INACTIVE</code>, the thread will not execute until it is explicitly started using a <i>threadActivate(2K)</i>. If <i>status</i> is <code>K_ACTIVE</code>, the thread can run immediately (assuming that the actor is in the <code>ACTIVE</code> state (see <i>actorStart(2K)</i>).</p> <p><i>schedparam</i> points to a descriptor of thread scheduling attributes (see <i>threadScheduler(2K)</i>). If <i>schedparam</i> is 0, the thread inherits the scheduling attributes of the parent thread.</p> <p><i>startinfo</i> points to a structure specifying the initial state of the new thread. The current implementation supports two types of <i>startinfo</i> structures: <i>KnDefaultStartInfo_f</i> and <i>KnStartInfo</i></p> <p>These structures are machine-dependent (<i>i.e.</i> may contain machine-dependent fields).</p> <p>For most conventional processors (eg. x86/Pentium, SPARC, etc.), <i>KnDefaultStartInfo_f</i> contains the following fields:</p> <pre>KnStartInfoType dsType ; unsigned dsSystemStackSize ; KnPc dsEntry ; void* dsUserStackPointer ; KnThreadPrivilege dsPrivilege ;</pre> <p>The <i>dsType</i> field must be set equal to <code>K_DEFAULT_START_INFO</code>.</p> <p>The <i>dsSystemStackSize</i> field provides a hint to the kernel of the supervisor stack size required by the thread. If <i>dsSystemStackSize</i> is equal to <code>K_DEFAULT_STACK_SIZE</code>, the default supervisor stack size will be used.</p> <p>The <i>dsEntry</i> and <i>dsUserStackPointer</i> fields specify initial values of the program counter and the user stack pointer of the thread when activated. The program counter and the stack pointer are the only machine registers the values of which may be directly provided upon thread creation. Other register values may be initialized by creating a thread in <code>K_INACTIVE</code> state and calling <i>threadContext</i> before resuming it (see <i>threadActivate(2K)</i>). Otherwise, these registers are initialized with arbitrary values.</p>

If the *dsPrivilege* field is equal to `K_USERTHREAD`, the thread is activated as a user mode thread with user privileges. If the *dsPrivilege* field is equal to `K_SUPRTHREAD`, the thread is activated at supervisor privilege level and has access to restricted machine instructions and kernel calls. In the latter case, the *dsUserStackPointer* field is ignored.

Similarly, the *KnStartInfo* contains the following fields:

```
KnStartInfoType  dsType ;
unsigned         dsSystemStackSize ;
KnPc            dsEntry ;
void*           dsUserStackPointer ;
KnThreadPrivilege dsPrivilege ;
unsigned int    dsSoftReg ;
void*           dsSystemStackTop ;
```

The *dsType* field must be `K_START_INFO` possibly or'ed with `K_START_INFO_SOFTREG` and/or `K_START_INFO_SYS_STACK` bits.

The meaning of *dsSystemStackSize*, *dsEntry*, *dsUserStackPointer* and *dsPrivilege* fields remains the same as in *KnDefaultStartInfo\_f* structure.

The value *dsSoftReg* field is only meaningful if `K_START_INFO_SOFTREG` bit is set in the *dsType* field. It specifies the initial value of the thread's software register corresponding to its privilege level. This is used to pass an argument to the created thread, which gets this argument using *threadLoadR(2K)*.

The value *dsSystemStackTop* field is only meaningful if `K_START_INFO_SYS_STACK` bit is set in the *dsType* field. It specifies the starting address of the memory area to be used as the created thread's system stack (while *dsSystemStackSize* gives, as usually, the size in bytes of this area).

#### WARNINGS

Note that in the current version, `K_SUSPENDED` is equivalent to `K_INACTIVE` and *threadResume(2K)* is equivalent to *threadActivate(2K)*. `K_SUSPENDED` and *threadResume* are supported for compatibility reasons.

`MEM_PROTECTED` for ARM and i386 as well as `MEM_VIRTUAL` on i386 and ppc60x implementations do not support supervisor stacks greater than a tunable which is defined per platform. To get the default value of the system stack size, it is necessary to invoke the configuration tool (*configurator*).

The supervisor stack must be always allocated by a supervisor thread only. Consequently only supervisor threads are allowed to change the system stack.

#### RETURN VALUE

Upon successful completion, `K_OK` is returned. Otherwise, a negative error code is returned.

#### ERRORS

[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.

[K_EFAULT]	Some of the provided data are outside the current actor's address space.
[K_EPRIV]	The current thread is not allowed to create a <i>SUPERVISOR</i> thread.
[K_ENOTIMP]	<i>schedparam</i> requires a non-supported scheduling policy.
[K_ENOMEM]	The system is out of resources.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`threadActivate(2K)`, `threadResume(2K)`, `actorCreate(2K)`,  
`threadContext(2K)`, `threadScheduler(2K)`

<b>NAME</b>	threadDelay – delay the current thread				
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadDelay(KnTimeVal *waitLimit);</pre>				
<b>FEATURES</b>	CORE				
<b>DESCRIPTION</b>	<p><i>threadDelay</i> delays the execution of the current thread according to the value of <i>*waitLimit</i>.</p> <p><i>waitLimit</i> is a pointer to a <i>KnTimeVal</i> structure containing a timeout interval as described in <i>sysTime(2K)</i>. Semantics and constant option values of <i>waitLimit</i> argument are fully described in <i>intro(2K)</i>. A <i>waitLimit</i> value of K_NOBLOCK (or zero) means that the thread is still ready to run after the call. In this case, <i>threadDelay</i> may have the effect of simply yielding the processor to another thread that is ready to run. The precise actions in this case depend on the scheduler module configured.</p>				
<b>ERRORS</b>	<p>[K_EINVAL]                    The <i>waitLimit</i> structure is not a valid <i>KnTimeVal</i>.</p> <p>[K_ETIMEOUT]                 The time out occurred.</p> <p>[K_EABORT]                    The thread has been aborted while delayed.</p>				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>threadAbort(2K)</i>				

<b>NAME</b>	threadDelete – delete a thread								
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadDelete(KnCap *actorcap, KnThreadLid threadli);</pre>								
<b>DESCRIPTION</b>	<p><i>threadDelete</i> deletes the thread the local identifier of which is <i>threadli</i> in the actor the capability of which is given by <i>actorcap</i>.</p> <p>If <i>actorcap</i> is <i>K_MYACTOR</i>, the deleted thread must belong to the current actor.</p> <p>In this case, if <i>threadli</i> is <i>K_MYSELF</i>, the current thread is deleted.</p>								
<b>RETURN VALUE</b>	Upon successful completion, a value of 0 is returned. Otherwise, a negative error code is returned.								
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td><i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.</td> </tr> <tr> <td style="vertical-align: top;">[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the provided data are outside the current actor's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_BUSY]</td> <td>The thread is already in the process of being deleted.</td> </tr> </table>	[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.	[K_EFAULT]	Some of the provided data are outside the current actor's address space.	[K_BUSY]	The thread is already in the process of being deleted.
[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.								
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.								
[K_EFAULT]	Some of the provided data are outside the current actor's address space.								
[K_BUSY]	The thread is already in the process of being deleted.								
<b>RESTRICTIONS</b>	The target actor and the current actor must be located on the same site.								
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Interface Stability	Evolving								
<b>SEE ALSO</b>	<code>threadCreate(2K)</code>								

<b>NAME</b>	threadLoadR, threadStoreR – Get the current thread's valid soft register value; Reset the current thread's valid soft register value
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chSoftCtx.h&gt; long threadLoadR(void);  long threadStoreR(long reg);</pre>
<b>DESCRIPTION</b>	<p>The execution context of a thread (its various machine registers) is extended by two software registers, one for each of the possible privilege levels of a thread ( <i>USER</i> or <i>SUPERVISOR</i> ). Like hardware registers, these two registers are saved and restored at each thread switch. These registers provide the means for maintaining per-thread data efficiently.</p> <p>The <i>threadLoadR</i> system call returns the value of the current thread's software register corresponding to the current privilege level of the thread.</p> <p>The <i>threadStoreR</i> system call sets the current thread's software register corresponding to the current privilege level of the thread to the value specified by <i>reg</i> , and returns its new value.</p> <p>The values of these registers may also be affected using <i>threadContext</i> (2K).</p>
<b>RETURN VALUE</b>	The <i>threadLoadR</i> call returns the current register value. The <i>threadStoreR</i> call returns the new register value.
<b>EXAMPLES</b>	<p>The software register facility can be used to store the thread's local identifier in its software register. However, this only impacts on performance, as <i>threadLoadR</i> does not imply a kernel call on most architectures, this is fully resolved in libraries which access a globally shared system context. However, <i>threadSelf</i> always implies a kernel call. In order to do this, call:</p> <pre>threadStoreR (threadSelf());</pre> <p>at the beginning of the thread's code; then invoke:</p> <pre>threadLoadR ;</pre> <p>instead of:</p> <pre>threadSelfC ;</pre> <p>A more common way to use this facility is to associate a per-thread structure with each thread of an actor, and to store the address of this structure in the software register. This will avoid systematically transferring this address as a parameter across the layers of the actor's code (the use of a global variable is not possible, especially on multiprocessor implementations). In this case, each actor's thread will perform a</p>

```
threadStoreR (myArea);
```

at the beginning of the thread's code, where myArea is the address of the per thread's structure. Note that there are several ways for the thread to acquire the

```
(myArea);
```

value: it may be provided to the thread's routine as an argument (by initializing its stack correctly before creation); it may also be computed by the created thread as the result of a hash function (the hash function taking the thread's local identifier as input, and returning an entry in the array of per-thread records). An alternative to this initialization process is for the creator to call *threadContext* (2K) before resuming the thread. In each case, the per-thread record is subsequently accessed by:

```
threadLoadR ();
```

Finally, note that the software register may be used differently by different code layers: the sequence S in:

```
oldVal = threadLoadR ();
threadStoreR (newVal);
        /* S */
threadStoreR (oldVal);
```

may make private use of the software register value.

#### ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

#### SEE ALSO

*threadContext*(2K)

**NAME** threadName – get and/or set the symbolic name of a thread

**SYNOPSIS**

```
#include <exec/chExec.h>
int threadName(KnCap *actorcap, int threadli, char *oldname, char *newname);
```

**FEATURES** CORE

**DESCRIPTION** The *threadName* system call gets and/or sets the symbolic name of the thread whose local identifier is *threadli*, in the actor whose capability is given by *actorcap*. If *actorcap* is K\_MYACTOR, the target thread must be in the current actor. In this case, if *threadli* is K\_MYSELF, the current thread is affected.

If *oldname* is not a NULL pointer, the thread’s symbolic name is copied to the caller address space at the location specified by *oldname*. A thread’s symbolic name has a maximum size of K\_THREADNAMEMAX (including the NULL character).

If *newname* is not a NULL pointer, the thread’s symbolic name is set to the new name pointed to by *newname* in the caller address space. The new name is truncated to a maximum size of K\_THREADNAMEMAX (including the NULL character).

**RETURN VALUE** Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K_EINVAL]	<i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EFAULT]	Some of the data provided are outside the current actor’s address space.

**RESTRICTIONS** The target actor and the current actor must be located on the same site.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *threadCreate(2K)*

<b>NAME</b>	threadSuspend, threadResume – Suspend a thread; Resume a thread
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadSuspend(KnCap * actorcap, int threadli);  int threadResume(KnCap * actorcap, int threadli);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>NOTE :</b> These calls are obsolete. Although they are maintained for compatibility reasons, they should be avoided wherever possible. See <i>threadActivate(2K)</i>, <i>threadStart(2K)</i>, and <i>threadSemInit(2K)</i> for documentation on other calls that can be used instead of <i>threadSuspend(2K)</i> and <i>threadResume(2K)</i>. The <i>threadSuspend</i> system call puts the thread whose local identifier is <i>threadli</i>, in the actor whose capability is given by <i>actorcap</i>, into the <i>SUSPENDED</i> state.</p> <p>A suspended thread may resume execution (enter the <i>ACTIVE</i> state) only if the thread is explicitly resumed using <i>threadResume(2K)</i>.</p> <p>The <i>threadSuspend</i> and <i>threadResume</i> system calls are based on a per-thread suspend counter, initialized to 0 at thread creation, incremented by <i>threadSuspend</i> and decremented by <i>threadResume</i>.</p> <p>The <i>threadResume</i> system call decrements the suspend counter of the thread whose local identifier is <i>threadli</i>, in the actor whose capability is given by <i>actorcap</i>. It only awakens the thread by setting its state to <i>ACTIVE</i>, if the suspend counter has become 0.</p> <p>If a thread which is in a <i>BLOCKED</i> state is suspended using <i>threadSuspend</i>, the thread will suspend just after it exits its <i>BLOCKED</i> state. If it is resumed using <i>threadResume</i> while still in the <i>BLOCKED</i> state, it will continue its execution normally after it exits its <i>BLOCKED</i> state (if its suspend counter has become 0). The same rule applies for a thread whose actor is in the <i>STOPPED</i> state.</p> <p>If <i>actorcap</i> is <i>K_MYACTOR</i>, the thread must be a thread of the current actor. In this case, if <i>threadli</i> is <i>K_MYSELF</i>, the current thread is used.</p>
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	<p>[<i>K_EINVAL</i>] <i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.</p> <p>[<i>K_EUNKNOWN</i>] <i>actorcap</i> does not specify a reachable actor.</p> <p>[<i>K_EFAULT</i>] Some of the data provided are outside the current actor's address space.</p>

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`threadCreate(2K)`

<b>NAME</b>	threadScheduler – get and/or set thread scheduling parameters
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; #include &lt;sched/chSched.h&gt; #include &lt;sched/chFifo.h&gt; #include &lt;sched/chRr.h&gt; #include &lt;sched/chRt.h&gt; int threadScheduler(KnCap *actorcap, KnThreadLid threadli, void *oldpara, void *newparam);</pre>
<b>FEATURES</b>	SCHED_FIFO, SCHED_CLASS
<b>DESCRIPTION</b>	<p><i>threadScheduler</i> returns and optionally modifies the scheduling parameters of the thread the local identifier of which is <i>threadli</i> in the actor the capability of which is given by <i>actorcap</i>. If <i>actorcap</i> is K_MYACTOR, the thread must belong to the current actor. In this case, if <i>threadli</i> is K_MYSELF, the current thread is affected.</p> <p><i>oldparam</i> and <i>newparam</i> point to descriptors of old and new scheduling parameters respectively. The type of the descriptor depends on the scheduler. Generally, there is one type of scheduling parameters per scheduling policy (i.e. scheduling class) implemented by the scheduler.</p> <p>Two priority-based schedulers are currently supported: SCHED_FIFO and SCHED_CLASS. The SCHED_FIFO scheduler implements the CLASS_FIFO scheduling policy, which manages parameters of type <i>KnFifoThParms</i>. The SCHED_CLASS scheduler implements the CLASS_FIFO, CLASS_RR and CLASS_RT scheduling policies, which respectively manage parameters of type <i>KnFifoThParms</i>, <i>KnRrThParms</i> and <i>KnRtThParms</i>. The SCHED_CLASS also implements the CLASS_TS that is reserved for the implementation of the Chorus/MiX SVR4 subsystem, and therefore should not be used. These policies, as well as the corresponding parameter structures, are described within specific sections below.</p> <p>If <i>oldparam</i> is not NULL, the thread's current scheduling parameters are stored in the area pointed to by <i>oldparam</i> in the caller's address space.</p> <p>If <i>newparam</i> is not NULL, it points to an area in the caller's address space containing the new scheduling parameters.</p> <p>In addition, <i>threadScheduler</i>, as well as the two available schedulers, support the notion of <i>default scheduling class</i>, which defines a normalized scheduling parameters structure, <i>KnThreadDefaultSched</i>. <i>newparms</i> can point to a variable of this type. <i>KnThreadDefaultSched</i> is a structure containing the following fields:</p> <pre>KnSchedClass      tdClass ; KnThreadPriority  tdPriority ;</pre>

where *tdClass* must be equal to `K_SCHED_DEFAULT`. *tdPriority* establishes the priority of the thread, which may vary between `K_PRIOMAX` (highest priority) and `K_PRIOMIN` (lowest priority). The two available schedulers currently map this (virtual) default scheduling class to the `CLASS_FIFO` policy. This means in particular that they interpret the *KnThreadDefaultSched* priority field in the same way the `CLASS_FIFO` interpretes the priority field of a *KnFifoThParms* structure.

