



ChorusOS man pages section 9DKI: Driver to Kernel Interface

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

Part No: 806-3342
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 57

Intro(9DKI) 63

dataCacheBlockFlush(9DKI) 71

dataCacheInvalidate(9DKI) 71

dataCacheBlockInvalidate(9DKI) 71

instCacheInvalidate(9DKI) 71

instCacheBlockInvalidate(9DKI) 71

dataCacheBlockFlush_powerpc(9DKI) 72

dataCacheInvalidate_powerpc(9DKI) 72

dataCacheBlockInvalidate_powerpc(9DKI) 72

instCacheInvalidate_powerpc(9DKI) 72

instCacheBlockInvalidate_powerpc(9DKI) 72

dataCacheBlockFlush(9DKI) 75

dataCacheInvalidate(9DKI) 75

dataCacheBlockInvalidate(9DKI) 75

instCacheInvalidate(9DKI) 75

instCacheBlockInvalidate(9DKI) 75

dataCacheBlockFlush_powerpc(9DKI) 76

dataCacheInvalidate_powerpc(9DKI) 76

dataCacheBlockInvalidate_powerpc(9DKI) 76
instCacheInvalidate_powerpc(9DKI) 76
instCacheBlockInvalidate_powerpc(9DKI) 76
dataCacheBlockFlush(9DKI) 79
dataCacheInvalidate(9DKI) 79
dataCacheBlockInvalidate(9DKI) 79
instCacheInvalidate(9DKI) 79
instCacheBlockInvalidate(9DKI) 79
dataCacheBlockFlush_powerpc(9DKI) 80
dataCacheInvalidate_powerpc(9DKI) 80
dataCacheBlockInvalidate_powerpc(9DKI) 80
instCacheInvalidate_powerpc(9DKI) 80
instCacheBlockInvalidate_powerpc(9DKI) 80
icacheInval(9DKI) 83
icacheLineInval(9DKI) 83
icacheBlockInval(9DKI) 83
dcacheFlush(9DKI) 83
dcacheLineFlush(9DKI) 83
dcacheBlockFlush(9DKI) 83
icacheInval_usparc(9DKI) 84
icacheLineInval_usparc(9DKI) 84
icacheBlockInval_usparc(9DKI) 84
dcacheFlush_usparc(9DKI) 84
dcacheLineFlush_usparc(9DKI) 84
dcacheBlockFlush_usparc(9DKI) 84
icacheInval(9DKI) 87
icacheLineInval(9DKI) 87
icacheBlockInval(9DKI) 87

dcacheFlush(9DKI)	87
dcacheLineFlush(9DKI)	87
dcacheBlockFlush(9DKI)	87
icacheInval_usparc(9DKI)	88
icacheLineInval_usparc(9DKI)	88
icacheBlockInval_usparc(9DKI)	88
dcacheFlush_usparc(9DKI)	88
dcacheLineFlush_usparc(9DKI)	88
dcacheBlockFlush_usparc(9DKI)	88
icacheInval(9DKI)	91
icacheLineInval(9DKI)	91
icacheBlockInval(9DKI)	91
dcacheFlush(9DKI)	91
dcacheLineFlush(9DKI)	91
dcacheBlockFlush(9DKI)	91
icacheInval_usparc(9DKI)	92
icacheLineInval_usparc(9DKI)	92
icacheBlockInval_usparc(9DKI)	92
dcacheFlush_usparc(9DKI)	92
dcacheLineFlush_usparc(9DKI)	92
dcacheBlockFlush_usparc(9DKI)	92
DISABLE_PREEMPT(9DKI)	95
ENABLE_PREEMPT(9DKI)	95
dtreeNodeRoot(9DKI)	97
dtreeNodeChild(9DKI)	97
dtreeNodePeer(9DKI)	97
dtreeNodeParent(9DKI)	97
dtreeNodeAlloc(9DKI)	97

dtreeNodeFree(9DKI) 97
dtreeNodeAttach(9DKI) 97
dtreeNodeDetach(9DKI) 97
dtreePropFind(9DKI) 97
dtreePropFindNext(9DKI) 97
dtreePropLength(9DKI) 97
dtreePropValue(9DKI) 97
dtreePropName(9DKI) 97
dtreePropAlloc(9DKI) 97
dtreePropFree(9DKI) 97
dtreePropAttach(9DKI) 97
dtreePropDetach(9DKI) 97
dtreeNodeAdd(9DKI) 97
dtreeNodeFind(9DKI) 97
dtreePropAdd(9DKI) 97
dtreePathLeng(9DKI) 97
dtreePathGet(9DKI) 97
dtreeNodeRoot(9DKI) 103
dtreeNodeChild(9DKI) 103
dtreeNodePeer(9DKI) 103
dtreeNodeParent(9DKI) 103
dtreeNodeAlloc(9DKI) 103
dtreeNodeFree(9DKI) 103
dtreeNodeAttach(9DKI) 103
dtreeNodeDetach(9DKI) 103
dtreePropFind(9DKI) 103
dtreePropFindNext(9DKI) 103
dtreePropLength(9DKI) 103

dtreePropValue(9DKI)	103
dtreePropName(9DKI)	103
dtreePropAlloc(9DKI)	103
dtreePropFree(9DKI)	103
dtreePropAttach(9DKI)	103
dtreePropDetach(9DKI)	103
dtreeNodeAdd(9DKI)	103
dtreeNodeFind(9DKI)	103
dtreePropAdd(9DKI)	103
dtreePathLeng(9DKI)	103
dtreePathGet(9DKI)	103
dtreeNodeRoot(9DKI)	109
dtreeNodeChild(9DKI)	109
dtreeNodePeer(9DKI)	109
dtreeNodeParent(9DKI)	109
dtreeNodeAlloc(9DKI)	109
dtreeNodeFree(9DKI)	109
dtreeNodeAttach(9DKI)	109
dtreeNodeDetach(9DKI)	109
dtreePropFind(9DKI)	109
dtreePropFindNext(9DKI)	109
dtreePropLength(9DKI)	109
dtreePropValue(9DKI)	109
dtreePropName(9DKI)	109
dtreePropAlloc(9DKI)	109
dtreePropFree(9DKI)	109
dtreePropAttach(9DKI)	109
dtreePropDetach(9DKI)	109

dtreeNodeAdd(9DKI) 109
dtreeNodeFind(9DKI) 109
dtreePropAdd(9DKI) 109
dtreePathLeng(9DKI) 109
dtreePathGet(9DKI) 109
dtreeNodeRoot(9DKI) 115
dtreeNodeChild(9DKI) 115
dtreeNodePeer(9DKI) 115
dtreeNodeParent(9DKI) 115
dtreeNodeAlloc(9DKI) 115
dtreeNodeFree(9DKI) 115
dtreeNodeAttach(9DKI) 115
dtreeNodeDetach(9DKI) 115
dtreePropFind(9DKI) 115
dtreePropFindNext(9DKI) 115
dtreePropLength(9DKI) 115
dtreePropValue(9DKI) 115
dtreePropName(9DKI) 115
dtreePropAlloc(9DKI) 115
dtreePropFree(9DKI) 115
dtreePropAttach(9DKI) 115
dtreePropDetach(9DKI) 115
dtreeNodeAdd(9DKI) 115
dtreeNodeFind(9DKI) 115
dtreePropAdd(9DKI) 115
dtreePathLeng(9DKI) 115
dtreePathGet(9DKI) 115
dtreeNodeRoot(9DKI) 121

dtreeNodeChild(9DKI)	121
dtreeNodePeer(9DKI)	121
dtreeNodeParent(9DKI)	121
dtreeNodeAlloc(9DKI)	121
dtreeNodeFree(9DKI)	121
dtreeNodeAttach(9DKI)	121
dtreeNodeDetach(9DKI)	121
dtreePropFind(9DKI)	121
dtreePropFindNext(9DKI)	121
dtreePropLength(9DKI)	121
dtreePropValue(9DKI)	121
dtreePropName(9DKI)	121
dtreePropAlloc(9DKI)	121
dtreePropFree(9DKI)	121
dtreePropAttach(9DKI)	121
dtreePropDetach(9DKI)	121
dtreeNodeAdd(9DKI)	121
dtreeNodeFind(9DKI)	121
dtreePropAdd(9DKI)	121
dtreePathLeng(9DKI)	121
dtreePathGet(9DKI)	121
dtreeNodeRoot(9DKI)	127
dtreeNodeChild(9DKI)	127
dtreeNodePeer(9DKI)	127
dtreeNodeParent(9DKI)	127
dtreeNodeAlloc(9DKI)	127
dtreeNodeFree(9DKI)	127
dtreeNodeAttach(9DKI)	127

dtreeNodeDetach(9DKI) 127
dtreePropFind(9DKI) 127
dtreePropFindNext(9DKI) 127
dtreePropLength(9DKI) 127
dtreePropValue(9DKI) 127
dtreePropName(9DKI) 127
dtreePropAlloc(9DKI) 127
dtreePropFree(9DKI) 127
dtreePropAttach(9DKI) 127
dtreePropDetach(9DKI) 127
dtreeNodeAdd(9DKI) 127
dtreeNodeFind(9DKI) 127
dtreePropAdd(9DKI) 127
dtreePathLeng(9DKI) 127
dtreePathGet(9DKI) 127
dtreeNodeRoot(9DKI) 133
dtreeNodeChild(9DKI) 133
dtreeNodePeer(9DKI) 133
dtreeNodeParent(9DKI) 133
dtreeNodeAlloc(9DKI) 133
dtreeNodeFree(9DKI) 133
dtreeNodeAttach(9DKI) 133
dtreeNodeDetach(9DKI) 133
dtreePropFind(9DKI) 133
dtreePropFindNext(9DKI) 133
dtreePropLength(9DKI) 133
dtreePropValue(9DKI) 133
dtreePropName(9DKI) 133

dtreePropAlloc(9DKI) 133
dtreePropFree(9DKI) 133
dtreePropAttach(9DKI) 133
dtreePropDetach(9DKI) 133
dtreeNodeAdd(9DKI) 133
dtreeNodeFind(9DKI) 133
dtreePropAdd(9DKI) 133
dtreePathLeng(9DKI) 133
dtreePathGet(9DKI) 133
dtreeNodeRoot(9DKI) 139
dtreeNodeChild(9DKI) 139
dtreeNodePeer(9DKI) 139
dtreeNodeParent(9DKI) 139
dtreeNodeAlloc(9DKI) 139
dtreeNodeFree(9DKI) 139
dtreeNodeAttach(9DKI) 139
dtreeNodeDetach(9DKI) 139
dtreePropFind(9DKI) 139
dtreePropFindNext(9DKI) 139
dtreePropLength(9DKI) 139
dtreePropValue(9DKI) 139
dtreePropName(9DKI) 139
dtreePropAlloc(9DKI) 139
dtreePropFree(9DKI) 139
dtreePropAttach(9DKI) 139
dtreePropDetach(9DKI) 139
dtreeNodeAdd(9DKI) 139
dtreeNodeFind(9DKI) 139

dtreePropAdd(9DKI) 139
dtreePathLeng(9DKI) 139
dtreePathGet(9DKI) 139
dtreeNodeRoot(9DKI) 145
dtreeNodeChild(9DKI) 145
dtreeNodePeer(9DKI) 145
dtreeNodeParent(9DKI) 145
dtreeNodeAlloc(9DKI) 145
dtreeNodeFree(9DKI) 145
dtreeNodeAttach(9DKI) 145
dtreeNodeDetach(9DKI) 145
dtreePropFind(9DKI) 145
dtreePropFindNext(9DKI) 145
dtreePropLength(9DKI) 145
dtreePropValue(9DKI) 145
dtreePropName(9DKI) 145
dtreePropAlloc(9DKI) 145
dtreePropFree(9DKI) 145
dtreePropAttach(9DKI) 145
dtreePropDetach(9DKI) 145
dtreeNodeAdd(9DKI) 145
dtreeNodeFind(9DKI) 145
dtreePropAdd(9DKI) 145
dtreePathLeng(9DKI) 145
dtreePathGet(9DKI) 145
dtreeNodeRoot(9DKI) 151
dtreeNodeChild(9DKI) 151
dtreeNodePeer(9DKI) 151

dtreeNodeParent(9DKI)	151
dtreeNodeAlloc(9DKI)	151
dtreeNodeFree(9DKI)	151
dtreeNodeAttach(9DKI)	151
dtreeNodeDetach(9DKI)	151
dtreePropFind(9DKI)	151
dtreePropFindNext(9DKI)	151
dtreePropLength(9DKI)	151
dtreePropValue(9DKI)	151
dtreePropName(9DKI)	151
dtreePropAlloc(9DKI)	151
dtreePropFree(9DKI)	151
dtreePropAttach(9DKI)	151
dtreePropDetach(9DKI)	151
dtreeNodeAdd(9DKI)	151
dtreeNodeFind(9DKI)	151
dtreePropAdd(9DKI)	151
dtreePathLeng(9DKI)	151
dtreePathGet(9DKI)	151
dtreeNodeRoot(9DKI)	157
dtreeNodeChild(9DKI)	157
dtreeNodePeer(9DKI)	157
dtreeNodeParent(9DKI)	157
dtreeNodeAlloc(9DKI)	157
dtreeNodeFree(9DKI)	157
dtreeNodeAttach(9DKI)	157
dtreeNodeDetach(9DKI)	157
dtreePropFind(9DKI)	157

dtreePropFindNext(9DKI) 157
dtreePropLength(9DKI) 157
dtreePropValue(9DKI) 157
dtreePropName(9DKI) 157
dtreePropAlloc(9DKI) 157
dtreePropFree(9DKI) 157
dtreePropAttach(9DKI) 157
dtreePropDetach(9DKI) 157
dtreeNodeAdd(9DKI) 157
dtreeNodeFind(9DKI) 157
dtreePropAdd(9DKI) 157
dtreePathLeng(9DKI) 157
dtreePathGet(9DKI) 157
dtreeNodeRoot(9DKI) 163
dtreeNodeChild(9DKI) 163
dtreeNodePeer(9DKI) 163
dtreeNodeParent(9DKI) 163
dtreeNodeAlloc(9DKI) 163
dtreeNodeFree(9DKI) 163
dtreeNodeAttach(9DKI) 163
dtreeNodeDetach(9DKI) 163
dtreePropFind(9DKI) 163
dtreePropFindNext(9DKI) 163
dtreePropLength(9DKI) 163
dtreePropValue(9DKI) 163
dtreePropName(9DKI) 163
dtreePropAlloc(9DKI) 163
dtreePropFree(9DKI) 163

dtreePropAttach(9DKI)	163
dtreePropDetach(9DKI)	163
dtreeNodeAdd(9DKI)	163
dtreeNodeFind(9DKI)	163
dtreePropAdd(9DKI)	163
dtreePathLeng(9DKI)	163
dtreePathGet(9DKI)	163
dtreeNodeRoot(9DKI)	169
dtreeNodeChild(9DKI)	169
dtreeNodePeer(9DKI)	169
dtreeNodeParent(9DKI)	169
dtreeNodeAlloc(9DKI)	169
dtreeNodeFree(9DKI)	169
dtreeNodeAttach(9DKI)	169
dtreeNodeDetach(9DKI)	169
dtreePropFind(9DKI)	169
dtreePropFindNext(9DKI)	169
dtreePropLength(9DKI)	169
dtreePropValue(9DKI)	169
dtreePropName(9DKI)	169
dtreePropAlloc(9DKI)	169
dtreePropFree(9DKI)	169
dtreePropAttach(9DKI)	169
dtreePropDetach(9DKI)	169
dtreeNodeAdd(9DKI)	169
dtreeNodeFind(9DKI)	169
dtreePropAdd(9DKI)	169
dtreePathLeng(9DKI)	169

dtreePathGet(9DKI) 169
dtreeNodeRoot(9DKI) 175
dtreeNodeChild(9DKI) 175
dtreeNodePeer(9DKI) 175
dtreeNodeParent(9DKI) 175
dtreeNodeAlloc(9DKI) 175
dtreeNodeFree(9DKI) 175
dtreeNodeAttach(9DKI) 175
dtreeNodeDetach(9DKI) 175
dtreePropFind(9DKI) 175
dtreePropFindNext(9DKI) 175
dtreePropLength(9DKI) 175
dtreePropValue(9DKI) 175
dtreePropName(9DKI) 175
dtreePropAlloc(9DKI) 175
dtreePropFree(9DKI) 175
dtreePropAttach(9DKI) 175
dtreePropDetach(9DKI) 175
dtreeNodeAdd(9DKI) 175
dtreeNodeFind(9DKI) 175
dtreePropAdd(9DKI) 175
dtreePathLeng(9DKI) 175
dtreePathGet(9DKI) 175
dtreeNodeRoot(9DKI) 181
dtreeNodeChild(9DKI) 181
dtreeNodePeer(9DKI) 181
dtreeNodeParent(9DKI) 181
dtreeNodeAlloc(9DKI) 181

dtreeNodeFree(9DKI)	181
dtreeNodeAttach(9DKI)	181
dtreeNodeDetach(9DKI)	181
dtreePropFind(9DKI)	181
dtreePropFindNext(9DKI)	181
dtreePropLength(9DKI)	181
dtreePropValue(9DKI)	181
dtreePropName(9DKI)	181
dtreePropAlloc(9DKI)	181
dtreePropFree(9DKI)	181
dtreePropAttach(9DKI)	181
dtreePropDetach(9DKI)	181
dtreeNodeAdd(9DKI)	181
dtreeNodeFind(9DKI)	181
dtreePropAdd(9DKI)	181
dtreePathLeng(9DKI)	181
dtreePathGet(9DKI)	181
dtreeNodeRoot(9DKI)	187
dtreeNodeChild(9DKI)	187
dtreeNodePeer(9DKI)	187
dtreeNodeParent(9DKI)	187
dtreeNodeAlloc(9DKI)	187
dtreeNodeFree(9DKI)	187
dtreeNodeAttach(9DKI)	187
dtreeNodeDetach(9DKI)	187
dtreePropFind(9DKI)	187
dtreePropFindNext(9DKI)	187
dtreePropLength(9DKI)	187

dtreePropValue(9DKI) 187
dtreePropName(9DKI) 187
dtreePropAlloc(9DKI) 187
dtreePropFree(9DKI) 187
dtreePropAttach(9DKI) 187
dtreePropDetach(9DKI) 187
dtreeNodeAdd(9DKI) 187
dtreeNodeFind(9DKI) 187
dtreePropAdd(9DKI) 187
dtreePathLeng(9DKI) 187
dtreePathGet(9DKI) 187
dtreeNodeRoot(9DKI) 193
dtreeNodeChild(9DKI) 193
dtreeNodePeer(9DKI) 193
dtreeNodeParent(9DKI) 193
dtreeNodeAlloc(9DKI) 193
dtreeNodeFree(9DKI) 193
dtreeNodeAttach(9DKI) 193
dtreeNodeDetach(9DKI) 193
dtreePropFind(9DKI) 193
dtreePropFindNext(9DKI) 193
dtreePropLength(9DKI) 193
dtreePropValue(9DKI) 193
dtreePropName(9DKI) 193
dtreePropAlloc(9DKI) 193
dtreePropFree(9DKI) 193
dtreePropAttach(9DKI) 193
dtreePropDetach(9DKI) 193

dtreeNodeAdd(9DKI)	193
dtreeNodeFind(9DKI)	193
dtreePropAdd(9DKI)	193
dtreePathLeng(9DKI)	193
dtreePathGet(9DKI)	193
dtreeNodeRoot(9DKI)	199
dtreeNodeChild(9DKI)	199
dtreeNodePeer(9DKI)	199
dtreeNodeParent(9DKI)	199
dtreeNodeAlloc(9DKI)	199
dtreeNodeFree(9DKI)	199
dtreeNodeAttach(9DKI)	199
dtreeNodeDetach(9DKI)	199
dtreePropFind(9DKI)	199
dtreePropFindNext(9DKI)	199
dtreePropLength(9DKI)	199
dtreePropValue(9DKI)	199
dtreePropName(9DKI)	199
dtreePropAlloc(9DKI)	199
dtreePropFree(9DKI)	199
dtreePropAttach(9DKI)	199
dtreePropDetach(9DKI)	199
dtreeNodeAdd(9DKI)	199
dtreeNodeFind(9DKI)	199
dtreePropAdd(9DKI)	199
dtreePathLeng(9DKI)	199
dtreePathGet(9DKI)	199
loadSwapEieio_16(9DKI)	205

storeSwapEieio_16(9DKI) 205
swapEieio_16(9DKI) 205
loadSwapEieio_32(9DKI) 205
storeSwapEieio_32(9DKI) 205
swapEieio_32(9DKI) 205
eieio(9DKI) 205
loadSwapEieio_16_powerpc(9DKI) 206
loadSwapSync_16_powerpc(9DKI) 206
storeSwapEieio_16_powerpc(9DKI) 206
storeSwapSync_16_powerpc(9DKI) 206
swapEieio_16_powerpc(9DKI) 206
loadSwapEieio_32_powerpc(9DKI) 206
loadSwapSync_32_powerpc(9DKI) 206
storeSwapEieio_32_powerpc(9DKI) 206
storeSwapSync_32_powerpc(9DKI) 206
swapEieio_32_powerpc(9DKI) 206
eieio_powerpc(9DKI) 206
ioSync_powerpc(9DKI) 206
DISABLE_PREEMPT(9DKI) 209
ENABLE_PREEMPT(9DKI) 209
icacheInval(9DKI) 211
icacheLineInval(9DKI) 211
icacheBlockInval(9DKI) 211
dcacheFlush(9DKI) 211
dcacheLineFlush(9DKI) 211
dcacheBlockFlush(9DKI) 211
icacheInval_usparc(9DKI) 212
icacheLineInval_usparc(9DKI) 212

icacheBlockInval_usparc(9DKI)	212
dcacheFlush_usparc(9DKI)	212
dcacheLineFlush_usparc(9DKI)	212
dcacheBlockFlush_usparc(9DKI)	212
icacheInval(9DKI)	215
icacheLineInval(9DKI)	215
icacheBlockInval(9DKI)	215
dcacheFlush(9DKI)	215
dcacheLineFlush(9DKI)	215
dcacheBlockFlush(9DKI)	215
icacheInval_usparc(9DKI)	216
icacheLineInval_usparc(9DKI)	216
icacheBlockInval_usparc(9DKI)	216
dcacheFlush_usparc(9DKI)	216
dcacheLineFlush_usparc(9DKI)	216
dcacheBlockFlush_usparc(9DKI)	216
icacheInval(9DKI)	219
icacheLineInval(9DKI)	219
icacheBlockInval(9DKI)	219
dcacheFlush(9DKI)	219
dcacheLineFlush(9DKI)	219
dcacheBlockFlush(9DKI)	219
icacheInval_usparc(9DKI)	220
icacheLineInval_usparc(9DKI)	220
icacheBlockInval_usparc(9DKI)	220
dcacheFlush_usparc(9DKI)	220
dcacheLineFlush_usparc(9DKI)	220
dcacheBlockFlush_usparc(9DKI)	220

imsIntrMask_f(9DKI) 223
imsIntrUnmask_f(9DKI) 223
imsIntrMask_f(9DKI) 224
imsIntrUnmask_f(9DKI) 224
dataCacheBlockFlush(9DKI) 225
dataCacheInvalidate(9DKI) 225
dataCacheBlockInvalidate(9DKI) 225
instCacheInvalidate(9DKI) 225
instCacheBlockInvalidate(9DKI) 225
dataCacheBlockFlush_powerpc(9DKI) 226
dataCacheInvalidate_powerpc(9DKI) 226
dataCacheBlockInvalidate_powerpc(9DKI) 226
instCacheInvalidate_powerpc(9DKI) 226
instCacheBlockInvalidate_powerpc(9DKI) 226
dataCacheBlockFlush(9DKI) 229
dataCacheInvalidate(9DKI) 229
dataCacheBlockInvalidate(9DKI) 229
instCacheInvalidate(9DKI) 229
instCacheBlockInvalidate(9DKI) 229
dataCacheBlockFlush_powerpc(9DKI) 230
dataCacheInvalidate_powerpc(9DKI) 230
dataCacheBlockInvalidate_powerpc(9DKI) 230
instCacheInvalidate_powerpc(9DKI) 230
instCacheBlockInvalidate_powerpc(9DKI) 230
ioLoad8(9DKI) 233
ioStore8(9DKI) 233
ioRead8(9DKI) 233
ioWrite8(9DKI) 233

ioLoad16(9DKI)	233
ioStore16(9DKI)	233
ioRead16(9DKI)	233
ioWrite16(9DKI)	233
ioLoad32(9DKI)	233
ioStore32(9DKI)	233
ioRead32(9DKI)	233
ioWrite32(9DKI)	233
ioLoad8_x86(9DKI)	234
ioStore8_x86(9DKI)	234
ioRead8_x86(9DKI)	234
ioWrite8_x86(9DKI)	234
ioLoad16_x86(9DKI)	234
ioStore16_x86(9DKI)	234
ioRead16_x86(9DKI)	234
ioWrite16_x86(9DKI)	234
ioLoad32_x86(9DKI)	234
ioStore32_x86(9DKI)	234
ioRead32_x86(9DKI)	234
ioWrite32_x86(9DKI)	234
ioLoad8(9DKI)	236
ioStore8(9DKI)	236
ioRead8(9DKI)	236
ioWrite8(9DKI)	236
ioLoad16(9DKI)	236
ioStore16(9DKI)	236
ioRead16(9DKI)	236
ioWrite16(9DKI)	236

ioLoad32(9DKI) 236
ioStore32(9DKI) 236
ioRead32(9DKI) 236
ioWrite32(9DKI) 236
ioLoad8_x86(9DKI) 237
ioStore8_x86(9DKI) 237
ioRead8_x86(9DKI) 237
ioWrite8_x86(9DKI) 237
ioLoad16_x86(9DKI) 237
ioStore16_x86(9DKI) 237
ioRead16_x86(9DKI) 237
ioWrite16_x86(9DKI) 237
ioLoad32_x86(9DKI) 237
ioStore32_x86(9DKI) 237
ioRead32_x86(9DKI) 237
ioWrite32_x86(9DKI) 237
ioLoad8(9DKI) 239
ioStore8(9DKI) 239
ioRead8(9DKI) 239
ioWrite8(9DKI) 239
ioLoad16(9DKI) 239
ioStore16(9DKI) 239
ioRead16(9DKI) 239
ioWrite16(9DKI) 239
ioLoad32(9DKI) 239
ioStore32(9DKI) 239
ioRead32(9DKI) 239
ioWrite32(9DKI) 239