**CLASS\_FIFO**

The `CLASS_FIFO` scheduling class implements a pure priority-based, preemptive, fifo policy. Thread's scheduling parameters are defined by the *KnFifoThParms* structure, the members of which are the following:

```
KnSchedClass    fifoClass ;
KnFifoPriority   fifoPriority ;
```

*fifoClass* must be equal to `K_SCHED_FIFO`. *fifoPriority* defines the priority of the thread, which may vary between `K_FIFO_PRIOMAX` (0, highest priority) and `K_FIFO_PRIOMIN` (255, lowest priority).

`CLASS_FIFO` uses the full range of priorities (256) in both the `SCHED_FIFO` and `SCHED_CLASS` schedulers.

A thread which becomes ready to run after being blocked is always inserted last at the end of its priority ready queue (fifo policy). This means that a running thread is only preempted if a thread of strictly highest priority becomes ready to run. A preempted thread is placed at the head of its priority ready queue. This means that it will be elected first when the preempting thread completes or is blocked.

**CLASS\_RR**

The `CLASS_RR` scheduling class is only available within the `SCHED_CLASS` scheduler. It implements a priority-based, preemptive, round-robin policy. Thread's scheduling parameters are defined by the *KnRrThParms* structure, the members of which are the following:

```
KnSchedClass    rrClass ;
KnRrPriority     rrPriority ;
```

*rrClass* must be equal to `K_SCHED_RR`. *rrPriority* defines the priority of the thread, which may vary between `K_RR_PRIOMAX` (0, highest priority) and `K_RR_PRIOMIN` (255, lowest priority).

`CLASS_RR` uses the full range of priorities (256) of the `SCHED_CLASS` scheduler.

The `SCHED_RR` policy is similar to the `SCHED_FIFO` policy, except that an elected thread is given a fixed time quantum. If the thread is still running at quantum expiration, it is placed at the end of its priority ready queue, and then may be preempted by threads of equal priority. The thread's quantum is reset when the thread becomes ready after being blocked; it is not reset when the thread is elected again after a preemption.

**CLASS\_RT**

The time quantum is the same for all priority levels and all threads. It is defined by the `K_RR_QUANTUM` value (100 milliseconds). It will be turned into a system tunable in a future release.

The `CLASS_RT` scheduling class is only available within the `SCHED_CLASS` scheduler. It implements the same policy as the real-time class of UNIX SVR4.0. Refer to the `priocntl(2)` manual entry of UNIX SVR4.0 for a complete description.

This policy is basically a round-robin policy, with per thread time quanta.

Thread's scheduling parameters are defined by the `KnRtThParms` structure, the members of which are the following:

```
KnSchedClass  rtClass ;
KnRtParms     *rtParms ;
```

`rtClass` must be equal to `K_SCHED_RT`.

The actual parameters are defined by the `KnRtParms` structure, the members of which are the following:

```
KnRtPriority   rtPriority ;
unsigned long  rtQSecs ;
long           rtQNSecs ;
```

The fields of this structure map the fields of the `rtparms_t` of UNIX SVR4. `rtPriority` defines the priority of the thread, which may vary between `K_RT_PRIOMAX` (159, highest priority) and `K_RT_PRIOMIN` (100, lowest priority). Note that the order of priorities is inverted compared to the `CLASS_FIFO` priorities. `CLASS_RT` uses only a sub-range of the `SCHED_CLASS` priorities. This sub-range corresponds to the range [96, 155] of `CLASS_FIFO` and `CLASS_RR`.

The `rtQSec` and `rtQNsec` define the thread's time quantum, respectively in seconds and nanoseconds.

The `rtQNsec` field can take the following special values:

`RT_SAMEQUANTUM`      The thread's quantum must not be modified, or if the thread enters the `CLASS_RT` class, the thread's quantum is set to the `RT_INFINITE` quantum.

`RT_INFINITE`            The thread's time quantum is infinite.

`RT_DEFAULT`            The thread's time quantum is set to the default for the thread's priority, as follows:

K_RT_PRIOMIN to K_RT_PRIOMIN+9	100 ticks
K_RT_PRIOMIN+10 to K_RT_PRIOMIN+19	80 ticks
K_RT_PRIOMIN+20 to K_RT_PRIOMIN+29	60 ticks
K_RT_PRIOMIN+30 to K_RT_PRIOMIN+39	40 ticks
K_RT_PRIOMIN+40 to K_RT_PRIOMIN+49	20 ticks
K_RT_PRIOMIN+50 to K_RT_PRIOMIN+59	10 ticks

The default tick period is 10 ms. In this case, the values in the table should be multiplied by ten in order to get the quantum in ms.

The CLASS\_RT manages per configurable priority default time quanta.

**CLASS\_TS**

The CLASS\_TS scheduling class is reserved and implements the same policy as the time-sharing class of UNIX SVR4.0.

**RETURN VALUE**

Upon successful completion, K\_OK is returned. Otherwise, a negative error code is returned.

**ERRORS**

- [K\_EINVAL] *actorcap* is an inconsistent actor capability.
- [K\_EUNKNOWN] *actorcap* does not specify a reachable actor.
- [K\_EINVAL] *newparam* points to a structure containing invalid scheduling parameters.
- [K\_ENOMEM] The system is out of resources.
- [K\_ENOTIMP] *newparm* specifies a non-supported scheduling policy.
- [K\_EFAULT] Some of the provided data are outside the current actor's address space.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

threadCreate(2K)

**NAME** threadSelf – get the current thread local identifier

**SYNOPSIS** #include <exec/chExec.h>  
KnThreadLid **threadSelf**(void);

**DESCRIPTION** *threadSelf* returns the current thread's local identifier in its actor.

**RETURN VALUE** The current thread's local identifier is returned.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *threadCreate(2K)*

<b>NAME</b>	threadSemInit, threadSemPost, threadSemWait – Initialize a thread semaphore; Signal a thread semaphore; Wait on a thread semaphore				
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadSemInit(KnThSem * sem);  int threadSemPost(KnThSem * sem);  int threadSemWait(KnThSem * sem, KnTimeVal * waitLimit);</pre>				
<b>DESCRIPTION</b>	<p>Thread semaphores are <i>KnThSem</i> structures allocated in user memory. A thread semaphore is a binary semaphore on which one, and only one, thread can wait, and which can be signalled to by multiple threads and/or interrupt handlers.</p> <p>A thread semaphore can be in two states: POSTED and UNPOSTED.</p> <p>The <i>threadSemInit</i> system call initializes the thread semaphore whose address is <i>sem</i> in the UNPOSTED state. The thread which performs <i>threadSemInit</i> on a given thread semaphore is the only thread which can wait on the semaphore (using <i>threadSemWait</i> ).</p> <p>The <i>threadSemWait</i> system call makes the current thread wait conditionally on the thread semaphore <i>sem</i>. If the thread semaphore is in the POSTED state, <i>threadSemWait</i> returns immediately and atomically changes the state of the thread semaphore to UNPOSTED. If the thread semaphore is in the UNPOSTED state, the thread is blocked according to the options described by <i>waitLimit</i> in <i>intro(2K)</i>. The <i>waitLimit</i> pointer points to a <i>KnTimeVal</i> structure containing a timeout interval as described in <i>sysTime(2K)</i>.</p> <p>The <i>threadSemPost</i> system call signals the thread semaphore <i>sem</i>. If the thread that initialized the thread semaphore is blocked behind the semaphore, <i>threadSemPost</i> awakens that thread without modifying the state of <i>sem</i>. If the thread is not blocked, <i>threadSemPost</i> sets the state of <i>sem</i> to POSTED. If the thread is not blocked and <i>sem</i> is already in the POSTED state, <i>threadSemPost</i> has no effect.</p> <p>The <i>threadSemPost</i> system call may be called within an interrupt handler, or with preemption disabled.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.				
<b>RESTRICTIONS</b>	<p>A user application and a supervisor application may not share a semaphore.</p> <p>However, two user applications may share a semaphore by mapping it in both user address spaces.</p>				
<b>ERRORS</b>	<table border="0"> <tr> <td style="padding-right: 20px;">[K_EABORT]</td> <td><i>threadSemWait</i> has been aborted.</td> </tr> <tr> <td>[K_EFAULT]</td> <td><i>sem</i> points outside the current actor's address space.</td> </tr> </table>	[K_EABORT]	<i>threadSemWait</i> has been aborted.	[K_EFAULT]	<i>sem</i> points outside the current actor's address space.
[K_EABORT]	<i>threadSemWait</i> has been aborted.				
[K_EFAULT]	<i>sem</i> points outside the current actor's address space.				

[K\_EINVAL] The thread semaphore structure has not been correctly initialized, or the *waitLimit* structure is not a valid *KnTimeVal*.

[K\_ETIMEOUT] The timeout occurred.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	threadSemInit, threadSemPost, threadSemWait – Initialize a thread semaphore; Signal a thread semaphore; Wait on a thread semaphore				
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadSemInit(KnThSem * sem);  int threadSemPost(KnThSem * sem);  int threadSemWait(KnThSem * sem, KnTimeVal * waitLimit);</pre>				
<b>DESCRIPTION</b>	<p>Thread semaphores are <i>KnThSem</i> structures allocated in user memory. A thread semaphore is a binary semaphore on which one, and only one, thread can wait, and which can be signalled to by multiple threads and/or interrupt handlers.</p> <p>A thread semaphore can be in two states: POSTED and UNPOSTED.</p> <p>The <i>threadSemInit</i> system call initializes the thread semaphore whose address is <i>sem</i> in the UNPOSTED state. The thread which performs <i>threadSemInit</i> on a given thread semaphore is the only thread which can wait on the semaphore (using <i>threadSemWait</i> ).</p> <p>The <i>threadSemWait</i> system call makes the current thread wait conditionally on the thread semaphore <i>sem</i>. If the thread semaphore is in the POSTED state, <i>threadSemWait</i> returns immediately and atomically changes the state of the thread semaphore to UNPOSTED. If the thread semaphore is in the UNPOSTED state, the thread is blocked according to the options described by <i>waitLimit</i> in <i>intro(2K)</i>. The <i>waitLimit</i> pointer points to a <i>KnTimeVal</i> structure containing a timeout interval as described in <i>sysTime(2K)</i>.</p> <p>The <i>threadSemPost</i> system call signals the thread semaphore <i>sem</i>. If the thread that initialized the thread semaphore is blocked behind the semaphore, <i>threadSemPost</i> awakens that thread without modifying the state of <i>sem</i>. If the thread is not blocked, <i>threadSemPost</i> sets the state of <i>sem</i> to POSTED. If the thread is not blocked and <i>sem</i> is already in the POSTED state, <i>threadSemPost</i> has no effect.</p> <p>The <i>threadSemPost</i> system call may be called within an interrupt handler, or with preemption disabled.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.				
<b>RESTRICTIONS</b>	<p>A user application and a supervisor application may not share a semaphore.</p> <p>However, two user applications may share a semaphore by mapping it in both user address spaces.</p>				
<b>ERRORS</b>	<table border="0"> <tr> <td style="padding-right: 20px;">[K_EABORT]</td> <td><i>threadSemWait</i> has been aborted.</td> </tr> <tr> <td>[K_EFAULT]</td> <td><i>sem</i> points outside the current actor's address space.</td> </tr> </table>	[K_EABORT]	<i>threadSemWait</i> has been aborted.	[K_EFAULT]	<i>sem</i> points outside the current actor's address space.
[K_EABORT]	<i>threadSemWait</i> has been aborted.				
[K_EFAULT]	<i>sem</i> points outside the current actor's address space.				

[K\_EINVAL] The thread semaphore structure has not been correctly initialized, or the *waitLimit* structure is not a valid *KnTimeVal*.

[K\_ETIMEOUT] The timeout occurred.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	threadSemInit, threadSemPost, threadSemWait – Initialize a thread semaphore; Signal a thread semaphore; Wait on a thread semaphore				
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadSemInit(KnThSem * sem);  int threadSemPost(KnThSem * sem);  int threadSemWait(KnThSem * sem, KnTimeVal * waitLimit);</pre>				
<b>DESCRIPTION</b>	<p>Thread semaphores are <i>KnThSem</i> structures allocated in user memory. A thread semaphore is a binary semaphore on which one, and only one, thread can wait, and which can be signalled to by multiple threads and/or interrupt handlers.</p> <p>A thread semaphore can be in two states: POSTED and UNPOSTED.</p> <p>The <i>threadSemInit</i> system call initializes the thread semaphore whose address is <i>sem</i> in the UNPOSTED state. The thread which performs <i>threadSemInit</i> on a given thread semaphore is the only thread which can wait on the semaphore (using <i>threadSemWait</i> ).</p> <p>The <i>threadSemWait</i> system call makes the current thread wait conditionally on the thread semaphore <i>sem</i>. If the thread semaphore is in the POSTED state, <i>threadSemWait</i> returns immediately and atomically changes the state of the thread semaphore to UNPOSTED. If the thread semaphore is in the UNPOSTED state, the thread is blocked according to the options described by <i>waitLimit</i> in <i>intro(2K)</i>. The <i>waitLimit</i> pointer points to a <i>KnTimeVal</i> structure containing a timeout interval as described in <i>sysTime(2K)</i>.</p> <p>The <i>threadSemPost</i> system call signals the thread semaphore <i>sem</i>. If the thread that initialized the thread semaphore is blocked behind the semaphore, <i>threadSemPost</i> awakens that thread without modifying the state of <i>sem</i>. If the thread is not blocked, <i>threadSemPost</i> sets the state of <i>sem</i> to POSTED. If the thread is not blocked and <i>sem</i> is already in the POSTED state, <i>threadSemPost</i> has no effect.</p> <p>The <i>threadSemPost</i> system call may be called within an interrupt handler, or with preemption disabled.</p>				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.				
<b>RESTRICTIONS</b>	<p>A user application and a supervisor application may not share a semaphore.</p> <p>However, two user applications may share a semaphore by mapping it in both user address spaces.</p>				
<b>ERRORS</b>	<table border="0"> <tr> <td style="padding-right: 20px;">[K_EABORT]</td> <td><i>threadSemWait</i> has been aborted.</td> </tr> <tr> <td>[K_EFAULT]</td> <td><i>sem</i> points outside the current actor's address space.</td> </tr> </table>	[K_EABORT]	<i>threadSemWait</i> has been aborted.	[K_EFAULT]	<i>sem</i> points outside the current actor's address space.
[K_EABORT]	<i>threadSemWait</i> has been aborted.				
[K_EFAULT]	<i>sem</i> points outside the current actor's address space.				

[K\_EINVAL] The thread semaphore structure has not been correctly initialized, or the *waitLimit* structure is not a valid *KnTimeVal*.

[K\_ETIMEOUT] The timeout occurred.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	threadStart, threadStop – Stop a thread; Start a thread
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadStop(KnCap * actorcap, int threadli);  int threadStart(KnCap * actorcap, int threadli);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>threadStop</i> system call stops the thread <i>threadli</i> , in the actor whose capability is given by <i>actorcap</i>.</p> <p>The thread must either be owned by the actor or have entered the actor through a safe LAP invocation (see <i>svLapCreate</i> (2K)).</p> <p>The effect of <i>threadStop</i> is to prevent the thread from running until the thread is restarted using <i>threadStart</i> (see <i>svLapCreate</i> (2K)) and becomes runnable (not suspended, and scheduled).</p> <p>The effect of <i>threadStop</i> is not instantaneous on threads executing system calls implemented via trap or cross-actor invocation, but it is guaranteed that a thread performing such a system call will not return from that call.</p> <p>The <i>threadStart</i> system call restarts the thread <i>threadli</i> , in the actor whose capability is given by <i>actorcap</i>.</p> <p>The thread must either be owned by the actor or have entered the actor through a safe LAP invocation (see <i>svLapCreate</i> (2K)).</p> <p>It may only be applied to a thread previously stopped by <i>threadStop</i> , or by <i>actorStop</i> .</p> <p>The <i>threadStop</i> and <i>threadStart</i> system calls are reserved for the usage of system or application debuggers, and should not be used inside applications.</p>
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	<p>[K_EINVAL]                    <i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.</p> <p>[K_EUNKNOWN]                <i>actorcap</i> does not specify a reachable actor.</p> <p>[K_EFAULT]                    Some of the data provided are outside the current actor's address space.</p>
<b>ATTRIBUTES</b>	See <i>attributes</i> (5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

svLapCreate(2K) , actorStop(2K)

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

<b>NAME</b>	threadStat – obtain the descriptions of the threads running in an actor
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadStat(KnCap *actorcap, unsigned int options, KnThreadStat *stat, unsigned int bufsize);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p>The <i>threadStat</i> system call obtains the descriptions of threads running in the actor whose capability is given by <i>actorcap</i>. If <i>actorcap</i> is <code>K_MYACTOR</code>, threads of the current actor are used.</p> <p>Only threads belonging to the actor or having entered it through a safe LAP invocation are returned. The system does not record information about threads entering actors through non-safe LAPs.</p> <p>The <i>options</i> argument is reserved for future use and must be set to 0.</p> <p>The <i>stat</i> argument points to a buffer in which the thread descriptions will be returned. The buffer is considered part of the caller (current user or supervisor) address space.</p> <p>The <i>bufsize</i> argument corresponds to the size of the stat buffer. It should be a multiple of the size of the <code>KnThreadStat</code> struct.</p> <p>On successful return the buffer will contain an array of descriptors, each one describing a single thread. For one thread, it is possible to retrieve several entries, each one describing a particular lap frame of the same thread executing in the actor (see <i>svLapCreate(2K)</i>).</p> <p>Each descriptor is a <i>KnThreadStat</i> structure with the following fields:</p> <pre>KnThreadLid    tsLid ; KnThreadStatus tsStatus ; int            tsFrame ; VmAddr        tsSvStack ; VmSize        tsSvStackSize ;</pre> <p>The <i>tsLid</i> field is the thread local identifier.</p> <p>The <i>tsStatus</i> status is the thread status (<code>K_INACTIVE</code> or <code>K_ACTIVE</code>, see <i>threadCreate(2K)</i> for details).</p> <p><i>tsFrame</i> is the lap frame level that appears in this actor for the considered thread (see <i>svLapCreate(2K)</i>). Level 0 means that the thread is owned by the actor</p> <p>The <i>tsSvStack</i> field is the address of the space in which the supervisor stack has been allocated. It is usually the lowest boundary of the stack area.</p> <p><i>tsSvStackSize</i> is the number of bytes allocated by the kernel for the thread's supervisor stack.</p>

**RETURN VALUE**

If successful, *threadStat* returns the number of threads running in the actor; otherwise a negative error code is returned.

As mentioned above, if several lap frames appear in the considered actor for one particular thread, all the lap frames are taken into account in the returned value.

**ERRORS**

[K\_EINVAL]                    *actorcap* is an inconsistent actor capability, or *options* specifies invalid options.

[K\_EUNKNOWN]                *actorcap* does not specify a reachable actor.

[K\_EFAULT]                    the *stat* argument is outside the caller's address space.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*threadCreate(2K)*, *svLapInvoke(2K)*

<b>NAME</b>	threadStart, threadStop – Stop a thread; Start a thread
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadStop(KnCap * actorcap, int threadli);  int threadStart(KnCap * actorcap, int threadli);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>threadStop</i> system call stops the thread <i>threadli</i> , in the actor whose capability is given by <i>actorcap</i>.</p> <p>The thread must either be owned by the actor or have entered the actor through a safe LAP invocation (see <i>svLapCreate</i> (2K)).</p> <p>The effect of <i>threadStop</i> is to prevent the thread from running until the thread is restarted using <i>threadStart</i> (see <i>svLapCreate</i> (2K)) and becomes runnable (not suspended, and scheduled).</p> <p>The effect of <i>threadStop</i> is not instantaneous on threads executing system calls implemented via trap or cross-actor invocation, but it is guaranteed that a thread performing such a system call will not return from that call.</p> <p>The <i>threadStart</i> system call restarts the thread <i>threadli</i> , in the actor whose capability is given by <i>actorcap</i>.</p> <p>The thread must either be owned by the actor or have entered the actor through a safe LAP invocation (see <i>svLapCreate</i> (2K)).</p> <p>It may only be applied to a thread previously stopped by <i>threadStop</i> , or by <i>actorStop</i> .</p> <p>The <i>threadStop</i> and <i>threadStart</i> system calls are reserved for the usage of system or application debuggers, and should not be used inside applications.</p>
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	<p>[K_EINVAL]                    <i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.</p> <p>[K_EUNKNOWN]                <i>actorcap</i> does not specify a reachable actor.</p> <p>[K_EFAULT]                    Some of the data provided are outside the current actor's address space.</p>
<b>ATTRIBUTES</b>	See <i>attributes</i> (5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

svLapCreate(2K) , actorStop(2K)

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

<b>NAME</b>	threadLoadR, threadStoreR – Get the current thread’s valid soft register value; Reset the current thread’s valid soft register value
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chSoftCtx.h&gt; long threadLoadR(void);  long threadStoreR(long reg);</pre>
<b>DESCRIPTION</b>	<p>The execution context of a thread (its various machine registers) is extended by two software registers, one for each of the possible privilege levels of a thread ( <i>USER</i> or <i>SUPERVISOR</i> ). Like hardware registers, these two registers are saved and restored at each thread switch. These registers provide the means for maintaining per-thread data efficiently.</p> <p>The <i>threadLoadR</i> system call returns the value of the current thread’s software register corresponding to the current privilege level of the thread.</p> <p>The <i>threadStoreR</i> system call sets the current thread’s software register corresponding to the current privilege level of the thread to the value specified by <i>reg</i> , and returns its new value.</p> <p>The values of these registers may also be affected using <i>threadContext</i> (2K).</p>
<b>RETURN VALUE</b>	The <i>threadLoadR</i> call returns the current register value. The <i>threadStoreR</i> call returns the new register value.
<b>EXAMPLES</b>	<p>The software register facility can be used to store the thread’s local identifier in its software register. However, this only impacts on performance, as <i>threadLoadR</i> does not imply a kernel call on most architectures, this is fully resolved in libraries which access a globally shared system context. However, <i>threadSelf</i> always implies a kernel call. In order to do this, call:</p> <pre>threadStoreR (threadSelf());</pre> <p>at the beginning of the thread’s code; then invoke:</p> <pre>threadLoadR ;</pre> <p>instead of:</p> <pre>threadSelfC ;</pre> <p>A more common way to use this facility is to associate a per-thread structure with each thread of an actor, and to store the address of this structure in the software register. This will avoid systematically transferring this address as a parameter across the layers of the actor’s code (the use of a global variable is not possible, especially on multiprocessor implementations). In this case, each actor’s thread will perform a</p>

```
threadStoreR (myArea);
```

at the beginning of the thread's code, where myArea is the address of the per thread's structure. Note that there are several ways for the thread to acquire the

```
(myArea);
```

value: it may be provided to the thread's routine as an argument (by initializing its stack correctly before creation); it may also be computed by the created thread as the result of a hash function (the hash function taking the thread's local identifier as input, and returning an entry in the array of per-thread records). An alternative to this initialization process is for the creator to call *threadContext* (2K) before resuming the thread. In each case, the per-thread record is subsequently accessed by:

```
threadLoadR ();
```

Finally, note that the software register may be used differently by different code layers: the sequence S in:

```
oldVal = threadLoadR ();
threadStoreR (newVal);
/* S */
threadStoreR (oldVal);
```

may make private use of the software register value.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*threadContext*(2K)

<b>NAME</b>	threadSuspend, threadResume – Suspend a thread; Resume a thread
<b>SYNOPSIS</b>	<pre>#include &lt;exec/chExec.h&gt; int threadSuspend(KnCap * actorcap, int threadli);  int threadResume(KnCap * actorcap, int threadli);</pre>
<b>FEATURES</b>	CORE
<b>DESCRIPTION</b>	<p><b>NOTE :</b> These calls are obsolete. Although they are maintained for compatibility reasons, they should be avoided wherever possible. See <i>threadActivate(2K)</i>, <i>threadStart(2K)</i>, and <i>threadSemInit(2K)</i> for documentation on other calls that can be used instead of <i>threadSuspend(2K)</i> and <i>threadResume(2K)</i>. The <i>threadSuspend</i> system call puts the thread whose local identifier is <i>threadli</i>, in the actor whose capability is given by <i>actorcap</i>, into the <i>SUSPENDED</i> state.</p> <p>A suspended thread may resume execution (enter the <i>ACTIVE</i> state) only if the thread is explicitly resumed using <i>threadResume(2K)</i>.</p> <p>The <i>threadSuspend</i> and <i>threadResume</i> system calls are based on a per-thread suspend counter, initialized to 0 at thread creation, incremented by <i>threadSuspend</i> and decremented by <i>threadResume</i>.</p> <p>The <i>threadResume</i> system call decrements the suspend counter of the thread whose local identifier is <i>threadli</i>, in the actor whose capability is given by <i>actorcap</i>. It only awakens the thread by setting its state to <i>ACTIVE</i>, if the suspend counter has become 0.</p> <p>If a thread which is in a <i>BLOCKED</i> state is suspended using <i>threadSuspend</i>, the thread will suspend just after it exits its <i>BLOCKED</i> state. If it is resumed using <i>threadResume</i> while still in the <i>BLOCKED</i> state, it will continue its execution normally after it exits its <i>BLOCKED</i> state (if its suspend counter has become 0). The same rule applies for a thread whose actor is in the <i>STOPPED</i> state.</p> <p>If <i>actorcap</i> is <i>K_MYACTOR</i>, the thread must be a thread of the current actor. In this case, if <i>threadli</i> is <i>K_MYSELF</i>, the current thread is used.</p>
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	<p>[K_EINVAL] <i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.</p> <p>[K_EUNKNOWN] <i>actorcap</i> does not specify a reachable actor.</p> <p>[K_EFAULT] Some of the data provided are outside the current actor's address space.</p>

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`threadCreate(2K)`

<b>NAME</b>	threadTimes – get thread execution time				
<b>SYNOPSIS</b>	<pre>#include &lt;vtimer/chVtimer.h&gt; int threadTimes(KnCap *actorcap, int threadli, KnTimeVal *internal, KnTimeVal *external);</pre>				
<b>FEATURES</b>	VTIMER				
<b>DESCRIPTION</b>	<p>The <i>threadTimes</i> system call returns the execution time of the thread whose local identifier is <i>threadli</i>, within the actor specified by <i>actorcap</i>.</p> <p>If <i>threadLi</i> is equal to <code>K_ALLACTORTHREADS</code>, <i>threadTimes</i> returns the total execution time of all the threads within the actor specified by <i>actorcap</i>.</p> <p>The <i>internal</i> time is the time spent executing in the actor in which the thread(s) was (were) originally created. The <i>external</i> time includes all execution time spent with the thread's <i>execution actor</i> set to another actor; in other words, while the thread(s) is (are) executing in a cross-actor invocation.</p> <p>The time spent in a kernel system call is added to the current execution time (internal or external) of the thread (a kernel system call isn't considered as a cross-actor invocation).</p>				
<b>RESTRICTIONS</b>	The target actor and the current actor must be located on the same site.				
<b>RETURN VALUES</b>	Upon successful completion, <i>threadTimes</i> returns <code>K_OK</code> . Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[<code>K_EINVAL</code>] <i>actorcap</i> is an inconsistent actor capability, or <i>threadli</i> is not a valid thread identifier in the specified actor.</p> <p>[<code>K_EUNKNOWN</code>] <i>actorcap</i> does not specify a reachable actor.</p> <p>[<code>K_EFAULT</code>] Some of the data provided are outside the current actor's address space.</p>				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				

**NAME** timerCreate – create a timer

**SYNOPSIS** `#include <etimer/chEtimer.h>`  
`int timerCreate(KnCap *actor, int clock, KnThreadPool *threadPool, void *cookie,`  
`int *timerLi);`

**FEATURES** TIMER

**DESCRIPTION** *timerCreate* creates a timer relative to the clock type *clock* in the actor specified by *actor*. If *actor* has the value K\_MYACTOR, the timer is created in the current actor. *timerCreate* returns the local id of the new timer in *timerLi*. The timer may then be armed using *timerSet(2K)*.

The only clock type currently supported is K\_CLOCK\_REALTIME. This clock corresponds to the system time as returned by *sysTime(2K)*.

*threadPool* must point to an initialized *KnThreadPool* object (see *timerThreadPoolInit(2K)*). At each timer expiration for the newly created timer, a thread blocked on the designated *KnThreadPool* object will be awakened so as to return from its *timerThreadPoolWait(2K)* system call.