ioLoad8_x86(9DKI)	240
ioStore8_x86(9DKI)	240
ioRead8_x86(9DKI)	240
ioWrite8_x86(9DKI)	240
ioLoad16_x86(9DKI)	240
ioStore16_x86(9DKI)	240
ioRead16_x86(9DKI)	240
ioWrite16_x86(9DKI)	240
ioLoad32_x86(9DKI)	240
ioStore32_x86(9DKI)	240
ioRead32_x86(9DKI)	240
ioWrite32_x86(9DKI)	240
ioLoad8(9DKI)	242
ioStore8(9DKI)	242
ioRead8(9DKI)	242
ioWrite8(9DKI)	242
ioLoad16(9DKI)	242
ioStore16(9DKI)	242
ioRead16(9DKI)	242
ioWrite16(9DKI)	242
ioLoad32(9DKI)	242
ioStore32(9DKI)	242
ioRead32(9DKI)	242
ioWrite32(9DKI)	242
ioLoad8_x86(9DKI)	243
ioStore8_x86(9DKI)	243
ioRead8_x86(9DKI)	243
ioWrite8_x86(9DKI)	243

ioLoad16_x86(9DKI) 243
ioStore16_x86(9DKI) 243
ioRead16_x86(9DKI) 243
ioWrite16_x86(9DKI) 243
ioLoad32_x86(9DKI) 243
ioStore32_x86(9DKI) 243
ioRead32_x86(9DKI) 243
ioWrite32_x86(9DKI) 243
ioLoad8(9DKI) 245
ioStore8(9DKI) 245
ioRead8(9DKI) 245
ioWrite8(9DKI) 245
ioLoad16(9DKI) 245
ioStore16(9DKI) 245
ioRead16(9DKI) 245
ioWrite16(9DKI) 245
ioLoad32(9DKI) 245
ioStore32(9DKI) 245
ioRead32(9DKI) 245
ioWrite32(9DKI) 245
ioLoad8_x86(9DKI) 246
ioStore8_x86(9DKI) 246
ioRead8_x86(9DKI) 246
ioWrite8_x86(9DKI) 246
ioLoad16_x86(9DKI) 246
ioStore16_x86(9DKI) 246
ioRead16_x86(9DKI) 246
ioWrite16_x86(9DKI) 246

ioLoad32_x86(9DKI)	246
ioStore32_x86(9DKI)	246
ioRead32_x86(9DKI)	246
ioWrite32_x86(9DKI)	246
ioLoad8(9DKI)	248
ioStore8(9DKI)	248
ioRead8(9DKI)	248
ioWrite8(9DKI)	248
ioLoad16(9DKI)	248
ioStore16(9DKI)	248
ioRead16(9DKI)	248
ioWrite16(9DKI)	248
ioLoad32(9DKI)	248
ioStore32(9DKI)	248
ioRead32(9DKI)	248
ioWrite32(9DKI)	248
ioLoad8_x86(9DKI)	249
ioStore8_x86(9DKI)	249
ioRead8_x86(9DKI)	249
ioWrite8_x86(9DKI)	249
ioLoad16_x86(9DKI)	249
ioStore16_x86(9DKI)	249
ioRead16_x86(9DKI)	249
ioWrite16_x86(9DKI)	249
ioLoad32_x86(9DKI)	249
ioStore32_x86(9DKI)	249
ioRead32_x86(9DKI)	249
ioWrite32_x86(9DKI)	249

ioLoad8(9DKI) 251
ioStore8(9DKI) 251
ioRead8(9DKI) 251
ioWrite8(9DKI) 251
ioLoad16(9DKI) 251
ioStore16(9DKI) 251
ioRead16(9DKI) 251
ioWrite16(9DKI) 251
ioLoad32(9DKI) 251
ioStore32(9DKI) 251
ioRead32(9DKI) 251
ioWrite32(9DKI) 251
ioLoad8_x86(9DKI) 252
ioStore8_x86(9DKI) 252
ioRead8_x86(9DKI) 252
ioWrite8_x86(9DKI) 252
ioLoad16_x86(9DKI) 252
ioStore16_x86(9DKI) 252
ioRead16_x86(9DKI) 252
ioWrite16_x86(9DKI) 252
ioLoad32_x86(9DKI) 252
ioStore32_x86(9DKI) 252
ioRead32_x86(9DKI) 252
ioWrite32_x86(9DKI) 252
ioLoad8(9DKI) 254
ioStore8(9DKI) 254
ioRead8(9DKI) 254
ioWrite8(9DKI) 254

ioLoad16(9DKI)	254
ioStore16(9DKI)	254
ioRead16(9DKI)	254
ioWrite16(9DKI)	254
ioLoad32(9DKI)	254
ioStore32(9DKI)	254
ioRead32(9DKI)	254
ioWrite32(9DKI)	254
ioLoad8_x86(9DKI)	255
ioStore8_x86(9DKI)	255
ioRead8_x86(9DKI)	255
ioWrite8_x86(9DKI)	255
ioLoad16_x86(9DKI)	255
ioStore16_x86(9DKI)	255
ioRead16_x86(9DKI)	255
ioWrite16_x86(9DKI)	255
ioLoad32_x86(9DKI)	255
ioStore32_x86(9DKI)	255
ioRead32_x86(9DKI)	255
ioWrite32_x86(9DKI)	255
ioLoad8(9DKI)	257
ioStore8(9DKI)	257
ioRead8(9DKI)	257
ioWrite8(9DKI)	257
ioLoad16(9DKI)	257
ioStore16(9DKI)	257
ioRead16(9DKI)	257
ioWrite16(9DKI)	257

ioLoad32(9DKI) 257
ioStore32(9DKI) 257
ioRead32(9DKI) 257
ioWrite32(9DKI) 257
ioLoad8_x86(9DKI) 258
ioStore8_x86(9DKI) 258
ioRead8_x86(9DKI) 258
ioWrite8_x86(9DKI) 258
ioLoad16_x86(9DKI) 258
ioStore16_x86(9DKI) 258
ioRead16_x86(9DKI) 258
ioWrite16_x86(9DKI) 258
ioLoad32_x86(9DKI) 258
ioStore32_x86(9DKI) 258
ioRead32_x86(9DKI) 258
ioWrite32_x86(9DKI) 258
ioLoad8(9DKI) 260
ioStore8(9DKI) 260
ioRead8(9DKI) 260
ioWrite8(9DKI) 260
ioLoad16(9DKI) 260
ioStore16(9DKI) 260
ioRead16(9DKI) 260
ioWrite16(9DKI) 260
ioLoad32(9DKI) 260
ioStore32(9DKI) 260
ioRead32(9DKI) 260
ioWrite32(9DKI) 260

ioLoad8_x86(9DKI)	261
ioStore8_x86(9DKI)	261
ioRead8_x86(9DKI)	261
ioWrite8_x86(9DKI)	261
ioLoad16_x86(9DKI)	261
ioStore16_x86(9DKI)	261
ioRead16_x86(9DKI)	261
ioWrite16_x86(9DKI)	261
ioLoad32_x86(9DKI)	261
ioStore32_x86(9DKI)	261
ioRead32_x86(9DKI)	261
ioWrite32_x86(9DKI)	261
ioLoad8(9DKI)	263
ioStore8(9DKI)	263
ioRead8(9DKI)	263
ioWrite8(9DKI)	263
ioLoad16(9DKI)	263
ioStore16(9DKI)	263
ioRead16(9DKI)	263
ioWrite16(9DKI)	263
ioLoad32(9DKI)	263
ioStore32(9DKI)	263
ioRead32(9DKI)	263
ioWrite32(9DKI)	263
ioLoad8_x86(9DKI)	264
ioStore8_x86(9DKI)	264
ioRead8_x86(9DKI)	264
ioWrite8_x86(9DKI)	264

ioLoad16_x86(9DKI) 264
ioStore16_x86(9DKI) 264
ioRead16_x86(9DKI) 264
ioWrite16_x86(9DKI) 264
ioLoad32_x86(9DKI) 264
ioStore32_x86(9DKI) 264
ioRead32_x86(9DKI) 264
ioWrite32_x86(9DKI) 264
ioLoad8(9DKI) 266
ioStore8(9DKI) 266
ioRead8(9DKI) 266
ioWrite8(9DKI) 266
ioLoad16(9DKI) 266
ioStore16(9DKI) 266
ioRead16(9DKI) 266
ioWrite16(9DKI) 266
ioLoad32(9DKI) 266
ioStore32(9DKI) 266
ioRead32(9DKI) 266
ioWrite32(9DKI) 266
ioLoad8_x86(9DKI) 267
ioStore8_x86(9DKI) 267
ioRead8_x86(9DKI) 267
ioWrite8_x86(9DKI) 267
ioLoad16_x86(9DKI) 267
ioStore16_x86(9DKI) 267
ioRead16_x86(9DKI) 267
ioWrite16_x86(9DKI) 267

ioLoad32_x86(9DKI)	267
ioStore32_x86(9DKI)	267
ioRead32_x86(9DKI)	267
ioWrite32_x86(9DKI)	267
loadSwap_16(9DKI)	269
storeSwap_16(9DKI)	269
swap_16(9DKI)	269
loadSwap_32(9DKI)	269
storeSwap_32(9DKI)	269
swap_32(9DKI)	269
loadSwap_64(9DKI)	269
storeSwap_64(9DKI)	269
swap_64(9DKI)	269
loadSwap_16(9DKI)	271
storeSwap_16(9DKI)	271
swap_16(9DKI)	271
loadSwap_32(9DKI)	271
storeSwap_32(9DKI)	271
swap_32(9DKI)	271
loadSwap_64(9DKI)	271
storeSwap_64(9DKI)	271
swap_64(9DKI)	271
loadSwap_16(9DKI)	273
storeSwap_16(9DKI)	273
swap_16(9DKI)	273
loadSwap_32(9DKI)	273
storeSwap_32(9DKI)	273
swap_32(9DKI)	273

loadSwap_64(9DKI) 273
storeSwap_64(9DKI) 273
swap_64(9DKI) 273
loadSwapEieio_16(9DKI) 275
storeSwapEieio_16(9DKI) 275
swapEieio_16(9DKI) 275
loadSwapEieio_32(9DKI) 275
storeSwapEieio_32(9DKI) 275
swapEieio_32(9DKI) 275
eieio(9DKI) 275
loadSwapEieio_16_powerpc(9DKI) 276
loadSwapSync_16_powerpc(9DKI) 276
storeSwapEieio_16_powerpc(9DKI) 276
storeSwapSync_16_powerpc(9DKI) 276
swapEieio_16_powerpc(9DKI) 276
loadSwapEieio_32_powerpc(9DKI) 276
loadSwapSync_32_powerpc(9DKI) 276
storeSwapEieio_32_powerpc(9DKI) 276
storeSwapSync_32_powerpc(9DKI) 276
swapEieio_32_powerpc(9DKI) 276
eieio_powerpc(9DKI) 276
ioSync_powerpc(9DKI) 276
loadSwapEieio_16(9DKI) 279
storeSwapEieio_16(9DKI) 279
swapEieio_16(9DKI) 279
loadSwapEieio_32(9DKI) 279
storeSwapEieio_32(9DKI) 279
swapEieio_32(9DKI) 279

eieio(9DKI)	279
loadSwapEieio_16_powerpc(9DKI)	280
loadSwapSync_16_powerpc(9DKI)	280
storeSwapEieio_16_powerpc(9DKI)	280
storeSwapSync_16_powerpc(9DKI)	280
swapEieio_16_powerpc(9DKI)	280
loadSwapEieio_32_powerpc(9DKI)	280
loadSwapSync_32_powerpc(9DKI)	280
storeSwapEieio_32_powerpc(9DKI)	280
storeSwapSync_32_powerpc(9DKI)	280
swapEieio_32_powerpc(9DKI)	280
eieio_powerpc(9DKI)	280
ioSync_powerpc(9DKI)	280
load_sync_8_usparc(9DKI)	283
store_sync_8_usparc(9DKI)	283
loadSwap_sync_16_usparc(9DKI)	283
storeSwap_sync_16_usparc(9DKI)	283
load_sync_16_usparc(9DKI)	283
store_sync_16_usparc(9DKI)	283
loadSwap_sync_32_usparc(9DKI)	283
storeSwap_sync_32_usparc(9DKI)	283
load_sync_32_usparc(9DKI)	283
store_sync_32_usparc(9DKI)	283
loadSwap_sync_64_usparc(9DKI)	283
storeSwap_sync_64_usparc(9DKI)	283
load_sync_64_usparc(9DKI)	283
store_sync_64_usparc(9DKI)	283
load_sync_8_usparc(9DKI)	286

store_sync_8_usparc(9DKI) 286
loadSwap_sync_16_usparc(9DKI) 286
storeSwap_sync_16_usparc(9DKI) 286
load_sync_16_usparc(9DKI) 286
store_sync_16_usparc(9DKI) 286
loadSwap_sync_32_usparc(9DKI) 286
storeSwap_sync_32_usparc(9DKI) 286
load_sync_32_usparc(9DKI) 286
store_sync_32_usparc(9DKI) 286
loadSwap_sync_64_usparc(9DKI) 286
storeSwap_sync_64_usparc(9DKI) 286
load_sync_64_usparc(9DKI) 286
store_sync_64_usparc(9DKI) 286
load_sync_8_usparc(9DKI) 289
store_sync_8_usparc(9DKI) 289
loadSwap_sync_16_usparc(9DKI) 289
storeSwap_sync_16_usparc(9DKI) 289
load_sync_16_usparc(9DKI) 289
store_sync_16_usparc(9DKI) 289
loadSwap_sync_32_usparc(9DKI) 289
storeSwap_sync_32_usparc(9DKI) 289
load_sync_32_usparc(9DKI) 289
store_sync_32_usparc(9DKI) 289
loadSwap_sync_64_usparc(9DKI) 289
storeSwap_sync_64_usparc(9DKI) 289
load_sync_64_usparc(9DKI) 289
store_sync_64_usparc(9DKI) 289
load_sync_8_usparc(9DKI) 292

store_sync_8_usparc(9DKI)	292
loadSwap_sync_16_usparc(9DKI)	292
storeSwap_sync_16_usparc(9DKI)	292
load_sync_16_usparc(9DKI)	292
store_sync_16_usparc(9DKI)	292
loadSwap_sync_32_usparc(9DKI)	292
storeSwap_sync_32_usparc(9DKI)	292
load_sync_32_usparc(9DKI)	292
store_sync_32_usparc(9DKI)	292
loadSwap_sync_64_usparc(9DKI)	292
storeSwap_sync_64_usparc(9DKI)	292
load_sync_64_usparc(9DKI)	292
store_sync_64_usparc(9DKI)	292
load_sync_8_usparc(9DKI)	295
store_sync_8_usparc(9DKI)	295
loadSwap_sync_16_usparc(9DKI)	295
storeSwap_sync_16_usparc(9DKI)	295
load_sync_16_usparc(9DKI)	295
store_sync_16_usparc(9DKI)	295
loadSwap_sync_32_usparc(9DKI)	295
storeSwap_sync_32_usparc(9DKI)	295
load_sync_32_usparc(9DKI)	295
store_sync_32_usparc(9DKI)	295
loadSwap_sync_64_usparc(9DKI)	295
storeSwap_sync_64_usparc(9DKI)	295
load_sync_64_usparc(9DKI)	295
store_sync_64_usparc(9DKI)	295
load_sync_8_usparc(9DKI)	298

store_sync_8_usparc(9DKI) 298
loadSwap_sync_16_usparc(9DKI) 298
storeSwap_sync_16_usparc(9DKI) 298
load_sync_16_usparc(9DKI) 298
store_sync_16_usparc(9DKI) 298
loadSwap_sync_32_usparc(9DKI) 298
storeSwap_sync_32_usparc(9DKI) 298
load_sync_32_usparc(9DKI) 298
store_sync_32_usparc(9DKI) 298
loadSwap_sync_64_usparc(9DKI) 298
storeSwap_sync_64_usparc(9DKI) 298
load_sync_64_usparc(9DKI) 298
store_sync_64_usparc(9DKI) 298
load_sync_8_usparc(9DKI) 301
store_sync_8_usparc(9DKI) 301
loadSwap_sync_16_usparc(9DKI) 301
storeSwap_sync_16_usparc(9DKI) 301
load_sync_16_usparc(9DKI) 301
store_sync_16_usparc(9DKI) 301
loadSwap_sync_32_usparc(9DKI) 301
storeSwap_sync_32_usparc(9DKI) 301
load_sync_32_usparc(9DKI) 301
store_sync_32_usparc(9DKI) 301
loadSwap_sync_64_usparc(9DKI) 301
storeSwap_sync_64_usparc(9DKI) 301
load_sync_64_usparc(9DKI) 301
store_sync_64_usparc(9DKI) 301
loadSwap_16(9DKI) 304

storeSwap_16(9DKI)	304
swap_16(9DKI)	304
loadSwap_32(9DKI)	304
storeSwap_32(9DKI)	304
swap_32(9DKI)	304
loadSwap_64(9DKI)	304
storeSwap_64(9DKI)	304
swap_64(9DKI)	304
loadSwap_16(9DKI)	306
storeSwap_16(9DKI)	306
swap_16(9DKI)	306
loadSwap_32(9DKI)	306
storeSwap_32(9DKI)	306
swap_32(9DKI)	306
loadSwap_64(9DKI)	306
storeSwap_64(9DKI)	306
swap_64(9DKI)	306
loadSwap_16(9DKI)	308
storeSwap_16(9DKI)	308
swap_16(9DKI)	308
loadSwap_32(9DKI)	308
storeSwap_32(9DKI)	308
swap_32(9DKI)	308
loadSwap_64(9DKI)	308
storeSwap_64(9DKI)	308
swap_64(9DKI)	308
loadSwapEieio_16(9DKI)	310
storeSwapEieio_16(9DKI)	310

swapEieio_16(9DKI) 310
loadSwapEieio_32(9DKI) 310
storeSwapEieio_32(9DKI) 310
swapEieio_32(9DKI) 310
eieio(9DKI) 310
loadSwapEieio_16_powerpc(9DKI) 311
loadSwapSync_16_powerpc(9DKI) 311
storeSwapEieio_16_powerpc(9DKI) 311
storeSwapSync_16_powerpc(9DKI) 311
swapEieio_16_powerpc(9DKI) 311
loadSwapEieio_32_powerpc(9DKI) 311
loadSwapSync_32_powerpc(9DKI) 311
storeSwapEieio_32_powerpc(9DKI) 311
storeSwapSync_32_powerpc(9DKI) 311
swapEieio_32_powerpc(9DKI) 311
eieio_powerpc(9DKI) 311
ioSync_powerpc(9DKI) 311
loadSwapEieio_16(9DKI) 314
storeSwapEieio_16(9DKI) 314
swapEieio_16(9DKI) 314
loadSwapEieio_32(9DKI) 314
storeSwapEieio_32(9DKI) 314
swapEieio_32(9DKI) 314
eieio(9DKI) 314
loadSwapEieio_16_powerpc(9DKI) 315
loadSwapSync_16_powerpc(9DKI) 315
storeSwapEieio_16_powerpc(9DKI) 315
storeSwapSync_16_powerpc(9DKI) 315

swapEieio_16_powerpc(9DKI)	315
loadSwapEieio_32_powerpc(9DKI)	315
loadSwapSync_32_powerpc(9DKI)	315
storeSwapEieio_32_powerpc(9DKI)	315
storeSwapSync_32_powerpc(9DKI)	315
swapEieio_32_powerpc(9DKI)	315
eieio_powerpc(9DKI)	315
ioSync_powerpc(9DKI)	315
load_sync_8_usparc(9DKI)	318
store_sync_8_usparc(9DKI)	318
loadSwap_sync_16_usparc(9DKI)	318
storeSwap_sync_16_usparc(9DKI)	318
load_sync_16_usparc(9DKI)	318
store_sync_16_usparc(9DKI)	318
loadSwap_sync_32_usparc(9DKI)	318
storeSwap_sync_32_usparc(9DKI)	318
load_sync_32_usparc(9DKI)	318
store_sync_32_usparc(9DKI)	318
loadSwap_sync_64_usparc(9DKI)	318
storeSwap_sync_64_usparc(9DKI)	318
load_sync_64_usparc(9DKI)	318
store_sync_64_usparc(9DKI)	318
load_sync_8_usparc(9DKI)	321
store_sync_8_usparc(9DKI)	321
loadSwap_sync_16_usparc(9DKI)	321
storeSwap_sync_16_usparc(9DKI)	321
load_sync_16_usparc(9DKI)	321
store_sync_16_usparc(9DKI)	321

loadSwap_sync_32_usparc(9DKI) 321
storeSwap_sync_32_usparc(9DKI) 321
load_sync_32_usparc(9DKI) 321
store_sync_32_usparc(9DKI) 321
loadSwap_sync_64_usparc(9DKI) 321
storeSwap_sync_64_usparc(9DKI) 321
load_sync_64_usparc(9DKI) 321
store_sync_64_usparc(9DKI) 321
load_sync_8_usparc(9DKI) 324
store_sync_8_usparc(9DKI) 324
loadSwap_sync_16_usparc(9DKI) 324
storeSwap_sync_16_usparc(9DKI) 324
load_sync_16_usparc(9DKI) 324
store_sync_16_usparc(9DKI) 324
loadSwap_sync_32_usparc(9DKI) 324
storeSwap_sync_32_usparc(9DKI) 324
load_sync_32_usparc(9DKI) 324
store_sync_32_usparc(9DKI) 324
loadSwap_sync_64_usparc(9DKI) 324
storeSwap_sync_64_usparc(9DKI) 324
load_sync_64_usparc(9DKI) 324
store_sync_64_usparc(9DKI) 324
load_sync_8_usparc(9DKI) 327
store_sync_8_usparc(9DKI) 327
loadSwap_sync_16_usparc(9DKI) 327
storeSwap_sync_16_usparc(9DKI) 327
load_sync_16_usparc(9DKI) 327
store_sync_16_usparc(9DKI) 327

loadSwap_sync_32_usparc(9DKI)	327
storeSwap_sync_32_usparc(9DKI)	327
load_sync_32_usparc(9DKI)	327
store_sync_32_usparc(9DKI)	327
loadSwap_sync_64_usparc(9DKI)	327
storeSwap_sync_64_usparc(9DKI)	327
load_sync_64_usparc(9DKI)	327
store_sync_64_usparc(9DKI)	327
load_sync_8_usparc(9DKI)	330
store_sync_8_usparc(9DKI)	330
loadSwap_sync_16_usparc(9DKI)	330
storeSwap_sync_16_usparc(9DKI)	330
load_sync_16_usparc(9DKI)	330
store_sync_16_usparc(9DKI)	330
loadSwap_sync_32_usparc(9DKI)	330
storeSwap_sync_32_usparc(9DKI)	330
load_sync_32_usparc(9DKI)	330
store_sync_32_usparc(9DKI)	330
loadSwap_sync_64_usparc(9DKI)	330
storeSwap_sync_64_usparc(9DKI)	330
load_sync_64_usparc(9DKI)	330
store_sync_64_usparc(9DKI)	330
load_sync_8_usparc(9DKI)	333
store_sync_8_usparc(9DKI)	333
loadSwap_sync_16_usparc(9DKI)	333
storeSwap_sync_16_usparc(9DKI)	333
load_sync_16_usparc(9DKI)	333
store_sync_16_usparc(9DKI)	333

loadSwap_sync_32_usparc(9DKI) 333
storeSwap_sync_32_usparc(9DKI) 333
load_sync_32_usparc(9DKI) 333
store_sync_32_usparc(9DKI) 333
loadSwap_sync_64_usparc(9DKI) 333
storeSwap_sync_64_usparc(9DKI) 333
load_sync_64_usparc(9DKI) 333
store_sync_64_usparc(9DKI) 333
load_sync_8_usparc(9DKI) 336
store_sync_8_usparc(9DKI) 336
loadSwap_sync_16_usparc(9DKI) 336
storeSwap_sync_16_usparc(9DKI) 336
load_sync_16_usparc(9DKI) 336
store_sync_16_usparc(9DKI) 336
loadSwap_sync_32_usparc(9DKI) 336
storeSwap_sync_32_usparc(9DKI) 336
load_sync_32_usparc(9DKI) 336
store_sync_32_usparc(9DKI) 336
loadSwap_sync_64_usparc(9DKI) 336
storeSwap_sync_64_usparc(9DKI) 336
load_sync_64_usparc(9DKI) 336
store_sync_64_usparc(9DKI) 336
svAsyncExcepAttach(9DKI) 339
svAsyncExcepDetach(9DKI) 339
svAsyncExcepAttach_usparc(9DKI) 340
svAsyncExcepDetach_usparc(9DKI) 340
svAsyncExcepAttach(9DKI) 343
svAsyncExcepDetach(9DKI) 343

svAsyncExcepAttach_usparc(9DKI)	344
svAsyncExcepDetach_usparc(9DKI)	344
svDeviceRegister(9DKI)	347
svDeviceAlloc(9DKI)	347
svDeviceFree(9DKI)	347
svDeviceUnregister(9DKI)	347
svDeviceEvent(9DKI)	347
svDeviceLookup(9DKI)	347
svDeviceEntry(9DKI)	347
svDeviceRelease(9DKI)	347
svDeviceRegister(9DKI)	354
svDeviceAlloc(9DKI)	354
svDeviceFree(9DKI)	354
svDeviceUnregister(9DKI)	354
svDeviceEvent(9DKI)	354
svDeviceLookup(9DKI)	354
svDeviceEntry(9DKI)	354
svDeviceRelease(9DKI)	354
svDeviceRegister(9DKI)	361
svDeviceAlloc(9DKI)	361
svDeviceFree(9DKI)	361
svDeviceUnregister(9DKI)	361
svDeviceEvent(9DKI)	361
svDeviceLookup(9DKI)	361
svDeviceEntry(9DKI)	361
svDeviceRelease(9DKI)	361
svDeviceRegister(9DKI)	368
svDeviceAlloc(9DKI)	368