The *cookie* value will be provided to that thread on return from *timerThreadPoolWait(2K)*. Library or application code may create and use these type of threads to execute timer notification handlers. The *cookie* value can be used to address a library object corresponding to the particular timer.

**RETURN VALUE** Upon successful completion, K\_OK is returned. Otherwise, a negative error code is returned.

**ERRORS**

[K_EINVAL]	<i>clock</i> is not equal to K_CLOCK_REALTIME, or <i>threadPool</i> has not been initialized.
[K_ENOMEM]	The maximum number of timers has been reached.
[K_EFAULT]	An address argument falls outside the caller's address space.
[K_EUNKNOWN]	<i>actor</i> does not specify a reachable actor.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *timerDelete(2K)*, *timerSet(2K)*, *timerThreadPoolInit(2K)*, *timerThreadPoolWait(2K)*

<b>NAME</b>	timerDelete – delete a timer						
<b>SYNOPSIS</b>	<pre>#include &lt;etimer/chEtimer.h&gt; int timerDelete(KnCap *actor, int timerLi);</pre>						
<b>FEATURES</b>	TIMER						
<b>DESCRIPTION</b>	The <i>timerDelete</i> system call deletes, and, if necessary, disarms the timer <i>timerLi</i> in the actor <i>actor</i> . If <i>actor</i> is equal to K_MYACTOR, the current actor is used						
<b>RETURN VALUE</b>	Upon successful completion, K_OK is returned. Otherwise, a negative error code is returned.						
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EINVAL]</td> <td><i>timerLi</i> is not a valid timer identifier within the current host actor.</td> </tr> <tr> <td>[K_EFAULT]</td> <td>An address argument falls outside the caller's address space.</td> </tr> <tr> <td>[K_EUNKNOWN]</td> <td><i>actor</i> does not specify a reachable actor.</td> </tr> </table>	[K_EINVAL]	<i>timerLi</i> is not a valid timer identifier within the current host actor.	[K_EFAULT]	An address argument falls outside the caller's address space.	[K_EUNKNOWN]	<i>actor</i> does not specify a reachable actor.
[K_EINVAL]	<i>timerLi</i> is not a valid timer identifier within the current host actor.						
[K_EFAULT]	An address argument falls outside the caller's address space.						
[K_EUNKNOWN]	<i>actor</i> does not specify a reachable actor.						
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
<b>SEE ALSO</b>	<code>timerCreate(2K)</code> , <code>timerSet(2K)</code> , <code>timerThreadPoolInit(2K)</code> , <code>timerThreadPoolWait(2K)</code>						

<b>NAME</b>	timerGetRes – get the timer resolution				
<b>SYNOPSIS</b>	<pre>#include &lt;etimer/chEtimer.h&gt; int timerGetRes(int clock, KnTimeVal *resolution);</pre>				
<b>FEATURES</b>	TIMER				
<b>DESCRIPTION</b>	<p>The <i>timerGetRes</i> system call obtains the resolution of the clock type <i>clock</i>. The value stored in <i>resolution</i> represents the smallest possible difference between two distinct time specifications related to <i>clock</i>.</p> <p>The only clock type currently supported is K_CLOCK_REALTIME. This clock corresponds to the system time as returned by <i>sysTime(2K)</i>.</p>				
<b>RETURN VALUE</b>	Upon successful completion, K_OK is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EINVAL]                    <i>clock</i> is not equal to K_CLOCK_REALTIME.</p> <p>[K_EFAULT]                    An address argument falls outside the caller's address space.</p>				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>timerCreate(2K)</i> , <i>timerSet(2K)</i>				

<b>NAME</b>	timerSet – start, cancel or query a timer				
<b>SYNOPSIS</b>	<pre>#include &lt;etimer/chEtimer.h&gt; int timerSet(KnCap *actor, int timerLi, int flags, KnTimer *newit, KnTimer *oldit);</pre>				
<b>FEATURES</b>	TIMER				
<b>DESCRIPTION</b>	<p><i>timerSet</i> starts, cancels, or queries the state of timer <i>timerLi</i> in actor <i>actor</i>. If <i>actor</i> has the value K_MYACTOR, the current actor is assumed.</p> <p>The <i>KnTimer</i> structure is defined as follows:</p> <pre>KnTimeVal ITmReload ; /* Reload value (periodic timer) */ KnTimeVal ITmValue ; /* Initial value */</pre> <p>If <i>newit</i> is non-NULL, the timer is armed accordingly. <i>ITmValue</i> contains the initial timeout value; <i>ITmReload</i> contains the subsequent interval for a periodic timer, or zero for a one-shot timer. If the K_TIMER_ABSOLUTE bit is set in <i>flags</i>, <i>ITmValue</i> is interpreted as an absolute time value (in terms of the values returned by <i>sysTime</i>). Otherwise, <i>ITmValue</i> is interpreted as a relative time interval.</p> <p>If the timer was already set, the previous setting is cancelled and replaced with the new specification. If <i>ITmValue</i> contains a zero value, the timer is disarmed.</p> <p>If <i>newit</i> is NULL, the current timer setting is unchanged.</p> <p>If <i>oldit</i> is non-NULL, the current setting of the timer (or zero if not set) is stored at the referenced location. A relative interval is always stored even if the timer was armed by absolute time.</p>				
<b>RETURN VALUE</b>	Upon successful completion, K_OK is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EINVAL] <i>timerLi</i> is not a valid timer identifier in the actor <i>actor</i>, or <i>ITmReload</i> or <i>ITmValue</i> are not valid KnTimeVal values.</p> <p>[K_EFAULT] An address argument falls outside the caller's address space.</p>				
<b>ATTRIBUTES</b>	See <a href="#">attributes(5)</a> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>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				
<b>SEE ALSO</b>	<a href="#">timerCreate(2K)</a> , <a href="#">timerDelete(2K)</a> , <a href="#">timerThreadPoolInit(2K)</a> , <a href="#">timerThreadPoolWait(2K)</a>				

<b>NAME</b>	timerThreadPoolInit – initialize a timer thread pool				
<b>SYNOPSIS</b>	<pre>#include &lt;etimer/chEtimer.h&gt; int timerThreadPoolInit(KnThreadPool *threadPool);</pre>				
<b>FEATURES</b>	TIMER				
<b>DESCRIPTION</b>	<p><i>timerThreadPoolInit</i> initializes the timer thread pool <i>threadPool</i>.</p> <p>The <i>threadPool</i> object is opaque; it must be pre-allocated by the caller but its fields are inaccessible to the caller. A <i>threadPool</i> object must be initialized before it can be accessed by a <i>timerCreate</i>.</p>				
<b>RETURN VALUE</b>	Upon successful completion, K_OK is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	[K_EFAULT]                      An address argument falls outside the caller's address space.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>timerCreate(2K)</i> , <i>timerThreadPoolWait(2K)</i>				

<b>NAME</b>	timerThreadPoolWait – wait for a timer expiration event				
<b>SYNOPSIS</b>	<pre>#include &lt;etimer/chEtimer.h&gt; int timerThreadPoolWait(KnThreadPool *threadPool, void **cookie, int *overrun, KnTimeVal *waitLimit);</pre>				
<b>FEATURES</b>	TIMER				
<b>DESCRIPTION</b>	<p><i>timerThreadPoolWait</i> waits for a subsequent timer expiration in a timer that was created using the thread pool <i>threadPool</i> (see <i>timerCreate(2K)</i>). When an expiration event occurs, and the current thread is selected among those waiting in the thread pool, it will return from <i>timerThreadPoolWait</i>. On return, <i>cookie</i> will contain the cookie value provided at the creation of the corresponding timer, and <i>overrun</i> will contain the overrun count corresponding to the current expiration.</p> <p>Until this type of timer event occurs, the current thread is blocked according to the <i>waitLimit</i> specification (see <i>intro(2K)</i>).</p> <p><i>threadPool</i> must have been previously initialized using <i>timerThreadPoolInit(2K)</i>.</p> <p>Only one thread at a time may be active on behalf of a single timer. Thus, during the interval between a thread's return from <i>timerThreadPoolWait</i> and its next invocation of <i>timerThreadPoolWait</i>, no other thread will be awakened for an expiration of the same timer; the expiration will be counted as a timer overrun.</p> <p>A timer expiration which occurs when there are no threads currently waiting on the associated thread pool is also counted as an overrun. A positive overrun count indicates that at least one timeout notification for the timer was delayed. An overrun count greater than one indicates that at least one timeout notification was lost, that is, no waiting threads could be awakened by the expiration event.</p> <p>Each time a thread returns from <i>timerThreadPoolWait</i>, the current overrun count for the corresponding timer is stored in <i>overrun</i>, and the timer's internal overrun count is reset to zero. Thus the overrun count delivered with each event notification represents only the number of delayed or lost expirations since the last notification.</p>				
<b>RETURN VALUE</b>	Upon successful completion, K_OK is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EINVAL]                      <i>threadPool</i> is not an initialized thread pool object.</p> <p>[K_EFAULT]                      An address argument falls outside the caller's address space.</p>				
<b>ATTRIBUTES</b>	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>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**

timerCreate(2K), timerThreadPoolInit(2K)

**NAME** | uiBuild, uiClear, uiEqual, uiGetSite, uiIsLocal, uiSite, uiValid – Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether a unique identifier has been cleared

**SYNOPSIS**

```
#include <ipc/chId.h>
int uiBuild(KnUniqueId * ui, unsigned int type, unsigned int site, unsigned int head, unsigned int tail);

int uiClear(KnUniqueId * Ui);

int uiEqual(KnUniqueId * ui1, KnUniqueId * ui2);

int uiGetSite(KnUniqueId * ui);

int uiIsLocal(KnUniqueId * ui);

int uiSite(KnUniqueId * ui, unsigned long site);

int uiValid(KnUniqueId * ui);
```

**FEATURES**

IPC

**DESCRIPTION**

The *uiBuild* system call builds a unique identifier from user-provided values. The built UI value is returned in the *KnUniqueId* structure pointed to by *ui*.

The *type* field defines the type of the UI as follows:

K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .
K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .
K_UISITE	The UI will be a site UI.
K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .

The *site* field is a site number. If a non-zero value is given, this field will be interpreted by the kernel as a location hint for the UI.

The *head* and *tail* fields represent the stamp of the UI. *tail* is a 32 bits value. *head* is a 13 bits value.

The *uiIsLocal* function checks whether the UI given by *ui* is recognized by the local site, and returns 1 if the UI is recognized, 0 otherwise.

On every site, the kernel declares a pre-defined site UI by default. The *uiSite* function returns the value of this pre-defined UI for the site whose site number is *site* (see *uiLocalSite* (2K)). Note that *uiSite* (*ui, site*) is equivalent to *uiBuild* (*ui, K\_UISITE, site, 0, 0*).

The *uiClear* function initializes *ui* with a standard null value.

The *uiEqual* function returns 1 if the two UI's passed as arguments are equal, 0 otherwise.

The *uiGetSite* function returns the site information contained within *ui*.

The *uiValid* function returns 0 if *ui* is equal to the null UI, 1 otherwise.

**RETURN VALUE**

The *uiIsLocal* function returns 1 if the UI checked is local, 0 otherwise. The *uiValid* function returns 1 if the checked UI is different from the null value, 0 otherwise. The *uiEqual* function returns 1 if the UI's are identical, 0 otherwise. Other calls return 0 upon successful completion, a negative error code otherwise.

The *uiClear*, *uiValid*, *uiEqual* and *uiGetSite* never return K\_EFAULT. If they are invoked with an invalid UI pointer, an exception occurs within the caller's context.

**ERRORS**

[K\_EFAULT]                      Some of the provided data are outside the current actor's address space ( *uiBuild* and *uiLocalSite* ).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*portDeclare*(2K) , *msgQueueBind*(2K) , *uiLocalSite*(2K)

**NAME** | uiBuild, uiClear, uiEqual, uiGetSite, uiIsLocal, uiSite, uiValid – Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether a unique identifier has been cleared

**SYNOPSIS**

```
#include <ipc/chId.h>
int uiBuild(KnUniqueId * ui, unsigned int type, unsigned int site, unsigned int head, unsigned int tail);

int uiClear(KnUniqueId * Ui);

int uiEqual(KnUniqueId * ui1, KnUniqueId * ui2);

int uiGetSite(KnUniqueId * ui);

int uiIsLocal(KnUniqueId * ui);

int uiSite(KnUniqueId * ui, unsigned long site);

int uiValid(KnUniqueId * ui);
```

**FEATURES**

IPC

**DESCRIPTION**

The *uiBuild* system call builds a unique identifier from user-provided values. The built UI value is returned in the *KnUniqueId* structure pointed to by *ui*.

The *type* field defines the type of the UI as follows:

K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .
K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .
K_UISITE	The UI will be a site UI.
K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .

The *site* field is a site number. If a non-zero value is given, this field will be interpreted by the kernel as a location hint for the UI.

The *head* and *tail* fields represent the stamp of the UI. *tail* is a 32 bits value. *head* is a 13 bits value.

The *uiIsLocal* function checks whether the UI given by *ui* is recognized by the local site, and returns 1 if the UI is recognized, 0 otherwise.

On every site, the kernel declares a pre-defined site UI by default. The *uiSite* function returns the value of this pre-defined UI for the site whose site number is *site* (see *uiLocalSite* (2K)). Note that *uiSite* (*ui, site*) is equivalent to *uiBuild* (*ui, K\_UISITE, site, 0, 0*).

The *uiClear* function initializes *ui* with a standard null value.

The *uiEqual* function returns 1 if the two UI's passed as arguments are equal, 0 otherwise.

The *uiGetSite* function returns the site information contained within *ui*.

The *uiValid* function returns 0 if *ui* is equal to the null UI, 1 otherwise.

**RETURN VALUE**

The *uiIsLocal* function returns 1 if the UI checked is local, 0 otherwise. The *uiValid* function returns 1 if the checked UI is different from the null value, 0 otherwise. The *uiEqual* function returns 1 if the UI's are identical, 0 otherwise. Other calls return 0 upon successful completion, a negative error code otherwise.

The *uiClear*, *uiValid*, *uiEqual* and *uiGetSite* never return K\_EFAULT. If they are invoked with an invalid UI pointer, an exception occurs within the caller's context.

**ERRORS**

[K\_EFAULT]                      Some of the provided data are outside the current actor's address space ( *uiBuild* and *uiLocalSite* ).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*portDeclare*(2K) , *msgQueueBind*(2K) , *uiLocalSite*(2K)

<b>NAME</b>	uiBuild, uiClear, uiEqual, uiGetSite, uiIsLocal, uiSite, uiValid – Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared								
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chId.h&gt; int uiBuild(KnUniqueId * ui, unsigned int type, unsigned int site, unsigned int head, unsigned int tail);  int uiClear(KnUniqueId * Ui);  int uiEqual(KnUniqueId * ui1, KnUniqueId * ui2);  int uiGetSite(KnUniqueId * ui);  int uiIsLocal(KnUniqueId * ui);  int uiSite(KnUniqueId * ui, unsigned long site);  int uiValid(KnUniqueId * ui);</pre>								
<b>FEATURES</b>	IPC								
<b>DESCRIPTION</b>	<p>The <i>uiBuild</i> system call builds a unique identifier from user-provided values. The built UI value is returned in the <i>KnUniqueId</i> structure pointed to by <i>ui</i>.</p> <p>The <i>type</i> field defines the type of the UI as follows:</p> <table border="0"> <tr> <td style="padding-right: 20px;">K_UIPORT</td> <td>The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i>.</td> </tr> <tr> <td>K_UIGROUP</td> <td>The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i>.</td> </tr> <tr> <td>K_UISITE</td> <td>The UI will be a site UI.</td> </tr> <tr> <td>K_UIMSGQUEUE</td> <td>The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i>.</td> </tr> </table> <p>The <i>site</i> field is a site number. If a non-zero value is given, this field will be interpreted by the kernel as a location hint for the UI.</p> <p>The <i>head</i> and <i>tail</i> fields represent the stamp of the UI. <i>tail</i> is a 32 bits value. <i>head</i> is a 13 bits value.</p> <p>The <i>uiIsLocal</i> function checks whether the UI given by <i>ui</i> is recognized by the local site, and returns 1 if the UI is recognized, 0 otherwise.</p>	K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .	K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .	K_UISITE	The UI will be a site UI.	K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .
K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .								
K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .								
K_UISITE	The UI will be a site UI.								
K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .								

On every site, the kernel declares a pre-defined site UI by default. The *uiSite* function returns the value of this pre-defined UI for the site whose site number is *site* (see *uiLocalSite* (2K)). Note that *uiSite* (*ui, site*) is equivalent to *uiBuild* (*ui, K\_UISITE, site, 0, 0*).

The *uiClear* function initializes *ui* with a standard null value.

The *uiEqual* function returns 1 if the two UI's passed as arguments are equal, 0 otherwise.

The *uiGetSite* function returns the site information contained within *ui*.

The *uiValid* function returns 0 if *ui* is equal to the null UI, 1 otherwise.

**RETURN VALUE**

The *uiIsLocal* function returns 1 if the UI checked is local, 0 otherwise. The *uiValid* function returns 1 if the checked UI is different from the null value, 0 otherwise. The *uiEqual* function returns 1 if the UI's are identical, 0 otherwise. Other calls return 0 upon successful completion, a negative error code otherwise.

The *uiClear*, *uiValid*, *uiEqual* and *uiGetSite* never return K\_EFAULT. If they are invoked with an invalid UI pointer, an exception occurs within the caller's context.

**ERRORS**

[K\_EFAULT]                      Some of the provided data are outside the current actor's address space ( *uiBuild* and *uiLocalSite* ).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*portDeclare*(2K) , *msgQueueBind*(2K) , *uiLocalSite*(2K)

**NAME** | uiBuild, uiClear, uiEqual, uiGetSite, uiIsLocal, uiSite, uiValid – Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether a unique identifier has been cleared

**SYNOPSIS**

```
#include <ipc/chId.h>
int uiBuild(KnUniqueId * ui, unsigned int type, unsigned int site, unsigned int head, unsigned int tail);

int uiClear(KnUniqueId * Ui);

int uiEqual(KnUniqueId * ui1, KnUniqueId * ui2);

int uiGetSite(KnUniqueId * ui);

int uiIsLocal(KnUniqueId * ui);

int uiSite(KnUniqueId * ui, unsigned long site);

int uiValid(KnUniqueId * ui);
```

**FEATURES**

IPC

**DESCRIPTION**

The *uiBuild* system call builds a unique identifier from user-provided values. The built UI value is returned in the *KnUniqueId* structure pointed to by *ui*.

The *type* field defines the type of the UI as follows:

K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .
K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .
K_UISITE	The UI will be a site UI.
K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .

The *site* field is a site number. If a non-zero value is given, this field will be interpreted by the kernel as a location hint for the UI.

The *head* and *tail* fields represent the stamp of the UI. *tail* is a 32 bits value. *head* is a 13 bits value.

The *uiIsLocal* function checks whether the UI given by *ui* is recognized by the local site, and returns 1 if the UI is recognized, 0 otherwise.

On every site, the kernel declares a pre-defined site UI by default. The *uiSite* function returns the value of this pre-defined UI for the site whose site number is *site* (see *uiLocalSite* (2K)). Note that *uiSite* (*ui*, *site*) is equivalent to *uiBuild* (*ui*, *K\_UISITE*, *site*, 0, 0).

The *uiClear* function initializes *ui* with a standard null value.

The *uiEqual* function returns 1 if the two UI's passed as arguments are equal, 0 otherwise.

The *uiGetSite* function returns the site information contained within *ui*.

The *uiValid* function returns 0 if *ui* is equal to the null UI, 1 otherwise.

**RETURN VALUE**

The *uiIsLocal* function returns 1 if the UI checked is local, 0 otherwise. The *uiValid* function returns 1 if the checked UI is different from the null value, 0 otherwise. The *uiEqual* function returns 1 if the UI's are identical, 0 otherwise. Other calls return 0 upon successful completion, a negative error code otherwise.

The *uiClear*, *uiValid*, *uiEqual* and *uiGetSite* never return *K\_EFAULT*. If they are invoked with an invalid UI pointer, an exception occurs within the caller's context.

**ERRORS**

[*K\_EFAULT*] Some of the provided data are outside the current actor's address space ( *uiBuild* and *uiLocalSite* ).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*portDeclare*(2K) , *msgQueueBind*(2K) , *uiLocalSite*(2K)

<b>NAME</b>	uiBuild, uiClear, uiEqual, uiGetSite, uiIsLocal, uiSite, uiValid – Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether a unique identifier has been cleared								
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chId.h&gt; int uiBuild(KnUniqueId * ui, unsigned int type, unsigned int site, unsigned int head, unsigned int tail);  int uiClear(KnUniqueId * Ui);  int uiEqual(KnUniqueId * ui1, KnUniqueId * ui2);  int uiGetSite(KnUniqueId * ui);  int uiIsLocal(KnUniqueId * ui);  int uiSite(KnUniqueId * ui, unsigned long site);  int uiValid(KnUniqueId * ui);</pre>								
<b>FEATURES</b>	IPC								
<b>DESCRIPTION</b>	<p>The <i>uiBuild</i> system call builds a unique identifier from user-provided values. The built UI value is returned in the <i>KnUniqueId</i> structure pointed to by <i>ui</i>.</p> <p>The <i>type</i> field defines the type of the UI as follows:</p> <table border="0"> <tr> <td style="padding-right: 20px;">K_UIPORT</td> <td>The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i>.</td> </tr> <tr> <td>K_UIGROUP</td> <td>The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i>.</td> </tr> <tr> <td>K_UISITE</td> <td>The UI will be a site UI.</td> </tr> <tr> <td>K_UIMSGQUEUE</td> <td>The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i>.</td> </tr> </table> <p>The <i>site</i> field is a site number. If a non-zero value is given, this field will be interpreted by the kernel as a location hint for the UI.</p> <p>The <i>head</i> and <i>tail</i> fields represent the stamp of the UI. <i>tail</i> is a 32 bits value. <i>head</i> is a 13 bits value.</p> <p>The <i>uiIsLocal</i> function checks whether the UI given by <i>ui</i> is recognized by the local site, and returns 1 if the UI is recognized, 0 otherwise.</p>	K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .	K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .	K_UISITE	The UI will be a site UI.	K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .
K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .								
K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .								
K_UISITE	The UI will be a site UI.								
K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .								

On every site, the kernel declares a pre-defined site UI by default. The *uiSite* function returns the value of this pre-defined UI for the site whose site number is *site* (see *uiLocalSite* (2K)). Note that *uiSite* (*ui, site*) is equivalent to *uiBuild* (*ui, K\_UISITE, site, 0, 0*).

The *uiClear* function initializes *ui* with a standard null value.

The *uiEqual* function returns 1 if the two UI's passed as arguments are equal, 0 otherwise.

The *uiGetSite* function returns the site information contained within *ui*.

The *uiValid* function returns 0 if *ui* is equal to the null UI, 1 otherwise.

**RETURN VALUE**

The *uiIsLocal* function returns 1 if the UI checked is local, 0 otherwise. The *uiValid* function returns 1 if the checked UI is different from the null value, 0 otherwise. The *uiEqual* function returns 1 if the UI's are identical, 0 otherwise. Other calls return 0 upon successful completion, a negative error code otherwise.

The *uiClear*, *uiValid*, *uiEqual* and *uiGetSite* never return K\_EFAULT. If they are invoked with an invalid UI pointer, an exception occurs within the caller's context.

**ERRORS**

[K\_EFAULT]                      Some of the provided data are outside the current actor's address space ( *uiBuild* and *uiLocalSite* ).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*portDeclare*(2K) , *msgQueueBind*(2K) , *uiLocalSite*(2K)

<b>NAME</b>	uiLocalSite – get the local site number
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chId.h&gt; int uiLocalSite(int *site);</pre>
<b>DESCRIPTION</b>	The <i>uiLocalSite</i> system call returns the local site number to the value pointed to by <i>site</i> .
<b>NOTES</b>	A site number is a 32 bit unsigned integer, but <i>uiLocalSite</i> returns the value of the local site number as a signed integer.
<b>RETURN VALUES</b>	Upon successful completion K_OK is returned. Otherwise, a negative error code is returned.
<b>ERRORS</b>	[K_EFAULT]                      Some of the data provided are outside the current actor's address space.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	uiBuild, uiClear, uiEqual, uiGetSite, uiIsLocal, uiSite, uiValid – Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared								
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chId.h&gt; int uiBuild(KnUniqueId * ui, unsigned int type, unsigned int site, unsigned int head, unsigned int tail);  int uiClear(KnUniqueId * Ui);  int uiEqual(KnUniqueId * ui1, KnUniqueId * ui2);  int uiGetSite(KnUniqueId * ui);  int uiIsLocal(KnUniqueId * ui);  int uiSite(KnUniqueId * ui, unsigned long site);  int uiValid(KnUniqueId * ui);</pre>								
<b>FEATURES</b>	IPC								
<b>DESCRIPTION</b>	<p>The <i>uiBuild</i> system call builds a unique identifier from user-provided values. The built UI value is returned in the <i>KnUniqueId</i> structure pointed to by <i>ui</i>.</p> <p>The <i>type</i> field defines the type of the UI as follows:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">K_UIPORT</td> <td>The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i>.</td> </tr> <tr> <td>K_UIGROUP</td> <td>The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i>.</td> </tr> <tr> <td>K_UISITE</td> <td>The UI will be a site UI.</td> </tr> <tr> <td>K_UIMSGQUEUE</td> <td>The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i>.</td> </tr> </table> <p>The <i>site</i> field is a site number. If a non-zero value is given, this field will be interpreted by the kernel as a location hint for the UI.</p> <p>The <i>head</i> and <i>tail</i> fields represent the stamp of the UI. <i>tail</i> is a 32 bits value. <i>head</i> is a 13 bits value.</p> <p>The <i>uiIsLocal</i> function checks whether the UI given by <i>ui</i> is recognized by the local site, and returns 1 if the UI is recognized, 0 otherwise.</p>	K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .	K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .	K_UISITE	The UI will be a site UI.	K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .
K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .								
K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .								
K_UISITE	The UI will be a site UI.								
K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .								

On every site, the kernel declares a pre-defined site UI by default. The *uiSite* function returns the value of this pre-defined UI for the site whose site number is *site* (see *uiLocalSite* (2K)). Note that *uiSite* (*ui*, *site*) is equivalent to *uiBuild* (*ui*, *K\_UISITE*, *site*, 0, 0).

The *uiClear* function initializes *ui* with a standard null value.

The *uiEqual* function returns 1 if the two UI's passed as arguments are equal, 0 otherwise.

The *uiGetSite* function returns the site information contained within *ui*.

The *uiValid* function returns 0 if *ui* is equal to the null UI, 1 otherwise.

**RETURN VALUE**

The *uiIsLocal* function returns 1 if the UI checked is local, 0 otherwise. The *uiValid* function returns 1 if the checked UI is different from the null value, 0 otherwise. The *uiEqual* function returns 1 if the UI's are identical, 0 otherwise. Other calls return 0 upon successful completion, a negative error code otherwise.

The *uiClear*, *uiValid*, *uiEqual* and *uiGetSite* never return *K\_EFAULT*. If they are invoked with an invalid UI pointer, an exception occurs within the caller's context.

**ERRORS**

[*K\_EFAULT*] Some of the provided data are outside the current actor's address space ( *uiBuild* and *uiLocalSite* ).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*portDeclare*(2K) , *msgQueueBind*(2K) , *uiLocalSite*(2K)

<b>NAME</b>	uiBuild, uiClear, uiEqual, uiGetSite, uiIsLocal, uiSite, uiValid – Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared								
<b>SYNOPSIS</b>	<pre>#include &lt;ipc/chId.h&gt; int uiBuild(KnUniqueId * ui, unsigned int type, unsigned int site, unsigned int head, unsigned int tail);  int uiClear(KnUniqueId * Ui);  int uiEqual(KnUniqueId * ui1, KnUniqueId * ui2);  int uiGetSite(KnUniqueId * ui);  int uiIsLocal(KnUniqueId * ui);  int uiSite(KnUniqueId * ui, unsigned long site);  int uiValid(KnUniqueId * ui);</pre>								
<b>FEATURES</b>	IPC								
<b>DESCRIPTION</b>	<p>The <i>uiBuild</i> system call builds a unique identifier from user-provided values. The built UI value is returned in the <i>KnUniqueId</i> structure pointed to by <i>ui</i>.</p> <p>The <i>type</i> field defines the type of the UI as follows:</p> <table border="0"> <tr> <td style="padding-right: 20px;">K_UIPORT</td> <td>The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i>.</td> </tr> <tr> <td>K_UIGROUP</td> <td>The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i>.</td> </tr> <tr> <td>K_UISITE</td> <td>The UI will be a site UI.</td> </tr> <tr> <td>K_UIMSGQUEUE</td> <td>The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i>.</td> </tr> </table> <p>The <i>site</i> field is a site number. If a non-zero value is given, this field will be interpreted by the kernel as a location hint for the UI.</p> <p>The <i>head</i> and <i>tail</i> fields represent the stamp of the UI. <i>tail</i> is a 32 bits value. <i>head</i> is a 13 bits value.</p> <p>The <i>uiIsLocal</i> function checks whether the UI given by <i>ui</i> is recognized by the local site, and returns 1 if the UI is recognized, 0 otherwise.</p>	K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .	K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .	K_UISITE	The UI will be a site UI.	K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .
K_UIPORT	The UI will be usable as a port UI for the IPC module. The UI can then be used in a call to <i>portDeclare(2K)</i> .								
K_UIGROUP	The UI will be usable as a group UI for the IPC module. The UI can then be used in a call to <i>grpPortInsert(2K)</i> .								
K_UISITE	The UI will be a site UI.								
K_UIMSGQUEUE	The UI will be usable as a message queue UI for the MIPC_G module. The UI can then be used in a call to <i>msgQueueBind(2K)</i> .								