svDeviceFree(9DKI) 368
svDeviceUnregister(9DKI) 368
svDeviceEvent(9DKI) 368
svDeviceLookup(9DKI) 368
svDeviceEntry(9DKI) 368
svDeviceRelease(9DKI) 368
svDeviceRegister(9DKI) 375
svDeviceAlloc(9DKI) 375
svDeviceFree(9DKI) 375
svDeviceUnregister(9DKI) 375
svDeviceEvent(9DKI) 375
svDeviceLookup(9DKI) 375
svDeviceEntry(9DKI) 375
svDeviceRelease(9DKI) 375
svDeviceRegister(9DKI) 382
svDeviceAlloc(9DKI) 382
svDeviceFree(9DKI) 382
svDeviceUnregister(9DKI) 382
svDeviceEvent(9DKI) 382
svDeviceLookup(9DKI) 382
svDeviceEntry(9DKI) 382
svDeviceRelease(9DKI) 382
svDeviceRegister(9DKI) 389
svDeviceAlloc(9DKI) 389
svDeviceFree(9DKI) 389
svDeviceUnregister(9DKI) 389
svDeviceEvent(9DKI) 389
svDeviceLookup(9DKI) 389

svDeviceEntry(9DKI)	389
svDeviceRelease(9DKI)	389
svDeviceRegister(9DKI)	396
svDeviceAlloc(9DKI)	396
svDeviceFree(9DKI)	396
svDeviceUnregister(9DKI)	396
svDeviceEvent(9DKI)	396
svDeviceLookup(9DKI)	396
svDeviceEntry(9DKI)	396
svDeviceRelease(9DKI)	396
svDkiOpen(9DKI)	403
svDkiClose(9DKI)	403
svDkiEvent(9DKI)	403
svDkiOpen(9DKI)	406
svDkiClose(9DKI)	406
svDkiEvent(9DKI)	406
svDkiOpen(9DKI)	409
svDkiClose(9DKI)	409
svDkiEvent(9DKI)	409
svDkiThreadCall(9DKI)	412
svDkiThreadTrigger(9DKI)	412
svDkiThreadCall(9DKI)	414
svDkiThreadTrigger(9DKI)	414
svDriverRegister(9DKI)	416
svDriverLookupFirst(9DKI)	416
svDriverLookupNext(9DKI)	416
svDriverRelease(9DKI)	416
svDriverEntry(9DKI)	416

svDriverCap(9DKI) 416
svDriverUnregister(9DKI) 416
svDriverRegister(9DKI) 422
svDriverLookupFirst(9DKI) 422
svDriverLookupNext(9DKI) 422
svDriverRelease(9DKI) 422
svDriverEntry(9DKI) 422
svDriverCap(9DKI) 422
svDriverUnregister(9DKI) 422
svDriverRegister(9DKI) 428
svDriverLookupFirst(9DKI) 428
svDriverLookupNext(9DKI) 428
svDriverRelease(9DKI) 428
svDriverEntry(9DKI) 428
svDriverCap(9DKI) 428
svDriverUnregister(9DKI) 428
svDriverRegister(9DKI) 434
svDriverLookupFirst(9DKI) 434
svDriverLookupNext(9DKI) 434
svDriverRelease(9DKI) 434
svDriverEntry(9DKI) 434
svDriverCap(9DKI) 434
svDriverUnregister(9DKI) 434
svDriverRegister(9DKI) 440
svDriverLookupFirst(9DKI) 440
svDriverLookupNext(9DKI) 440
svDriverRelease(9DKI) 440
svDriverEntry(9DKI) 440

svDriverCap(9DKI)	440
svDriverUnregister(9DKI)	440
svDriverRegister(9DKI)	446
svDriverLookupFirst(9DKI)	446
svDriverLookupNext(9DKI)	446
svDriverRelease(9DKI)	446
svDriverEntry(9DKI)	446
svDriverCap(9DKI)	446
svDriverUnregister(9DKI)	446
svDriverRegister(9DKI)	452
svDriverLookupFirst(9DKI)	452
svDriverLookupNext(9DKI)	452
svDriverRelease(9DKI)	452
svDriverEntry(9DKI)	452
svDriverCap(9DKI)	452
svDriverUnregister(9DKI)	452
svIntrAttach(9DKI)	458
svIntrDetach(9DKI)	458
svSoftIntrAttach(9DKI)	458
svSoftIntrDetach(9DKI)	458
svTimerIntrAttach(9DKI)	458
svTimerIntrDetach(9DKI)	458
svIntrCtxGet(9DKI)	458
svIntrAttach_powerpc(9DKI)	459
svIntrDetach_powerpc(9DKI)	459
svIntrCtxGet_powerpc(9DKI)	459
svIntrAttach_usparc(9DKI)	462
svIntrDetach_usparc(9DKI)	462

svSoftIntrAttach_usparc(9DKI) 462
svSoftIntrDetach_usparc(9DKI) 462
svTimerIntrAttach_uparc(9DKI) 462
svTimerIntrDetach_uparc(9DKI) 462
svIntrCtxGet_usparc(9DKI) 462
svIntrAttach_x86(9DKI) 468
svIntrDetach_x86(9DKI) 468
svIntrCtxGet_x86(9DKI) 468
svIntrAttach(9DKI) 472
svIntrDetach(9DKI) 472
svSoftIntrAttach(9DKI) 472
svSoftIntrDetach(9DKI) 472
svTimerIntrAttach(9DKI) 472
svTimerIntrDetach(9DKI) 472
svIntrCtxGet(9DKI) 472
svIntrAttach_powerpc(9DKI) 473
svIntrDetach_powerpc(9DKI) 473
svIntrCtxGet_powerpc(9DKI) 473
svIntrAttach_usparc(9DKI) 476
svIntrDetach_usparc(9DKI) 476
svSoftIntrAttach_usparc(9DKI) 476
svSoftIntrDetach_usparc(9DKI) 476
svTimerIntrAttach_uparc(9DKI) 476
svTimerIntrDetach_uparc(9DKI) 476
svIntrCtxGet_usparc(9DKI) 476
svIntrAttach_x86(9DKI) 482
svIntrDetach_x86(9DKI) 482
svIntrCtxGet_x86(9DKI) 482

svIntrAttach(9DKI) 486
svIntrDetach(9DKI) 486
svSoftIntrAttach(9DKI) 486
svSoftIntrDetach(9DKI) 486
svTimerIntrAttach(9DKI) 486
svTimerIntrDetach(9DKI) 486
svIntrCtxGet(9DKI) 486
svIntrAttach_powerpc(9DKI) 487
svIntrDetach_powerpc(9DKI) 487
svIntrCtxGet_powerpc(9DKI) 487
svIntrAttach_usparc(9DKI) 490
svIntrDetach_usparc(9DKI) 490
svSoftIntrAttach_usparc(9DKI) 490
svSoftIntrDetach_usparc(9DKI) 490
svTimerIntrAttach_usparc(9DKI) 490
svTimerIntrDetach_usparc(9DKI) 490
svIntrCtxGet_usparc(9DKI) 490
svIntrAttach_x86(9DKI) 496
svIntrDetach_x86(9DKI) 496
svIntrCtxGet_x86(9DKI) 496
svMemAlloc(9DKI) 500
svMemFree(9DKI) 500
svMemAlloc(9DKI) 502
svMemFree(9DKI) 502
svPhysAlloc(9DKI) 504
svPhysFree(9DKI) 504
svPhysAlloc(9DKI) 507
svPhysFree(9DKI) 507

svPhysMap(9DKI) 510
svPhysUnmap(9DKI) 510
vmMapToPhys(9DKI) 510
svPhysMap_powerpc(9DKI) 511
svPhysUnmap_powerpc(9DKI) 511
vmMapToPhys_powerpc(9DKI) 511
svPhysMap_usparc(9DKI) 515
svPhysUnmap_usparc(9DKI) 515
vmMapToPhys_usparc(9DKI) 515
svPhysMap_x86(9DKI) 519
svPhysUnmap_x86(9DKI) 519
vmMapToPhys_x86(9DKI) 519
svPhysMap(9DKI) 522
svPhysUnmap(9DKI) 522
vmMapToPhys(9DKI) 522
svPhysMap_powerpc(9DKI) 523
svPhysUnmap_powerpc(9DKI) 523
vmMapToPhys_powerpc(9DKI) 523
svPhysMap_usparc(9DKI) 527
svPhysUnmap_usparc(9DKI) 527
vmMapToPhys_usparc(9DKI) 527
svPhysMap_x86(9DKI) 531
svPhysUnmap_x86(9DKI) 531
vmMapToPhys_x86(9DKI) 531
svIntrAttach_usparc(9DKI) 534
svIntrDetach_usparc(9DKI) 534
svSoftIntrAttach_usparc(9DKI) 534
svSoftIntrDetach_usparc(9DKI) 534

svTimerIntrAttach_uparc(9DKI)	534
svTimerIntrDetach_uparc(9DKI)	534
svIntrCtxGet_usparc(9DKI)	534
svIntrAttach_uparc(9DKI)	540
svIntrDetach_uparc(9DKI)	540
svSoftIntrAttach_uparc(9DKI)	540
svSoftIntrDetach_uparc(9DKI)	540
svTimerIntrAttach_uparc(9DKI)	540
svTimerIntrDetach_uparc(9DKI)	540
svIntrCtxGet_usparc(9DKI)	540
svTimeoutSet(9DKI)	546
svTimeoutCancel(9DKI)	546
svTimeoutGetRes(9DKI)	546
svTimeoutSet(9DKI)	548
svTimeoutCancel(9DKI)	548
svTimeoutGetRes(9DKI)	548
svTimeoutSet(9DKI)	550
svTimeoutCancel(9DKI)	550
svTimeoutGetRes(9DKI)	550
loadSwap_16(9DKI)	552
storeSwap_16(9DKI)	552
swap_16(9DKI)	552
loadSwap_32(9DKI)	552
storeSwap_32(9DKI)	552
swap_32(9DKI)	552
loadSwap_64(9DKI)	552
storeSwap_64(9DKI)	552
swap_64(9DKI)	552

loadSwap_16(9DKI) 554
storeSwap_16(9DKI) 554
swap_16(9DKI) 554
loadSwap_32(9DKI) 554
storeSwap_32(9DKI) 554
swap_32(9DKI) 554
loadSwap_64(9DKI) 554
storeSwap_64(9DKI) 554
swap_64(9DKI) 554
loadSwap_16(9DKI) 556
storeSwap_16(9DKI) 556
swap_16(9DKI) 556
loadSwap_32(9DKI) 556
storeSwap_32(9DKI) 556
swap_32(9DKI) 556
loadSwap_64(9DKI) 556
storeSwap_64(9DKI) 556
swap_64(9DKI) 556
loadSwapEieio_16(9DKI) 558
storeSwapEieio_16(9DKI) 558
swapEieio_16(9DKI) 558
loadSwapEieio_32(9DKI) 558
storeSwapEieio_32(9DKI) 558
swapEieio_32(9DKI) 558
eieio(9DKI) 558
loadSwapEieio_16_powerpc(9DKI) 559
loadSwapSync_16_powerpc(9DKI) 559
storeSwapEieio_16_powerpc(9DKI) 559

storeSwapSync_16_powerpc(9DKI)	559
swapEieio_16_powerpc(9DKI)	559
loadSwapEieio_32_powerpc(9DKI)	559
loadSwapSync_32_powerpc(9DKI)	559
storeSwapEieio_32_powerpc(9DKI)	559
storeSwapSync_32_powerpc(9DKI)	559
swapEieio_32_powerpc(9DKI)	559
eieio_powerpc(9DKI)	559
ioSync_powerpc(9DKI)	559
loadSwapEieio_16(9DKI)	562
storeSwapEieio_16(9DKI)	562
swapEieio_16(9DKI)	562
loadSwapEieio_32(9DKI)	562
storeSwapEieio_32(9DKI)	562
swapEieio_32(9DKI)	562
eieio(9DKI)	562
loadSwapEieio_16_powerpc(9DKI)	563
loadSwapSync_16_powerpc(9DKI)	563
storeSwapEieio_16_powerpc(9DKI)	563
storeSwapSync_16_powerpc(9DKI)	563
swapEieio_16_powerpc(9DKI)	563
loadSwapEieio_32_powerpc(9DKI)	563
loadSwapSync_32_powerpc(9DKI)	563
storeSwapEieio_32_powerpc(9DKI)	563
storeSwapSync_32_powerpc(9DKI)	563
swapEieio_32_powerpc(9DKI)	563
eieio_powerpc(9DKI)	563
ioSync_powerpc(9DKI)	563

usecBusyWait(9DKI) 566
svPhysMap(9DKI) 567
svPhysUnmap(9DKI) 567
vmMapToPhys(9DKI) 567
svPhysMap_powerpc(9DKI) 568
svPhysUnmap_powerpc(9DKI) 568
vmMapToPhys_powerpc(9DKI) 568
svPhysMap_usparc(9DKI) 572
svPhysUnmap_usparc(9DKI) 572
vmMapToPhys_usparc(9DKI) 572
svPhysMap_x86(9DKI) 576
svPhysUnmap_x86(9DKI) 576
vmMapToPhys_x86(9DKI) 576
Index 578

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

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.