On every site, the kernel declares a pre-defined site UI by default. The *uiSite* function returns the value of this pre-defined UI for the site whose site number is *site* (see *uiLocalSite* (2K)). Note that *uiSite* (*ui*, *site*) is equivalent to *uiBuild* (*ui*, *K\_UISITE*, *site*, 0, 0).

The *uiClear* function initializes *ui* with a standard null value.

The *uiEqual* function returns 1 if the two UI's passed as arguments are equal, 0 otherwise.

The *uiGetSite* function returns the site information contained within *ui*.

The *uiValid* function returns 0 if *ui* is equal to the null UI, 1 otherwise.

**RETURN VALUE**

The *uiIsLocal* function returns 1 if the UI checked is local, 0 otherwise. The *uiValid* function returns 1 if the checked UI is different from the null value, 0 otherwise. The *uiEqual* function returns 1 if the UI's are identical, 0 otherwise. Other calls return 0 upon successful completion, a negative error code otherwise.

The *uiClear*, *uiValid*, *uiEqual* and *uiGetSite* never return *K\_EFAULT*. If they are invoked with an invalid UI pointer, an exception occurs within the caller's context.

**ERRORS**

[*K\_EFAULT*] Some of the provided data are outside the current actor's address space ( *uiBuild* and *uiLocalSite* ).

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*portDeclare*(2K) , *msgQueueBind*(2K) , *uiLocalSite*(2K)

<b>NAME</b>	univTime, univTimeSet, univTimeAdjust – Get time-of-day; Set time-of-day; Adjust time-of-day univTimeGetRes; Get time-of-day resolution
<b>SYNOPSIS</b>	<pre>#include &lt;date/chDate.h&gt; int univTime(KnTimeVal * curtime);  int univTimeSet(KnTimeVal * oldtime, KnTimeVal * newtime);  int univTimeAdjust(KnTimeVal * olddelta, KnTimeVal * newdelta);  int univTimeGetRes(KnTimeVal * resolution);</pre>
<b>FEATURES</b>	DATE
<b>DESCRIPTION</b>	<p>The kernel maintains a Universal Time (time-of-day) variable using a <i>KnTimeVal</i> structure whose members are defined in <i>sysTime(2K)</i>.</p> <p>If not initialized from a hardware time-of-day chip during system initialization, this variable will hold an invalid value until <i>univTimeSet</i> is called.</p> <p>The <i>univTime</i> function returns the time of day in Universal Time to <i>curtime</i>. If the time of day is not yet set, <i>curtime</i> will be set to an invalid value (<i>tmSec = 0, tmNSec = K_INVALID_NSEC</i>).</p> <p>The <i>univTimeSet</i> and <i>univTimeAdjust</i> functions allow you to get and/or set the kernel time-of-day value, or to request a progressive adjustment of that value.</p> <p>The <i>univTimeGetRes</i> function obtains the resolution of the Universal Time system calls. The time value returned in <i>resolution</i> represents the smallest possible difference between two distinct values of the universal time.</p> <p>All the <i>univTime</i> arguments are pointers to <i>KnTimeVal</i> structures. The input arguments are <i>newtime</i> and <i>newdelta</i>. The output arguments are <i>curtime</i>, <i>oldtime</i>, <i>olddelta</i> and <i>resolution</i>. If <i>newtime</i> is not 0, the kernel variable is assigned with the contents of the structure pointed to by <i>newtime</i>. If <i>oldtime</i> is not 0, the current value of the kernel variable is given in the structure pointed to by <i>oldtime</i> (if <i>newtime</i> is not 0, the returned value corresponds to the value of the kernel variable before its modification). If <i>newdelta</i> is not 0, the kernel will add (if the structure pointed to by <i>newdelta</i> represents a positive value) or subtract (if <i>newdelta</i> represents a negative value) the number of (seconds, nanoseconds) represented by <i>newdelta</i>. This adjustment is done progressively: a small portion of the adjustment which is a fraction of the system clock resolution is performed at each clock interrupt until the full adjustment is made. The time value always increases. If <i>olddelta</i> is not 0, the value of the remaining adjustment is given in the structure pointed to by <i>olddelta</i>.</p> <p>Only <i>SUPERVISOR</i> threads (see <i>threadCreate (2K)</i>) or threads belonging to <i>SYSTEM</i> actors (see <i>actorCreate (2K)</i>) are allowed to set the time-of-day or to request an adjustment.</p>

**RETURN VALUE** Upon successful completion, the Universal Time system calls return zero. Otherwise, a negative error code is returned.

**ERRORS**

[K_EFAULT]	Some of the data provided are outside the current actor's address space.
[K_EINVAL]	The time value referenced by <i>newtime</i> or <i>newdelta</i> is invalid.
[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `sysTime(2K)`

<b>NAME</b>	univTime, univTimeSet, univTimeAdjust – Get time-of-day; Set time-of-day; Adjust time-of-day univTimeGetRes; Get time-of-day resolution
<b>SYNOPSIS</b>	<pre>#include &lt;date/chDate.h&gt; int univTime(KnTimeVal * curtime);  int univTimeSet(KnTimeVal * oldtime, KnTimeVal * newtime);  int univTimeAdjust(KnTimeVal * olddelta, KnTimeVal * newdelta);  int univTimeGetRes(KnTimeVal * resolution);</pre>
<b>FEATURES</b>	DATE
<b>DESCRIPTION</b>	<p>The kernel maintains a Universal Time (time-of-day) variable using a <i>KnTimeVal</i> structure whose members are defined in <i>sysTime(2K)</i>.</p> <p>If not initialized from a hardware time-of-day chip during system initialization, this variable will hold an invalid value until <i>univTimeSet</i> is called.</p> <p>The <i>univTime</i> function returns the time of day in Universal Time to <i>curtime</i>. If the time of day is not yet set, <i>curtime</i> will be set to an invalid value (<i>tmSec = 0, tmNSec = K_INVALID_NSEC</i>).</p> <p>The <i>univTimeSet</i> and <i>univTimeAdjust</i> functions allow you to get and/or set the kernel time-of-day value, or to request a progressive adjustment of that value.</p> <p>The <i>univTimeGetRes</i> function obtains the resolution of the Universal Time system calls. The time value returned in <i>resolution</i> represents the smallest possible difference between two distinct values of the universal time.</p> <p>All the <i>univTime</i> arguments are pointers to <i>KnTimeVal</i> structures. The input arguments are <i>newtime</i> and <i>newdelta</i>. The output arguments are <i>curtime</i>, <i>oldtime</i>, <i>olddelta</i> and <i>resolution</i>. If <i>newtime</i> is not 0, the kernel variable is assigned with the contents of the structure pointed to by <i>newtime</i>. If <i>oldtime</i> is not 0, the current value of the kernel variable is given in the structure pointed to by <i>oldtime</i> (if <i>newtime</i> is not 0, the returned value corresponds to the value of the kernel variable before its modification). If <i>newdelta</i> is not 0, the kernel will add (if the structure pointed to by <i>newdelta</i> represents a positive value) or subtract (if <i>newdelta</i> represents a negative value) the number of (seconds, nanoseconds) represented by <i>newdelta</i>. This adjustment is done progressively: a small portion of the adjustment which is a fraction of the system clock resolution is performed at each clock interrupt until the full adjustment is made. The time value always increases. If <i>olddelta</i> is not 0, the value of the remaining adjustment is given in the structure pointed to by <i>olddelta</i>.</p> <p>Only <i>SUPERVISOR</i> threads (see <i>threadCreate(2K)</i>) or threads belonging to <i>SYSTEM</i> actors (see <i>actorCreate(2K)</i>) are allowed to set the time-of-day or to request an adjustment.</p>

**RETURN VALUE** Upon successful completion, the Universal Time system calls return zero. Otherwise, a negative error code is returned.

**ERRORS**

[K_EFAULT]	Some of the data provided are outside the current actor's address space.
[K_EINVAL]	The time value referenced by <i>newtime</i> or <i>newdelta</i> is invalid.
[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `sysTime(2K)`

<b>NAME</b>	univTime, univTimeSet, univTimeAdjust – Get time-of-day; Set time-of-day; Adjust time-of-day univTimeGetRes; Get time-of-day resolution
<b>SYNOPSIS</b>	<pre>#include &lt;date/chDate.h&gt; int univTime(KnTimeVal * curtime);  int univTimeSet(KnTimeVal * oldtime, KnTimeVal * newtime);  int univTimeAdjust(KnTimeVal * olddelta, KnTimeVal * newdelta);  int univTimeGetRes(KnTimeVal * resolution);</pre>
<b>FEATURES</b>	DATE
<b>DESCRIPTION</b>	<p>The kernel maintains a Universal Time (time-of-day) variable using a <i>KnTimeVal</i> structure whose members are defined in <i>sysTime(2K)</i>.</p> <p>If not initialized from a hardware time-of-day chip during system initialization, this variable will hold an invalid value until <i>univTimeSet</i> is called.</p> <p>The <i>univTime</i> function returns the time of day in Universal Time to <i>curtime</i>. If the time of day is not yet set, <i>curtime</i> will be set to an invalid value (<i>tmSec = 0, tmNSec = K_INVALID_NSEC</i>).</p> <p>The <i>univTimeSet</i> and <i>univTimeAdjust</i> functions allow you to get and/or set the kernel time-of-day value, or to request a progressive adjustment of that value.</p> <p>The <i>univTimeGetRes</i> function obtains the resolution of the Universal Time system calls. The time value returned in <i>resolution</i> represents the smallest possible difference between two distinct values of the universal time.</p> <p>All the <i>univTime</i> arguments are pointers to <i>KnTimeVal</i> structures. The input arguments are <i>newtime</i> and <i>newdelta</i>. The output arguments are <i>curtime</i>, <i>oldtime</i>, <i>olddelta</i> and <i>resolution</i>. If <i>newtime</i> is not 0, the kernel variable is assigned with the contents of the structure pointed to by <i>newtime</i>. If <i>oldtime</i> is not 0, the current value of the kernel variable is given in the structure pointed to by <i>oldtime</i> (if <i>newtime</i> is not 0, the returned value corresponds to the value of the kernel variable before its modification). If <i>newdelta</i> is not 0, the kernel will add (if the structure pointed to by <i>newdelta</i> represents a positive value) or subtract (if <i>newdelta</i> represents a negative value) the number of (seconds, nanoseconds) represented by <i>newdelta</i>. This adjustment is done progressively: a small portion of the adjustment which is a fraction of the system clock resolution is performed at each clock interrupt until the full adjustment is made. The time value always increases. If <i>olddelta</i> is not 0, the value of the remaining adjustment is given in the structure pointed to by <i>olddelta</i>.</p> <p>Only <i>SUPERVISOR</i> threads (see <i>threadCreate (2K)</i>) or threads belonging to <i>SYSTEM</i> actors (see <i>actorCreate (2K)</i>) are allowed to set the time-of-day or to request an adjustment.</p>

**RETURN VALUE** Upon successful completion, the Universal Time system calls return zero. Otherwise, a negative error code is returned.

**ERRORS**

[K_EFAULT]	Some of the data provided are outside the current actor's address space.
[K_EINVAL]	The time value referenced by <i>newtime</i> or <i>newdelta</i> is invalid.
[K_EPRIV]	The current thread is neither a supervisor thread nor a thread of a system actor.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `sysTime(2K)`

<b>NAME</b>	virtualTimeGetRes – get virtual time resolution				
<b>SYNOPSIS</b>	#include <vtimer/chVtimer.h> int virtualTimeGetRes(KnTimeVal *resolution);				
<b>FEATURES</b>	VTIMER				
<b>DESCRIPTION</b>	The <i>virtualTimeGetRes</i> function obtains the resolution of virtual time specifications and returned values in the <i>threadTimes(2K)</i> and <i>svVirtualTimeoutSet(2K)</i> . The time value returned in <i>resolution</i> represents the smallest possible difference between the two distinct virtual time values.				
<b>RETURN VALUE</b>	Upon successful completion a value of 0 is returned. Otherwise, a negative error code is returned.				
<b>ERRORS</b>	[K_EFAULT] <i>resolution</i> contains an invalid memory address.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>svVirtualTimeoutSet(2K)</i> , <i>threadTimes(2K)</i>				

<b>NAME</b>	vmCopy – copy data between actor address spaces										
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int vmCopy(KnCap *srcactorcap, VmAddr srcaddress, KnCap *dstactorcap, VmAddr dstaddress, VmSize size, VmFlags flags);</pre>										
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL										
<b>DESCRIPTION</b>	<p>The <i>vmCopy</i> function copies data from a source actor address space to a target actor address space. The source data is specified by <i>srcactorcap</i> - the source actor capability, <i>srcaddress</i> - the data start address in the source actor address space and <i>size</i> - the data size. The target is described by <i>dstactorcap</i> - the target actor capability and <i>dstaddress</i> - the start address in the target actor address space.</p> <p>If <i>srcactorcap</i> and/or <i>dstactorcap</i> is K_MYACTOR, the current actor is used. If <i>srcactorcap</i> and/or <i>dstactorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The <i>flags</i> argument has to be zero and will be used for future interface extensions.</p>										
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space.</td> </tr> <tr> <td>[K_EINVAL]</td> <td>An inconsistent actor capability.</td> </tr> <tr> <td>[K_EUNKNOWN]</td> <td><i>tgtactorcap</i> or <i>srcactorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td>[K_EFAULT]</td> <td>Some or all of the addresses in the source or the target address range are invalid.</td> </tr> <tr> <td>[K_EFAIL]</td> <td>A remote transaction has failed during a remote <i>vmCopy</i>.</td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.	[K_EINVAL]	An inconsistent actor capability.	[K_EUNKNOWN]	<i>tgtactorcap</i> or <i>srcactorcap</i> does not specify a reachable actor.	[K_EFAULT]	Some or all of the addresses in the source or the target address range are invalid.	[K_EFAIL]	A remote transaction has failed during a remote <i>vmCopy</i> .
[K_EFAULT]	Some of the arguments provided are outside the caller's address space.										
[K_EINVAL]	An inconsistent actor capability.										
[K_EUNKNOWN]	<i>tgtactorcap</i> or <i>srcactorcap</i> does not specify a reachable actor.										
[K_EFAULT]	Some or all of the addresses in the source or the target address range are invalid.										
[K_EFAIL]	A remote transaction has failed during a remote <i>vmCopy</i> .										
<b>ATTRIBUTES</b>	<p>See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving						
ATTRIBUTE TYPE	ATTRIBUTE VALUE										
Interface Stability	Evolving										

**NAME** vmFree – free memory

**SYNOPSIS** #include <mem/chMem.h>  
int vmFree(KnCap \*actorcap, VmAddr address, VmSize size);

**FEATURES** MEM\_VIRTUAL

**DESCRIPTION** The *vmFree* function frees memory in an actor address space from *address* to *address + size - 1*. The actor is specified by its capability *actorcap*. On—demand physical memory corresponding to the specified address space range is deallocated without region destruction. The *address* and *size* fields must be page-aligned.

**RETURN VALUE** If successful, K\_OK is returned, otherwise a negative error code is returned.

**ERRORS**

- [K\_EFAULT] Some of the arguments provided are outside the caller’s address space.
- [K\_EINVAL] An inconsistent actor capability was provided.
- [K\_EUNKNOWN] *actorcap* does not specify a reachable actor.
- [K\_EADDR] Some or all the addresses from the target address range are invalid .
- [K\_EROUND] *address* isn’t page aligned.
- [K\_EROUND] *size* isn’t page-aligned.
- [K\_EBUSY] Tried to free no—demand (mapped to a region with the K\_NODEMAND attribute) physical memory.

**RESTRICTIONS** The target actor and the current actor must be located on the same site.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	vmLock, vmUnLock – Fix data in memory; Data in memory								
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int vmLock(KnCap * actorcap, VmAddr address, VmSize size, VmFlags flags); int vmUnLock(KnCap * actorcap, VmAddr address, VmSize size);</pre>								
<b>FEATURES</b>	MEM_PROTECTED, MEM_VIRTUAL								
<b>DESCRIPTION</b>	<p>The <i>vmLock</i> function locks the translations of the target virtual address range of the target actor address space. The (<i>vmUnLock</i>) function unlocks the translations of the target virtual address range of the target actor address space.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>actorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The target address range is specified using its starting <i>address</i> and its <i>size</i>. Both <i>address</i> and <i>size</i> must be page aligned.</p> <p>The <i>vmLock</i> function causes each virtual address from the target address range to be mapped to physical memory with the required protections. If the K_WRITABLE flag is set in the <i>flags</i> argument, write access is required; otherwise, read access is required. Then <i>vmLock</i> gets a reader lock on the translations from the target address range. The reader lock guarantees that when a thread tries to modify either the physical address or the protections of a locked mapping (using <i>rgnFree</i>, <i>rgnSetProt</i> or <i>lcFlush</i>, for example), it blocks until the corresponding <i>vmUnLock</i> is encountered.</p> <p>If the K_WRITABLE flag is set in the <i>flags</i> argument, <i>vmLock</i> marks the locked pages as modified. This could be necessary if, for example after locking, the pages are modified by a driver using the corresponding physical addresses.</p> <p>An address range locked by a number of non-overlapped <i>vmLock</i> operations may be unlocked by any number of non-overlapped <i>vmUnLock</i> operations.</p>								
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.								
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td>An inconsistent actor capability has been provided.</td> </tr> <tr> <td style="vertical-align: top;">[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td style="vertical-align: top;">[K_EADDR]</td> <td>Some or all addresses from the target address range are invalid.</td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.	[K_EINVAL]	An inconsistent actor capability has been provided.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.	[K_EADDR]	Some or all addresses from the target address range are invalid.
[K_EFAULT]	Some of the arguments provided are outside the caller's address space.								
[K_EINVAL]	An inconsistent actor capability has been provided.								
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.								
[K_EADDR]	Some or all addresses from the target address range are invalid.								

[K\_EPROT] Write access is required but some or all the addresses from the target address range are write—p roTECTED.

[K\_EROUND] *address* isn't page aligned.

[K\_EROUND] *size* isn't page-aligned.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`rgnFree(2K)` , `lcFlush(2K)`

<b>NAME</b>	vmPageSize – get the minimum allocatable memory block size				
<b>SYNOPSIS</b>	#include <mem/chMem.h> int vmPageSize(void);				
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL				
<b>DESCRIPTION</b>	The <i>vmPageSize</i> system call returns the minimum allocatable memory block size. In case of virtual memory management, this size corresponds to the virtual page size.				
<b>RETURN VALUE</b>	The minimum allocatable memory block size, in bytes.				
<b>ATTRIBUTES</b>	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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	svPagesAllocate(2K)				

**NAME** vmPhysAddr – get the physical address corresponding to a virtual address

**SYNOPSIS** #include <mem/chMem.h>  
 int **vmPhysAddr**(KnCap \*actorcap, VmAddr virtaddress, PhAddr \*physaddr);

**FEATURES** MEM\_PROTECTED, MEM\_VIRTUAL

**DESCRIPTION** The *vmPhysAddr* system call returns to *physaddr* the physical address mapped to the virtual address *virtaddress* in the target actor address space.

The target actor is specified by *actorcap* - a pointer to the actor capability. If *actorcap* is K\_MYACTOR, the address space of the current actor is used. If *actorcap* is K\_SVACTOR, the supervisor address space is used.

The caller must insure that a physical address was effectively mapped to the target virtual address before invoking *vmPhysAddr*, and that the same physical address will always be mapped to the target virtual address during the invocation.

The *vmPhysAddr* call can be called from interrupt level.

**RETURN VALUE** If successful K\_OK is returned, otherwise a negative error code is returned.

**ERRORS** [K\_EFAULT] Some of the arguments provided are outside the caller's address space.  
 [K\_EINVAL] An inconsistent actor capability was provided.  
 [K\_EUNKNOWN] *actorcap* does not specify a reachable actor.

**RESTRICTIONS** The target actor and the current actor must be located on the same site.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

<b>NAME</b>	vmSetPar – set the memory management parameters
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int vmSetPar(KnVmPar *par);</pre>
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL
<b>DESCRIPTION</b>	<p><i>vmSetPar</i> sets the virtual memory management parameters on the current site.</p> <p>The <code>KnVmPar</code> structure describes the site memory management parameters:</p> <pre>KnUniqueId  defaultMapper ; VmSize      redMark ; VmSize      orangeMark ; VmSize      greenMark ; VmAddr      usrStartAddr ; VmAddr      usrEndAddr ; VmAddr      svStartAddr ; VmAddr      svEndAddr ; VmOffset    minOffset ; VmOffset    maxOffset ;</pre> <p><code>defaultMapper</code> field specifies the default mapper UI.</p> <p>When the size of free memory on the site goes over the <code>orangeMark</code> the memory management system starts the swap out daemon(s).</p> <p>When the size of free memory on the site comes back to the <code>greenMark</code>, the swap out daemons stop.</p> <p>When the size of free memory on the site goes over the <code>redMark</code> the memory management system tries to stop all threads not involved in swap out activity.</p> <p><code>usrStartAddr</code> and <code>usrEndAddr</code> specify the boundary of user address space.</p> <p><code>svStartAddr</code> and <code>svEndAddr</code> specify the boundary of the system address space.</p> <p><code>minOffset</code> and <code>maxOffset</code> specifies the valid range of segment offsets.</p>
<b>RETURN VALUE</b>	If successful <code>K_OK</code> is returned, otherwise a negative error code is returned.
<b>ERRORS</b>	<p>[<code>K_EINVAL</code>] The values of marks are inconsistent.</p> <p>[<code>K_EFAULT</code>] Some of the provided arguments are outside the caller's address space.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

vmStat(2K)

<b>NAME</b>	vmStat – get memory management statistics				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int vmStat(KnVmStat *stat);</pre>				
<b>FEATURES</b>	MEM_FLAT, MEM_PROTECTED, MEM_VIRTUAL				
<b>DESCRIPTION</b>	<p>The <i>vmStat</i> system call gets memory management statistics on the current site.</p> <p>The <i>KnVmStat</i> contains the following fields:</p> <pre>KnVmPar par ; VmSize hostMem ; VmSize freeMem ; VmSize lockMem ;</pre> <p>The <i>hostMem</i> field gives the site's physical memory size.</p> <p>The <i>freeMem</i> field gives the amount of currently unallocated physical memory.</p> <p>The <i>lockMem</i> field gives the amount of currently allocated and non-swappable physical memory.</p> <p>The <i>par</i> field gives the memory management parameters as described in <i>vmSetPar(2K)</i>.</p>				
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.				
<b>ERRORS</b>	[K_EFAULT]                      Some of the arguments provided are outside the caller's address space.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> 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>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>vmSetPar(2K)</i>				

<b>NAME</b>	vmLock, vmUnlock – Fix data in memory; Data in memory								
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int vmLock(KnCap * actorcap, VmAddr address, VmSize size, VmFlags flags); int vmUnlock(KnCap * actorcap, VmAddr address, VmSize size);</pre>								
<b>FEATURES</b>	MEM_PROTECTED, MEM_VIRTUAL								
<b>DESCRIPTION</b>	<p>The <i>vmLock</i> function locks the translations of the target virtual address range of the target actor address space. The (<i>vmUnlock</i>) function unlocks the translations of the target virtual address range of the target actor address space.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>actorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The target address range is specified using its starting <i>address</i> and its <i>size</i>. Both <i>address</i> and <i>size</i> must be page aligned.</p> <p>The <i>vmLock</i> function causes each virtual address from the target address range to be mapped to physical memory with the required protections. If the K_WRITABLE flag is set in the <i>flags</i> argument, write access is required; otherwise, read access is required. Then <i>vmLock</i> gets a reader lock on the translations from the target address range. The reader lock guarantees that when a thread tries to modify either the physical address or the protections of a locked mapping (using <i>rgnFree</i>, <i>rgnSetProt</i> or <i>lcFlush</i>, for example), it blocks until the corresponding <i>vmUnlock</i> is encountered.</p> <p>If the K_WRITABLE flag is set in the <i>flags</i> argument, <i>vmLock</i> marks the locked pages as modified. This could be necessary if, for example after locking, the pages are modified by a driver using the corresponding physical addresses.</p> <p>An address range locked by a number of non-overlapped <i>vmLock</i> operations may be unlocked by any number of non-overlapped <i>vmUnlock</i> operations.</p>								
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.								
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space.</td> </tr> <tr> <td>[K_EINVAL]</td> <td>An inconsistent actor capability has been provided.</td> </tr> <tr> <td>[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td>[K_EADDR]</td> <td>Some or all addresses from the target address range are invalid.</td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.	[K_EINVAL]	An inconsistent actor capability has been provided.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.	[K_EADDR]	Some or all addresses from the target address range are invalid.
[K_EFAULT]	Some of the arguments provided are outside the caller's address space.								
[K_EINVAL]	An inconsistent actor capability has been provided.								
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.								
[K_EADDR]	Some or all addresses from the target address range are invalid.								

[K\_EPROT] Write access is required but some or all the addresses from the target address range are write—p roTECTED.

[K\_EROUND] *address* isn't page aligned.

[K\_EROUND] *size* isn't page-aligned.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`rgnFree(2K)` , `lcFlush(2K)`



# Index

---

## A

acap — get a c\_actor capability 52  
aconf — get configurable system variables 53  
acreate — creates a c\_actor 54  
acred — get/set c\_actor credentials 58  
actorCreate — create an actor 59  
actorDelete — delete an actor 61  
actorName — get and/or set the symbolic name of an actor 62  
actorPi — get and/or set the protection identifier of an actor 63  
actorPrivilege — get and/or set the privilege of an actor 64  
actorSelf — get the current actor capability 65  
actorStart — Stop an actor; Start an actor 66, 69  
actorStat — get the status of all actors on a site 68  
actorStop — Stop an actor; Start an actor 66, 69  
afexec — create a new c\_actor 71, 75, 79, 83, 87, 91, 95  
afexecl — create a new c\_actor 71, 75, 79, 83, 87, 91, 95  
afexeclc — create a new c\_actor 71, 75, 79, 83, 87, 91, 95  
afexeclp — create a new c\_actor 71, 75, 79, 83, 87, 91, 95  
afexecv — create a new c\_actor 71, 75, 79, 83, 87, 91, 95  
afexecvc — create a new c\_actor 71, 75, 79, 83, 87, 91, 95  
afexecvp — create a new c\_actor 71, 75, 79, 83, 87, 91, 95

agetalparam — Manage the loading of an actor 99, 104, 107, 110  
agetId — get c\_actor's ID 102  
akill — kill an actor 103  
aload — Manage the loading of an actor 99, 104, 107, 110  
alParamBuild — Manage the loading of an actor 99, 104, 107, 110  
alParamUnpack — Manage the loading of an actor 99, 104, 107, 110  
astart — activates a c\_actor 113  
astat — list all active c\_actors 114  
atrace — c\_actor trace 115  
await — wait for c\_actor to terminate or stop 118, 120  
awaits — wait for c\_actor to terminate or stop 118, 120

## D

dladdr — translate address to symbolic information 122  
dlclose — close a dynamic object 124  
dlderror — get diagnostic information 125  
dlopen — gain access to a dynamic object file 126  
dlsym — get the address of a symbol in a dynamic object 129

## E

ethIpcStackAttach — attach an IPC stack to an Ethernet device 131

ethOsiStackAttach — attach an OSI stack to an Ethernet device 132  
eventClear — initialize an event flag set; Clear events in an event flag set; Post events in an event flag set; Wait for events in an event flag set 134, 136, 138, 140  
eventInit — initialize an event flag set; Clear events in an event flag set; Post events in an event flag set; Wait for events in an event flag set 134, 136, 138, 140  
eventPost — initialize an event flag set; Clear events in an event flag set; Post events in an event flag set; Wait for events in an event flag set 134, 136, 138, 140  
eventWait — initialize an event flag set; Clear events in an event flag set; Post events in an event flag set; Wait for events in an event flag set 134, 136, 138, 140  
\_exit — terminate a c\_actor 142

## G

grpAllocate — allocate a port group capability 143  
grpPortInsert — insert a port into a port group; remove a port from a port group 145–146  
grpPortRemove — insert a port into a port group; remove a port from a port group 145–146

## I

intro — introduction to ChorusOS error codes and system calls 33  
ipcCall — send an RPC request and wait for the reply 147  
ipcGetData — get the current message body 149  
ipcReceive — receive a message 150  
ipcReply — reply to the current message 153  
ipcRestore — Save the current message; Restore a saved message as the current message 154–155

ipcSave — Save the current message; Restore a saved message as the current message 154–155  
ipcSend — send a message 156  
ipcSysInfo — get system information about the current message 159  
ipcTarget — build a message target 161

## L

lapDescDup — create a lap; reset a lap descriptor; test if a lap descriptor has been initialized; duplicate a lap descriptor 163, 165, 167, 310  
lapDescIsZero — create a lap; reset a lap descriptor; test if a lap descriptor has been initialized; duplicate a lap descriptor 163, 165, 167, 310  
lapDescZero — create a lap; reset a lap descriptor; test if a lap descriptor has been initialized; duplicate a lap descriptor 163, 165, 167, 310  
lapInvoke — invoke a lap handler 169  
lapResolve — bind a symbolic name to a lap descriptor; unbind the symbolic name bound to a lap descriptor; get a lap descriptor from a lap symbolic name 170, 308, 313

## M

msgAllocate — allocates a message from a message space 172  
msgFree — free a message of a message space 174  
msgGet — retrieves the first message of a message queue 175  
msgPut — post a message to a message queue 177  
msgRemove — remove a message from a message queue 178  
msgSpaceCreate — create a message space 179  
msgSpaceOpen — open a message space 181

mutexGet — initialize a mutex; acquire a mutex; release a mutex; try to acquire a mutex 182, 184, 186, 188  
mutexInit — initialize a mutex; acquire a mutex; release a mutex; try to acquire a mutex 182, 184, 186, 188  
mutexRel — initialize a mutex; acquire a mutex; release a mutex; try to acquire a mutex 182, 184, 186, 188  
mutexTry — initialize a mutex; acquire a mutex; release a mutex; try to acquire a mutex 182, 184, 186, 188