Driver to Kernel Interface

NAME	Intro – driver kernel interface introduction				
SYNOPSIS	<pre>#include <dki/dki.h> #include <dki/f_dki.h></pre>				
FEATURES	DKI				
DESCRIPTION	Provides Driver/Kernel interface services				
EXTENDED DESCRIPTION	<p>The Drivers/Kernel Interface or DKI, defines all services provided by the ChorusOS microkernel in order to write driver components.</p> <p>The DKI is composed of:</p> <ul style="list-style-type: none"> ■ A "Common DKI API", which defines services common to all platforms and processors, and which are usable by all kinds of drivers, notwithstanding the layer level. ■ Various "Processor family specific DKI" APIs, which are defined and available only for a given processor family and which should be used only by the lowest-level drivers. That is, drivers for busses and devices which are directly connected to the processor's local bus. <p>Currently, "Processor family specific DKI" APIs are defined for the following processors:</p> <ul style="list-style-type: none"> ■ PowerPC 60x (603x, 604x, MPC750) ■ Intel (ix86 ...) 				
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">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				
SEE ALSO	<p>Common DKI:</p> <p><code>svDkiOpen(9DKI)</code>, <code>dreeNodeRoot(9DKI)</code>, <code>svDriverRegister(9DKI)</code>, <code>svDeviceRegister(9DKI)</code>, <code>svMemAlloc(9DKI)</code>, <code>svPhysAlloc(9DKI)</code>, <code>svPhysMap(9DKI)</code>, <code>svDkiThreadCall(9DKI)</code>, <code>svTimeoutSet(9DKI)</code>, <code>usecBusyWait(9DKI)</code>, <code>loadSwap(9DKI)</code>, <code>DISABLE_PREEMPT(9DKI)</code>, <code>imsIntrMask_f(9DKI)</code></p> <p>PowerPC DKI:</p> <p><code>loadSwapEieio_16(9DKI)</code>, <code>svIntrAttach(9DKI)</code>, <code>svPhysMap(9DKI)</code>, <code>dataCacheBlockFlush(9DKI)</code>,</p>				

Intel ix86:

svIntrAttach(9DKI), svPhysMap(9DKI), ioLoad(9DKI),

Name	Description
DISABLE_PREEMPT(9DKI)	thread preemption disabling; thread preemption enabling
ENABLE_PREEMPT(9DKI)	See DISABLE_PREEMPT(9DKI)
dataCacheBlockFlush(9DKI)	cache management
dataCacheBlockFlush_powerpc(9DKI)	PowerPC cache management
dataCacheBlockInvalidate(9DKI)	See dataCacheBlockFlush(9DKI)
dataCacheBlockInvalidate_powerpc(9DKI)	See dataCacheBlockFlush_powerpc(9DKI)
dataCacheInvalidate(9DKI)	See dataCacheBlockFlush(9DKI)
dataCacheInvalidate_powerpc(9DKI)	See dataCacheBlockFlush_powerpc(9DKI)
dcacheBlockFlush(9DKI)	See icacheInval(9DKI)
dcacheBlockFlush_usparc(9DKI)	See icacheInval_usparc(9DKI)
dcacheFlush(9DKI)	See icacheInval(9DKI)
dcacheFlush_usparc(9DKI)	See icacheInval_usparc(9DKI)
dcacheLineFlush(9DKI)	See icacheInval(9DKI)
dcacheLineFlush_usparc(9DKI)	See icacheInval_usparc(9DKI)
dtreeNodeAdd(9DKI)	See dtreeNodeRoot(9DKI)
dtreeNodeAlloc(9DKI)	See dtreeNodeRoot(9DKI)
dtreeNodeAttach(9DKI)	See dtreeNodeRoot(9DKI)
dtreeNodeChild(9DKI)	See dtreeNodeRoot(9DKI)
dtreeNodeDetach(9DKI)	See dtreeNodeRoot(9DKI)
dtreeNodeFind(9DKI)	See dtreeNodeRoot(9DKI)

<code>dtreeNodeFree(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodeParent(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePeer(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodeRoot(9DKI)</code>	device tree operations
<code>dtreeNodePathGet(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePathLeng(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropAdd(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropAlloc(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropAttach(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropDetach(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropFind(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropFindNext(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropFree(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropLength(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropName(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>dtreeNodePropValue(9DKI)</code>	See <code>dtreeNodeRoot(9DKI)</code>
<code>eieio(9DKI)</code>	See <code>loadSwapEieio_16(9DKI)</code>
<code>eieio_powerpc(9DKI)</code>	See <code>loadSwapEieio_16_powerpc(9DKI)</code>
<code>icacheBlockInval(9DKI)</code>	See <code>icacheInval(9DKI)</code>
<code>icacheBlockInval_usparc(9DKI)</code>	See <code>icacheInval_usparc(9DKI)</code>
<code>icacheInval(9DKI)</code>	cache management
<code>icacheInval_usparc(9DKI)</code>	UltraSPARC cache management
<code>icacheLineInval(9DKI)</code>	See <code>icacheInval(9DKI)</code>
<code>icacheLineInval_usparc(9DKI)</code>	See <code>icacheInval_usparc(9DKI)</code>
<code>instCacheBlockInvalidate(9DKI)</code>	See <code>dataCacheBlockFlush(9DKI)</code>

	See
instCacheBlockInvalidate_powerpc(9DKI)	dataCacheBlockFlush_powerpc(9DKI)
instCacheInvalidate(9DKI)	See dataCacheBlockFlush(9DKI)
instCacheInvalidate_powerpc(9DKI)	See dataCacheBlockFlush_powerpc(9DKI)
ioLoad16(9DKI)	See ioLoad8(9DKI)
ioLoad16_x86(9DKI)	See ioLoad8_x86(9DKI)
ioLoad32(9DKI)	See ioLoad8(9DKI)
ioLoad32_x86(9DKI)	See ioLoad8_x86(9DKI)
ioLoad8(9DKI)	I/O services
ioLoad8_x86(9DKI)	Intel x86 specific I/O services
ioRead16(9DKI)	See ioLoad8(9DKI)
ioRead16_x86(9DKI)	See ioLoad8_x86(9DKI)
ioRead32(9DKI)	See ioLoad8(9DKI)
ioRead32_x86(9DKI)	See ioLoad8_x86(9DKI)
ioRead8(9DKI)	See ioLoad8(9DKI)
ioRead8_x86(9DKI)	See ioLoad8_x86(9DKI)
ioStore16(9DKI)	See ioLoad8(9DKI)
ioStore16_x86(9DKI)	See ioLoad8_x86(9DKI)
ioStore32(9DKI)	See ioLoad8(9DKI)
ioStore32_x86(9DKI)	See ioLoad8_x86(9DKI)
ioStore8(9DKI)	See ioLoad8(9DKI)
ioStore8_x86(9DKI)	See ioLoad8_x86(9DKI)
ioSync_powerpc(9DKI)	See loadSwapEieio_16_powerpc(9DKI)
ioWrite16(9DKI)	See ioLoad8(9DKI)
ioWrite16_x86(9DKI)	See ioLoad8_x86(9DKI)
ioWrite32(9DKI)	See ioLoad8(9DKI)
ioWrite32_x86(9DKI)	See ioLoad8_x86(9DKI)

<code>ioWrite8(9DKI)</code>	See <code>ioLoad8(9DKI)</code>
<code>ioWrite8_x86(9DKI)</code>	See <code>ioLoad8_x86(9DKI)</code>
<code>loadSwapEieio_16(9DKI)</code>	i/o services
<code>loadSwapEieio_16_powerpc(9DKI)</code>	PowerPC specific i/o services
<code>loadSwapEieio_32(9DKI)</code>	See <code>loadSwapEieio_16(9DKI)</code>
<code>loadSwapEieio_32_powerpc(9DKI)</code>	See <code>loadSwapEieio_16_powerpc(9DKI)</code>
<code>loadSwapSync_16_powerpc(9DKI)</code>	See <code>loadSwapEieio_16_powerpc(9DKI)</code>
<code>loadSwapSync_32_powerpc(9DKI)</code>	See <code>loadSwapEieio_16_powerpc(9DKI)</code>
<code>loadSwap_16(9DKI)</code>	specific i/o services
<code>loadSwap_32(9DKI)</code>	See <code>loadSwap_16(9DKI)</code>
<code>loadSwap_64(9DKI)</code>	See <code>loadSwap_16(9DKI)</code>
<code>loadSwap_sync_16_usparc(9DKI)</code>	See <code>load_sync_8_usparc(9DKI)</code>
<code>loadSwap_sync_32_usparc(9DKI)</code>	See <code>load_sync_8_usparc(9DKI)</code>
<code>loadSwap_sync_64_usparc(9DKI)</code>	See <code>load_sync_8_usparc(9DKI)</code>
<code>load_sync_16_usparc(9DKI)</code>	See <code>load_sync_8_usparc(9DKI)</code>
<code>load_sync_32_usparc(9DKI)</code>	See <code>load_sync_8_usparc(9DKI)</code>
<code>load_sync_64_usparc(9DKI)</code>	See <code>load_sync_8_usparc(9DKI)</code>
<code>load_sync_8_usparc(9DKI)</code>	UltraSparc specific i/o services
<code>storeSwapEieio_16(9DKI)</code>	See <code>loadSwapEieio_16(9DKI)</code>
<code>storeSwapEieio_16_powerpc(9DKI)</code>	See <code>loadSwapEieio_16_powerpc(9DKI)</code>
<code>storeSwapEieio_32(9DKI)</code>	See <code>loadSwapEieio_16(9DKI)</code>

storeSwapEieio_32_powerpc(9DKI)	See loadSwapEieio_16_powerpc(9DKI)
storeSwapSync_16_powerpc(9DKI)	See loadSwapEieio_16_powerpc(9DKI)
storeSwapSync_32_powerpc(9DKI)	See loadSwapEieio_16_powerpc(9DKI)
storeSwap_16(9DKI)	See loadSwap_16(9DKI)
storeSwap_32(9DKI)	See loadSwap_16(9DKI)
storeSwap_64(9DKI)	See loadSwap_16(9DKI)
storeSwap_sync_16_usparc(9DKI)	See load_sync_8_usparc(9DKI)
storeSwap_sync_32_usparc(9DKI)	See load_sync_8_usparc(9DKI)
storeSwap_sync_64_usparc(9DKI)	See load_sync_8_usparc(9DKI)
store_sync_16_usparc(9DKI)	See load_sync_8_usparc(9DKI)
store_sync_32_usparc(9DKI)	See load_sync_8_usparc(9DKI)
store_sync_64_usparc(9DKI)	See load_sync_8_usparc(9DKI)
store_sync_8_usparc(9DKI)	See load_sync_8_usparc(9DKI)
svAsyncExcepAttach(9DKI)	asynchronous exceptions management
svAsyncExcepAttach_usparc(9DKI)	UltraSPARC asynchronous exceptions management
svAsyncExcepDetach(9DKI)	See svAsyncExcepAttach(9DKI)
svAsyncExcepDetach_usparc(9DKI)	See svAsyncExcepAttach_usparc(9DKI)
svDeviceAlloc(9DKI)	See svDeviceRegister(9DKI)
svDeviceEntry(9DKI)	See svDeviceRegister(9DKI)
svDeviceEvent(9DKI)	See svDeviceRegister(9DKI)

svDeviceFree(9DKI)	See svDeviceRegister(9DKI)
svDeviceLookup(9DKI)	See svDeviceRegister(9DKI)
svDeviceRegister(9DKI)	device registry operations
svDeviceRelease(9DKI)	See svDeviceRegister(9DKI)
svDeviceUnregister(9DKI)	See svDeviceRegister(9DKI)
svDkiClose(9DKI)	See svDkiOpen(9DKI)
svDkiEvent(9DKI)	See svDkiOpen(9DKI)
svDkiOpen(9DKI)	system event management
svDkiThreadCall(9DKI)	call a routine in the DKI thread context
svDkiThreadTrigger(9DKI)	See svDkiThreadCall(9DKI)
svDriverCap(9DKI)	See svDriverRegister(9DKI)
svDriverEntry(9DKI)	See svDriverRegister(9DKI)
svDriverLookupFirst(9DKI)	See svDriverRegister(9DKI)
svDriverLookupNext(9DKI)	See svDriverRegister(9DKI)
svDriverRegister(9DKI)	driver registry operations
svDriverRelease(9DKI)	See svDriverRegister(9DKI)
svDriverUnregister(9DKI)	See svDriverRegister(9DKI)
svMemAlloc(9DKI)	A general purpose memory allocator
svMemFree(9DKI)	See svMemAlloc(9DKI)
svPhysAlloc(9DKI)	A special purpose physical memory allocator
svPhysFree(9DKI)	See svPhysAlloc(9DKI)
svPhysMap(9DKI)	physical to virtual memory mapping
svPhysMap_powerpc(9DKI)	PowerPC physical to virtual memory mapping
svPhysMap_usparc(9DKI)	UltraSPARC physical to virtual memory mapping

svPhysMap_x86(9DKI)	Intel x86 physical to virtual memory mapping
svPhysUnmap(9DKI)	See svPhysMap(9DKI)
svPhysUnmap_powerpc(9DKI)	See svPhysMap_powerpc(9DKI)
svPhysUnmap_usparc(9DKI)	See svPhysMap_usparc(9DKI)
svPhysUnmap_x86(9DKI)	See svPhysMap_x86(9DKI)
svTimeoutCancel(9DKI)	See svTimeoutSet(9DKI)
svTimeoutGetRes(9DKI)	See svTimeoutSet(9DKI)
svTimeoutSet(9DKI)	timeout operations
swapEieio_16(9DKI)	See loadSwapEieio_16(9DKI)
swapEieio_16_powerpc(9DKI)	See loadSwapEieio_16_powerpc(9DKI)
swapEieio_32(9DKI)	See loadSwapEieio_16(9DKI)
swapEieio_32_powerpc(9DKI)	See loadSwapEieio_16_powerpc(9DKI)
swap_16(9DKI)	See loadSwap_16(9DKI)
swap_32(9DKI)	See loadSwap_16(9DKI)
swap_64(9DKI)	See loadSwap_16(9DKI)
usecBusyWait(9DKI)	precise busy wait service
vmMapToPhys(9DKI)	See svPhysMap(9DKI)
vmMapToPhys_powerpc(9DKI)	See svPhysMap_powerpc(9DKI)
vmMapToPhys_usparc(9DKI)	See svPhysMap_usparc(9DKI)
vmMapToPhys_x86(9DKI)	See svPhysMap_x86(9DKI)

NAME dataCacheBlockFlush, dataCacheInvalidate, dataCacheBlockInvalidate, instCacheInvalidate, instCacheBlockInvalidate – cache management

FEATURES DKI

DESCRIPTION See the architecture specific man pages:

- dataCacheBlockFlush_powerpc(9DKI)
- dataCacheInvalidate_powerpc(9DKI)
- dataCacheBlockInvalidate_powerpc(9DKI)
- instCacheInvalidate_powerpc(9DKI)
- instCacheBlockInvalidate_powerpc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	dataCacheBlockFlush_powerpc, dataCacheInvalidate_powerpc, dataCacheBlockInvalidate_powerpc, instCacheInvalidate_powerpc, instCacheBlockInvalidate_powerpc – PowerPC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void dataCacheFlush(void); void dataCacheBlockFlush(VmAddr addr, VmSize size); void dataCacheInvalidate(void); void dataCacheBlockInvalidate(VmAddr addr, VmSize size); void instCacheInvalidate(void); void instCacheBlockInvalidate(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	Provides PowerPC cache management services.
EXTENDED DESCRIPTION	<p>The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.</p> <p>Typically, PowerPC family processors have separate instruction and data caches which are virtually indexed and physically tagged. However, the PowerPC architecture does not specify the type or existence of a cache. This allows for various different cache types (unified, or no cache at all). In any case, cache management services should behave consistently. They should do nothing if there is no cache, and data or instruction invalidation routines should be equivalent for a unified cache.</p> <hr/> <p>Note - The data cache may operate in either write-through or copy-back mode, on a per line basis, depending on the cached memory attributes. The data/instruction cache size and data/instruction line size are processor implementation specific.</p> <hr/> <p>The data/instruction cache configuration is available from the PPC_PROP_CACHE property attached to the NODE_CPU node. (The cpu node may be found in the device tree as a child node of the root node. The cpu node name is NODE_CPU.)</p> <p>The cache configuration property value is the PpcPropCache structure.</p> <p>The <i>blockNumber</i> field specifies the number of cache blocks in each data or instruction cache.</p> <p>The <i>blockSize</i> field specifies the cache block size in bytes.</p> <p>The <i>blockSizeShift</i> field specifies the number of bits to shift right/left to divide/multiply by the cache block size (cache block size is always a power of 2).</p>

`dataCacheFlush` globally flushes and invalidates the data cache.

`dataCacheBlockFlush` flushes and invalidates a given range of addresses within the CPU data cache. The range being flushed is specified by the virtual start address (within the current MMU context) and the range size.

`dataCacheInvalidate` globally invalidates the data cache. Note that all blocks in the data cache are marked as invalid without writing back any modified lines to memory. This function does nothing if the data cache is disabled.

`dataCacheBlockInvalidate` invalidates a given range of addresses within the CPU data cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that invalidated blocks are not written back to memory. Note also that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

`instCacheInvalidate` globally invalidates the instruction cache; that is, all blocks in the instruction cache are marked as invalid. This function does nothing if the instruction cache is disabled.

`instCacheBlockInvalidate` invalidates a given range of addresses within the CPU instruction cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dataCacheFlush</code>	+	+	+	-
<code>dataCacheBlockFlush</code>	+	+	+	-
<code>dataCacheInvalidate</code>	+	+	+	-
<code>dataCacheBlockInvalidate</code>	+	+	+	-
<code>instCacheInvalidate</code>	+	+	+	-
<code>instCacheBlockInvalidate</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME dataCacheBlockFlush, dataCacheInvalidate, dataCacheBlockInvalidate, instCacheInvalidate, instCacheBlockInvalidate – cache management

FEATURES DKI

DESCRIPTION See the architecture specific man pages:

- dataCacheBlockFlush_powerpc(9DKI)
- dataCacheInvalidate_powerpc(9DKI)
- dataCacheBlockInvalidate_powerpc(9DKI)
- instCacheInvalidate_powerpc(9DKI)
- instCacheBlockInvalidate_powerpc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	dataCacheBlockFlush_powerpc, dataCacheInvalidate_powerpc, dataCacheBlockInvalidate_powerpc, instCacheInvalidate_powerpc, instCacheBlockInvalidate_powerpc – PowerPC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void dataCacheFlush(void); void dataCacheBlockFlush(VmAddr addr, VmSize size); void dataCacheInvalidate(void); void dataCacheBlockInvalidate(VmAddr addr, VmSize size); void instCacheInvalidate(void); void instCacheBlockInvalidate(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	Provides PowerPC cache management services.
EXTENDED DESCRIPTION	<p>The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.</p> <p>Typically, PowerPC family processors have separate instruction and data caches which are virtually indexed and physically tagged. However, the PowerPC architecture does not specify the type or existence of a cache. This allows for various different cache types (unified, or no cache at all). In any case, cache management services should behave consistently. They should do nothing if there is no cache, and data or instruction invalidation routines should be equivalent for a unified cache.</p> <hr/> <p>Note - The data cache may operate in either write-through or copy-back mode, on a per line basis, depending on the cached memory attributes. The data/instruction cache size and data/instruction line size are processor implementation specific.</p> <hr/> <p>The data/instruction cache configuration is available from the PPC_PROP_CACHE property attached to the NODE_CPU node. (The cpu node may be found in the device tree as a child node of the root node. The cpu node name is NODE_CPU.)</p> <p>The cache configuration property value is the PpcPropCache structure.</p> <p>The <i>blockNumber</i> field specifies the number of cache blocks in each data or instruction cache.</p> <p>The <i>blockSize</i> field specifies the cache block size in bytes.</p> <p>The <i>blockSizeShift</i> field specifies the number of bits to shift right/left to divide/multiply by the cache block size (cache block size is always a power of 2).</p>

`dataCacheFlush` globally flushes and invalidates the data cache.

`dataCacheBlockFlush` flushes and invalidates a given range of addresses within the CPU data cache. The range being flushed is specified by the virtual start address (within the current MMU context) and the range size.

`dataCacheInvalidate` globally invalidates the data cache. Note that all blocks in the data cache are marked as invalid without writing back any modified lines to memory. This function does nothing if the data cache is disabled.

`dataCacheBlockInvalidate` invalidates a given range of addresses within the CPU data cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that invalidated blocks are not written back to memory. Note also that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

`instCacheInvalidate` globally invalidates the instruction cache; that is, all blocks in the instruction cache are marked as invalid. This function does nothing if the instruction cache is disabled.

`instCacheBlockInvalidate` invalidates a given range of addresses within the CPU instruction cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dataCacheFlush</code>	+	+	+	-
<code>dataCacheBlockFlush</code>	+	+	+	-
<code>dataCacheInvalidate</code>	+	+	+	-
<code>dataCacheBlockInvalidate</code>	+	+	+	-
<code>instCacheInvalidate</code>	+	+	+	-
<code>instCacheBlockInvalidate</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME dataCacheBlockFlush, dataCacheInvalidate, dataCacheBlockInvalidate, instCacheInvalidate, instCacheBlockInvalidate – cache management

FEATURES DKI

DESCRIPTION See the architecture specific man pages:

- dataCacheBlockFlush_powerpc(9DKI)
- dataCacheInvalidate_powerpc(9DKI)
- dataCacheBlockInvalidate_powerpc(9DKI)
- instCacheInvalidate_powerpc(9DKI)
- instCacheBlockInvalidate_powerpc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	dataCacheBlockFlush_powerpc, dataCacheInvalidate_powerpc, dataCacheBlockInvalidate_powerpc, instCacheInvalidate_powerpc, instCacheBlockInvalidate_powerpc – PowerPC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void dataCacheFlush(void); void dataCacheBlockFlush(VmAddr addr, VmSize size); void dataCacheInvalidate(void); void dataCacheBlockInvalidate(VmAddr addr, VmSize size); void instCacheInvalidate(void); void instCacheBlockInvalidate(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	Provides PowerPC cache management services.
EXTENDED DESCRIPTION	<p>The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.</p> <p>Typically, PowerPC family processors have separate instruction and data caches which are virtually indexed and physically tagged. However, the PowerPC architecture does not specify the type or existence of a cache. This allows for various different cache types (unified, or no cache at all). In any case, cache management services should behave consistently. They should do nothing if there is no cache, and data or instruction invalidation routines should be equivalent for a unified cache.</p> <hr/> <p>Note - The data cache may operate in either write-through or copy-back mode, on a per line basis, depending on the cached memory attributes. The data/instruction cache size and data/instruction line size are processor implementation specific.</p> <hr/> <p>The data/instruction cache configuration is available from the PPC_PROP_CACHE property attached to the NODE_CPU node. (The cpu node may be found in the device tree as a child node of the root node. The cpu node name is NODE_CPU.)</p> <p>The cache configuration property value is the PpcPropCache structure.</p> <p>The <i>blockNumber</i> field specifies the number of cache blocks in each data or instruction cache.</p> <p>The <i>blockSize</i> field specifies the cache block size in bytes.</p> <p>The <i>blockSizeShift</i> field specifies the number of bits to shift right/left to divide/multiply by the cache block size (cache block size is always a power of 2).</p>

`dataCacheFlush` globally flushes and invalidates the data cache.

`dataCacheBlockFlush` flushes and invalidates a given range of addresses within the CPU data cache. The range being flushed is specified by the virtual start address (within the current MMU context) and the range size.

`dataCacheInvalidate` globally invalidates the data cache. Note that all blocks in the data cache are marked as invalid without writing back any modified lines to memory. This function does nothing if the data cache is disabled.

`dataCacheBlockInvalidate` invalidates a given range of addresses within the CPU data cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that invalidated blocks are not written back to memory. Note also that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

`instCacheInvalidate` globally invalidates the instruction cache; that is, all blocks in the instruction cache are marked as invalid. This function does nothing if the instruction cache is disabled.

`instCacheBlockInvalidate` invalidates a given range of addresses within the CPU instruction cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dataCacheFlush</code>	+	+	+	-
<code>dataCacheBlockFlush</code>	+	+	+	-
<code>dataCacheInvalidate</code>	+	+	+	-
<code>dataCacheBlockInvalidate</code>	+	+	+	-
<code>instCacheInvalidate</code>	+	+	+	-
<code>instCacheBlockInvalidate</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME | icacheInval, icacheLineInval, icacheBlockInval, dcacheFlush, dcacheLineFlush, dcacheBlockFlush – cache management

FEATURES | DKI

DESCRIPTION | The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.