## P

padGet — return actor-specific values associated with keys 190  
padKeyCreate — create a private key for an actor 191  
padKeyDelete — delete an actor private key 193  
padSet — set the actor's key to a specific value 194  
portCreate — create a port; declare a port 195, 197  
portDeclare — create a port; declare a port 195, 197  
portDelete — delete a port 199  
portEnable — Enable a port portDisable; Disable a port 200  
portGetSeqNum — Migrate a port; Get a port sequence number 201, 204  
portLi — Get the unique identifier of a port, given its local identifier; Get the local identifier of a port, given its unique identifier 203, 208  
portMigrate — Migrate a port; Get a port sequence number 201, 204  
portPi — get and/or set the protection identifier of a port 206  
portUi — Get the unique identifier of a port, given its local identifier; Get the local identifier of a port, given its unique identifier 203, 208

ptdErrnoAddr — return a thread-specific errno address 209  
ptdGet — return thread-specific value associated with key 210  
ptdKeyCreate — create a thread-specific data key 211  
ptdKeyDelete — delete a thread-specific data key 213  
ptdRemoteGet — return a thread-specific data value for another thread 214  
ptdRemoteSet — set a thread-specific data value for another thread 215  
ptdSet — set a thread-specific value 216  
ptdThreadDelete — delete all thread-specific values and call destructors 217  
ptdThreadId — return the thread ID 218

## R

rgnAllocate — allocate a region in an actor address space 219  
rgnDup — duplicate an actor address space 223  
rgnFree — deallocate regions of an actor address space 225  
rgnInitFromActor — allocate a region in an actor address space and initialise it from another region 227  
rgnMapFromActor — create a region in an actor address space and map another region to it 229  
rgnPhysMap — create a region in an actor address space and map (on demand) to it physical memory specified by the caller 231  
rgnSetInherit — Change inheritance options associated with a region; Change paging options associated with a region; Change inheritance options associated with a region; Change opaque values associated with a region 233, 235, 237

rgnSetOpaque — Change inheritance options associated with a region; Change paging options associated with a region; Change inheritance options associated with a region; Change opaque values associated with a region 233, 235, 237  
 rgnSetPaging — Change inheritance options associated with a region; Change paging options associated with a region; Change inheritance options associated with a region; Change opaque values associated with a region 233, 235, 237  
 rgnSetProtect — change protection options associated with a region 239  
 rgnStat — get the statistics of a region of an actor address space 241  
 rtMutexGet — Initialize a realtime mutex; Acquire a realtime mutex; Release a realtime mutex; Try to acquire a realtime mutex 244, 246, 248, 250  
 rtMutexInit — Initialize a realtime mutex; Acquire a realtime mutex; Release a realtime mutex; Try to acquire a realtime mutex 244, 246, 248, 250  
 rtMutexRel — Initialize a realtime mutex; Acquire a realtime mutex; Release a realtime mutex; Try to acquire a realtime mutex 244, 246, 248, 250  
 rtMutexTry — Initialize a realtime mutex; Acquire a realtime mutex; Release a realtime mutex; Try to acquire a realtime mutex 244, 246, 248, 250

## S

schedAdmin — scheduling classes administration 252

semInit — initialize a semaphore; wait on a semaphore; signal a semaphore 253, 255, 257  
 semP — initialize a semaphore; wait on a semaphore; signal a semaphore 253, 255, 257  
 semV — initialize a semaphore; wait on a semaphore; signal a semaphore 253, 255, 257  
 svAbortHandler — Define an exception handler; Define an abort handler 259, 305  
 svActorAbortHandler — Connect an actor abort handler; Disconnect an actor abort handler; Get an actor abort handler 261, 263, 265, 267  
 svActorAbortHandlerConnect — Connect an actor abort handler; Disconnect an actor abort handler; Get an actor abort handler 261, 263, 265, 267  
 svActorAbortHandlerDisconnect — Connect an actor abort handler; Disconnect an actor abort handler; Get an actor abort handler 261, 263, 265, 267  
 svActorAbortHandlerGetConnected — Connect an actor abort handler; Disconnect an actor abort handler; Get an actor abort handler 261, 263, 265, 267  
 svActorExcHandler — Connect an actor exception handler; Disconnect an actor exception handler; Get an actor exception handler 269, 272, 275, 278  
 svActorExcHandlerConnect — Connect an actor exception handler; Disconnect an actor exception handler; Get an actor exception handler 269, 272, 275, 278  
 svActorExcHandlerDisconnect — Connect an actor exception handler; Disconnect an actor exception handler; Get an actor

exception handler 269, 272, 275, 278

svActorExceptionHandlerGetConnected — Connect an actor exception handler; Disconnect an actor exception handler; Get an actor exception handler 269, 272, 275, 278

svActorStopHandler — Actor stop handler management 281, 284, 287, 290

svActorStopHandlerConnect — Connect an actor stop handler 281, 284, 287, 290

svActorStopHandlerDisconnect — Disconnect an actor stop handler 281, 284, 287, 290

svActorStopHandlerGetConnected — Get an actor stop handler 281, 284, 287, 290

svActorVirtualTimeout — Set an actor's virtual timeout; Cancel an actor's virtual timeout 293, 295, 297

svActorVirtualTimeoutCancel — Set an actor's virtual timeout; Cancel an actor's virtual timeout 293, 295, 297

svActorVirtualTimeoutSet — Set an actor's virtual timeout; Cancel an actor's virtual timeout 293, 295, 297

svCopyIn — Copy from trap caller space; Copy string from trap caller space; Copy to trap caller space 299, 301, 303

svCopyInString — Copy from trap caller space; Copy string from trap caller space; Copy to trap caller space 299, 301, 303

svCopyOut — Copy from trap caller space; Copy string from trap caller space; Copy to trap caller space 299, 301, 303

svExceptionHandler — Define an exception handler; Define an abort handler 259, 305

svGetInvoker — get handler invoker 307

svLapBind — bind a symbolic name to a lap descriptor; unbind the symbolic name bound to a lap descriptor; get a lap descriptor from a lap symbolic name 170, 308, 313

svLapCreate — create a lap; reset a lap descriptor; test if a lap descriptor has been initialized; duplicate a lap descriptor 163, 165, 167, 310

svLapDelete — delete a lap 312

svLapUnbind — bind a symbolic name to a lap descriptor; unbind the symbolic name bound to a lap descriptor; get a lap descriptor from a lap symbolic name 170, 308, 313

svMaskAll — Disable interrupt processing; Enable interrupt processing; Reenable interrupt processing 315, 380–381

svMaskedLockGet — Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock 316, 319, 322, 325, 340, 343, 346, 349

svMaskedLockInit — Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock 316, 319, 322, 325, 340, 343, 346, 349

svMaskedLockRel — Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and

- enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock 316, 319, 322, 325, 340, 343, 346, 349
- svMaskedLockTry — Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock 316, 319, 322, 325, 340, 343, 346, 349
- svMemRead — Copy from supervisor caller space; Copy to supervisor caller space 328–329
- svMemWrite — Copy from supervisor caller space; Copy to supervisor caller space 328–329
- svMsgHandler — Connect/disconnect a message handler; Prepare a reply to a handled message 330, 333
- svMsgHdlReply — Connect/disconnect a message handler; Prepare a reply to a handled message 330, 333
- svPagesAllocate — supervisor address space memory allocator 336, 338
- svPagesFree — supervisor address space memory allocator 336, 338
- svSpinLockGet — Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock 316, 319, 322, 325, 340, 343, 346, 349
- svSpinLockInit — Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock 316, 319, 322, 325, 340, 343, 346, 349
- svSpinLockRel — Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock 316, 319, 322, 325, 340, 343, 346, 349
- svSpinLockTry — Initialize a spin lock; Disable interrupts and acquire a spin lock; Try to disable interrupts and acquire a spin lock; Release a spin lock and enable interrupts; Initialize a spin lock; Acquire a spin lock; Try to acquire a spin lock; Release a spin lock 316, 319, 322, 325, 340, 343, 346, 349
- svSysCtx — get system context table address 352
- svSysPanic — trigger the invocation of the panic handler 353
- svSysTimeout — Request a timeout; Cancel a timeout; Get timeout resolution 354, 356, 358, 374
- svSysTimeoutCancel — Cancel a timeout 354, 356, 358, 374
- svSysTimeoutSet — Request a timeout 354, 356, 358, 374
- svSysTrapHandler — Connect a trap handler; Disconnect a trap handler; Get a trap handler 360, 362, 364, 366
- svSysTrapHandlerConnect — Connect a trap handler; Disconnect a trap handler; Get a trap handler 360, 362, 364, 366
- svSysTrapHandlerDisconnect — Connect a trap handler; Disconnect

- a trap handler; Get a trap handler 360, 362, 364, 366
- svSysTrapHandlerGetConnected — Connect a trap handler; Disconnect a trap handler; Get a trap handler 360, 362, 364, 366
- svThreadVirtualTimeout — Set a thread's virtual timeout; Cancel a thread's virtual timeout 368, 370, 372
- svThreadVirtualTimeoutCancel — Set a thread's virtual timeout; Cancel a thread's virtual timeout 368, 370, 372
- svThreadVirtualTimeoutSet — Set a thread's virtual timeout; Cancel a thread's virtual timeout 368, 370, 372
- svTimeoutGetRes — Get timeout resolution 354, 356, 358, 374
- svTrapConnect — Connect a trap handler; Disconnect a trap handler 376, 378
- svTrapDisConnect — Connect a trap handler; Disconnect a trap handler 376, 378
- svUnmask — Disable interrupt processing; Enable interrupt processing; Reenable interrupt processing 315, 380–381
- svUnmaskAll — Disable interrupt processing; Enable interrupt processing; Reenable interrupt processing 315, 380–381
- svVirtualTimeoutCancel — Set a virtual timeout; Cancel a virtual timeout 382, 384
- svVirtualTimeoutSet — Set a virtual timeout; Cancel a virtual timeout 382, 384
- sysBench — kernel benchmark utility 386
- sysGetConf — Get Chorus module configuration value 388
- sysGetEnv — Get a value from the Chorus configuration environment 389
- sysLog — log a message in the kernel's cyclical buffer 390
- sysPoll — Read characters from the system console; Write characters to the system console; Poll characters from the system console 391–392, 413
- sysRead — Read characters from the system console; Write characters to the system console; Poll characters from the system console 391–392, 413
- sysReboot — request a reboot of the local site 393
- sysSetEnv — Set a value in the ChorusOS configuration environment 394
- sysShutdown — shut down the system 395
- sysTime — get system time; get system time resolution 396–397
- sysTimeGetRes — get system time; get system time resolution 396–397
- sysTimer — system timer management 398, 400, 402, 404, 406, 408, 410
- sysTimerGetCounterFrequency — system timer management 398, 400, 402, 404, 406, 408, 410
- sysTimerGetCounterPeriod — system timer management 398, 400, 402, 404, 406, 408, 410
- sysTimerReadCounter — system timer management 398, 400, 402, 404, 406, 408, 410
- sysTimerStartFreerun — system timer management 398, 400, 402, 404, 406, 408, 410
- sysTimerStartPeriodic — system timer management 398, 400, 402, 404, 406, 408, 410
- sysTimerStop — system timer management 398, 400, 402, 404, 406, 408, 410
- sysUnsetEnv — delete a value from the ChorusOS configuration environment 412
- sysWrite — Read characters from the system console; Write characters

to the system console; Poll characters from the system console 391-392, 413

## T

`threadAbort` — Abort a thread; Check whether the current thread has been aborted 414, 416

`threadAborted` — Abort a thread; Check whether the current thread has been aborted 414, 416

`threadActivate` — make a thread active 418

`threadBind` — bind a thread to a processor 419

`threadContext` — get and/or set the context of a thread 421

`threadCreate` — create a thread 423

`threadDelay` — delay the current thread 426

`threadDelete` — delete a thread 427

`threadLoadR` — Get the current thread's valid soft register value; Reset the current thread's valid soft register value 428, 451

`threadName` — get and/or set the symbolic name of a thread 430

`threadResume` — Suspend a thread; Resume a thread 431, 453

`threadScheduler` — get and/or set thread scheduling parameters 433

`threadSelf` — get the current thread local identifier 438

`threadSemInit` — Initialize a thread semaphore; Signal a thread semaphore; Wait on a thread semaphore 439, 441, 443

`threadSemPost` — Initialize a thread semaphore; Signal a thread semaphore; Wait on a thread semaphore 439, 441, 443

`threadSemWait` — Initialize a thread semaphore; Signal a thread semaphore; Wait on a thread semaphore 439, 441, 443

`threadStart` — Stop a thread; Start a thread 445, 449

`threadStat` — obtain the descriptions of the threads running in actor 447

`threadStop` — Stop a thread; Start a thread 445, 449

`threadStoreR` — Get the current thread's valid soft register value; Reset the current thread's valid soft register value 428, 451

`threadSuspend` — Suspend a thread; Resume a thread 431, 453

`threadTimes` — get thread execution time 455

`timerCreate` — create a timer 456

`timerDelete` — delete a timer 457

`timerGetRes` — get the timer resolution 458

`timerSet` — start, cancel or query a timer 459

`timerThreadPoolInit` — initialize a timer thread pool 460

`timerThreadPoolWait` — wait for a timer expiration event 461

## U

`uiBuild` — Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared 463, 465, 467, 469, 471, 474, 476

`uiClear` — Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared 463, 465, 467, 469, 471, 474, 476

`uiEqual` — Build a user-defined unique identifier; Clear a unique identifier; Compare two

unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared 463, 465, 467, 469, 471, 474, 476

uiGetSite — Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared 463, 465, 467, 469, 471, 474, 476

uiIsLocal — Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared 463, 465, 467, 469, 471, 474, 476

uiLocalSite — get the local site number 473

uiSite — Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared 463, 465, 467, 469, 471, 474, 476

uiValid — Build a user-defined unique identifier; Clear a unique identifier; Compare two unique identifiers; Extract the site number from a unique identifier; Check a unique identifier's locality; Get a pre-defined site unique identifier; Check whether an unique identifier has been cleared 463, 465, 467, 469, 471, 474, 476

univTime — Get time-of-day; Set time-of-day; Adjust time-of-day univTimeGetRes; Get time-of-day resolution 478, 480, 482

univTimeAdjust — Get time-of-day; Set time-of-day; Adjust time-of-day univTimeGetRes; Get time-of-day resolution 478, 480, 482

univTimeSet — Get time-of-day; Set time-of-day; Adjust time-of-day univTimeGetRes; Get time-of-day resolution 478, 480, 482

## V

virtualTimeGetRes — get virtual time resolution 484

vmCopy — copy data between actor address spaces 485

vmFree — free memory 486

vmLock — Fix data in memory; Data in memory 487, 494

vmPageSize — get the minimum allocatable memory block size 489

vmPhysAddr — get the physical address corresponding to a virtual address 490

vmSetPar — set the memory management parameters 491

vmStat — get memory management  
statistics 493

vmUnLock — Fix data in memory; Data in  
memory 487, 494