Refer to architecture specific man pages:

- icacheInval_usparc(9DKI)
- icacheLineInval_usparc(9DKI)
- icacheBlockInval_usparc(9DKI)
- dcacheFlush_usparc(9DKI)
- dcacheLineFlush_usparc(9DKI)
- dcacheBlockFlush_usparc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	icacheInval_usparc, icacheLineInval_usparc, icacheBlockInval_usparc, dcacheFlush_usparc, dcacheLineFlush_usparc, dcacheBlockFlush_usparc – UltraSPARC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void icacheInval(void); void icacheLineInval(VmAddr addr); void icacheBlockInval(VmAddr addr, VmSize size); void dcacheFlush(void); void dcacheLineFlush(VmAddr addr); void dcacheBlockFlush(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.
EXTENDED DESCRIPTION	<p>UltraSPARC family processors have split instruction and data caches which are directly mapped and virtually indexed. The instruction cache is two-way associative, directly mapped. It is composed of two banks with an LRU replacement mechanism. The data cache is one-way associative, directly mapped and operates in the write-through mode.</p> <p>The data/instruction caches configuration is available as the "i-cache" and "d-cache" properties attached to the "cpu". The cpu node may be found in the device tree as a child node of the root node. The cpu node name is "cpu".</p> <p>The cache configuration property value is the SparcPropCache structure shown below.</p> <pre>typedef struct { uint32_f csize; /* cache size */ uint32_f bsize; /* cache bank size */ uint32_f lsize; /* cache line size */ uint32_f nbanks; /* number of banks (csize = bsize * nbanks) */ uint32_f type; /* cache type */ } SparcPropCache;</pre> <p>The <i>csize</i> field specifies the cache size in bytes.</p> <p>The <i>bsize</i> field specifies the cache bank size in bytes.</p> <p>The <i>lsize</i> field specifies the cache line size in bytes.</p> <p>The <i>nbanks</i> field specifies the number of banks.</p>

The *type* field specifies the cache properties. It is composed of the following bit-fields and flags:

CACHE_TYPE_IDX	Bit-field specifying the cache indexing mode: CACHE_TYPE_IDX_PHYS — physically indexed CACHE_TYPE_IDX_VIRT — virtually indexed
CACHE_TYPE_TAG	Bit-field specifying the cache tagging mode: CACHE_TYPE_TAG_PHYS — physically tagged CACHE_TYPE_TAG_VIRT — virtually tagged
CACHE_TYPE_MODE	Bit-field specifying the data cache mode: CACHE_TYPE_MODE_WT — write-through CACHE_TYPE_MODE_CB — copy-back
CACHE_TYPE_CFS	Flag specifying that the entire cache flush/invalidation is supported.
CACHE_TYPE_LFS	Flag specifying that line cache flush/invalidation is supported.
CACHE_TYPE_MATCH	Flag specifying that the tag match criteria is used for the line flush/invalidation

Note that, typically, a host bus driver does not need to examine the caches' properties because they are already taken into account by the cache management service routines described below.

`icacheInval` invalidates the entire CPU Instruction Cache.

`icacheLineInval` invalidates a given cache line within the CPU Instruction Cache. The cache line being invalidated is specified by a virtual address.

`icacheBlockInval` invalidates a given range within the CPU Instruction Cache. The range being invalidated is specified by the virtual start address and range size.

`dcacheFlush` flushes and invalidates the entire CPU Data Cache.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
icacheInval	+	+	+	-
icacheLineInval	+	+	+	-
icacheBlockInval	+	+	+	-
dcacheFlush	+	+	+	-
dcacheLineFlush	+	+	+	-
dcacheBlockFlush	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME | icacheInval, icacheLineInval, icacheBlockInval, dcacheFlush, dcacheLineFlush, dcacheBlockFlush – cache management

FEATURES | DKI

DESCRIPTION | The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.

Refer to architecture specific man pages:

- icacheInval_usparc(9DKI)
- icacheLineInval_usparc(9DKI)
- icacheBlockInval_usparc(9DKI)
- dcacheFlush_usparc(9DKI)
- dcacheLineFlush_usparc(9DKI)
- dcacheBlockFlush_usparc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	icacheInval_usparc, icacheLineInval_usparc, icacheBlockInval_usparc, dcacheFlush_usparc, dcacheLineFlush_usparc, dcacheBlockFlush_usparc – UltraSPARC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void icacheInval(void); void icacheLineInval(VmAddr addr); void icacheBlockInval(VmAddr addr, VmSize size); void dcacheFlush(void); void dcacheLineFlush(VmAddr addr); void dcacheBlockFlush(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.
EXTENDED DESCRIPTION	<p>UltraSPARC family processors have split instruction and data caches which are directly mapped and virtually indexed. The instruction cache is two-way associative, directly mapped. It is composed of two banks with an LRU replacement mechanism. The data cache is one-way associative, directly mapped and operates in the write-through mode.</p> <p>The data/instruction caches configuration is available as the "i-cache" and "d-cache" properties attached to the "cpu". The cpu node may be found in the device tree as a child node of the root node. The cpu node name is "cpu".</p> <p>The cache configuration property value is the SparcPropCache structure shown below.</p> <pre>typedef struct { uint32_f csize; /* cache size */ uint32_f bsize; /* cache bank size */ uint32_f lsize; /* cache line size */ uint32_f nbanks; /* number of banks (csize = bsize * nbanks) */ uint32_f type; /* cache type */ } SparcPropCache;</pre> <p>The <i>csize</i> field specifies the cache size in bytes.</p> <p>The <i>bsize</i> field specifies the cache bank size in bytes.</p> <p>The <i>lsize</i> field specifies the cache line size in bytes.</p> <p>The <i>nbanks</i> field specifies the number of banks.</p>

The *type* field specifies the cache properties. It is composed of the following bit-fields and flags:

CACHE_TYPE_IDX	Bit-field specifying the cache indexing mode: CACHE_TYPE_IDX_PHYS — physically indexed CACHE_TYPE_IDX_VIRT — virtually indexed
CACHE_TYPE_TAG	Bit-field specifying the cache tagging mode: CACHE_TYPE_TAG_PHYS — physically tagged CACHE_TYPE_TAG_VIRT — virtually tagged
CACHE_TYPE_MODE	Bit-field specifying the data cache mode: CACHE_TYPE_MODE_WT — write-through CACHE_TYPE_MODE_CB — copy-back
CACHE_TYPE_CFS	Flag specifying that the entire cache flush/invalidation is supported.
CACHE_TYPE_LFS	Flag specifying that line cache flush/invalidation is supported.
CACHE_TYPE_MATCH	Flag specifying that the tag match criteria is used for the line flush/invalidation

Note that, typically, a host bus driver does not need to examine the caches' properties because they are already taken into account by the cache management service routines described below.

`icacheInval` invalidates the entire CPU Instruction Cache.

`icacheLineInval` invalidates a given cache line within the CPU Instruction Cache. The cache line being invalidated is specified by a virtual address.

`icacheBlockInval` invalidates a given range within the CPU Instruction Cache. The range being invalidated is specified by the virtual start address and range size.

`dcacheFlush` flushes and invalidates the entire CPU Data Cache.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

**ALLOWED
CALLING
CONTEXTS**

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
icacheInval	+	+	+	-
icacheLineInval	+	+	+	-
icacheBlockInval	+	+	+	-
dcacheFlush	+	+	+	-
dcacheLineFlush	+	+	+	-
dcacheBlockFlush	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME | icacheInval, icacheLineInval, icacheBlockInval, dcacheFlush, dcacheLineFlush, dcacheBlockFlush – cache management

FEATURES | DKI

DESCRIPTION | The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.

Refer to architecture specific man pages:

- icacheInval_usparc(9DKI)
- icacheLineInval_usparc(9DKI)
- icacheBlockInval_usparc(9DKI)
- dcacheFlush_usparc(9DKI)
- dcacheLineFlush_usparc(9DKI)
- dcacheBlockFlush_usparc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	icacheInval_usparc, icacheLineInval_usparc, icacheBlockInval_usparc, dcacheFlush_usparc, dcacheLineFlush_usparc, dcacheBlockFlush_usparc – UltraSPARC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void icacheInval(void); void icacheLineInval(VmAddr addr); void icacheBlockInval(VmAddr addr, VmSize size); void dcacheFlush(void); void dcacheLineFlush(VmAddr addr); void dcacheBlockFlush(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.
EXTENDED DESCRIPTION	<p>UltraSPARC family processors have split instruction and data caches which are directly mapped and virtually indexed. The instruction cache is two-way associative, directly mapped. It is composed of two banks with an LRU replacement mechanism. The data cache is one-way associative, directly mapped and operates in the write-through mode.</p> <p>The data/instruction caches configuration is available as the "i-cache" and "d-cache" properties attached to the "cpu". The cpu node may be found in the device tree as a child node of the root node. The cpu node name is "cpu".</p> <p>The cache configuration property value is the SparcPropCache structure shown below.</p> <pre>typedef struct { uint32_f csize; /* cache size */ uint32_f bsize; /* cache bank size */ uint32_f lsize; /* cache line size */ uint32_f nbanks; /* number of banks (csize = bsize * nbanks) */ uint32_f type; /* cache type */ } SparcPropCache;</pre> <p>The <i>csize</i> field specifies the cache size in bytes.</p> <p>The <i>bsize</i> field specifies the cache bank size in bytes.</p> <p>The <i>lsize</i> field specifies the cache line size in bytes.</p> <p>The <i>nbanks</i> field specifies the number of banks.</p>

The *type* field specifies the cache properties. It is composed of the following bit-fields and flags:

CACHE_TYPE_IDX	Bit-field specifying the cache indexing mode: CACHE_TYPE_IDX_PHYS — physically indexed CACHE_TYPE_IDX_VIRT — virtually indexed
CACHE_TYPE_TAG	Bit-field specifying the cache tagging mode: CACHE_TYPE_TAG_PHYS — physically tagged CACHE_TYPE_TAG_VIRT — virtually tagged
CACHE_TYPE_MODE	Bit-field specifying the data cache mode: CACHE_TYPE_MODE_WT — write-through CACHE_TYPE_MODE_CB — copy-back
CACHE_TYPE_CFS	Flag specifying that the entire cache flush/invalidation is supported.
CACHE_TYPE_LFS	Flag specifying that line cache flush/invalidation is supported.
CACHE_TYPE_MATCH	Flag specifying that the tag match criteria is used for the line flush/invalidation

Note that, typically, a host bus driver does not need to examine the caches' properties because they are already taken into account by the cache management service routines described below.

`icacheInval` invalidates the entire CPU Instruction Cache.

`icacheLineInval` invalidates a given cache line within the CPU Instruction Cache. The cache line being invalidated is specified by a virtual address.

`icacheBlockInval` invalidates a given range within the CPU Instruction Cache. The range being invalidated is specified by the virtual start address and range size.

`dcacheFlush` flushes and invalidates the entire CPU Data Cache.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

**ALLOWED
CALLING
CONTEXTS**

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
icacheInval	+	+	+	-
icacheLineInval	+	+	+	-
icacheBlockInval	+	+	+	-
dcacheFlush	+	+	+	-
dcacheLineFlush	+	+	+	-
dcacheBlockFlush	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME	DISABLE_PREEMPT, ENABLE_PREEMPT – thread preemption disabling; thread preemption enabling															
SYNOPSIS	#include <dki/dki.h> DISABLE_PREEMPT(); ENABLE_PREEMPT();															
FEATURES	DKI															
DESCRIPTION	DKI provides a means for a driver to disable/enable the preemption of the current thread. These services can be used by a driver to prevent the current thread being preempted while interrupts are masked at bus/device level.															
	<hr/> Note - These services are implemented as macros. <hr/>															
EXTENDED DESCRIPTION																
Thread Preemption Disabling	The <code>DISABLE_PREEMPT()</code> macro disables preemption of the thread which is currently executing. This macro increments a per-processor preemption mask count. When this count is not zero, the scheduler is locked. This occurs as when there is a preemption request, the scheduler simply raises a pending preemption flag deferring the real thread preemption until the preemption mask count drops to zero.															
Thread Preemption Enabling	The <code>ENABLE_PREEMPT()</code> macro enables preemption of the thread which is currently executing and was previously disabled by <code>DISABLE_PREEMPT()</code> as outlined above. This macro decrements the preemption mask count and, if it drops to zero, checks whether the pending preemption flag is raised and thus if the current thread should be preempted.															
	<hr/> Note - As <code>DISABLE_PREEMPT()</code> / <code>ENABLE_PREEMPT()</code> rely on the preemption mask count, a driver may issue nested calls to these services. <hr/>															
Allowed Calling Contexts	The following table specifies the allowed calling contexts for each service:															
	<table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="text-align: left;">Services</th> <th>Base level</th> <th>DKI thread</th> <th>Interrupt</th> <th>Blocking</th> </tr> </thead> <tbody> <tr> <td style="text-align: left;">DISABLE_PREEMPT</td> <td>+</td> <td>+</td> <td>+</td> <td>-</td> </tr> <tr> <td style="text-align: left;">ENABLE_PREEMPT</td> <td>+</td> <td>+</td> <td>+</td> <td>-</td> </tr> </tbody> </table>	Services	Base level	DKI thread	Interrupt	Blocking	DISABLE_PREEMPT	+	+	+	-	ENABLE_PREEMPT	+	+	+	-
Services	Base level	DKI thread	Interrupt	Blocking												
DISABLE_PREEMPT	+	+	+	-												
ENABLE_PREEMPT	+	+	+	-												
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:															

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreeNodePropFind, dtreeNodePropFindNext, dtreeNodePropLength, dtreeNodePropValue, dtreeNodePropName, dtreeNodePropAlloc, dtreeNodePropFree, dtreeNodePropAttach, dtreeNodePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreeNodePropAdd, dtreeNodePathLeng, dtreeNodePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreeNodePropFind(DevNode node, char * name); DevProperty dtreeNodePropFindNext(DevProperty prop, char * name); unsigned int dtreeNodePropLength(DevProperty prop); void * dtreeNodePropValue(DevProperty prop); char * dtreeNodePropName(DevProperty prop); DevProperty dtreeNodePropAlloc(char * name, int length); void dtreeNodePropFree(DevProperty prop); void dtreeNodePropAttach(DevNode node, DevProperty prop); void dtreeNodePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreeNodePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreeNodePathLeng(DevNode node, int length); DevProperty dtreeNodePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreeNodePathLeng` returns the pathname length of a device node.

`dtreeNodePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreeNodePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreeNodePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreeNodePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreeNodePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreeNodePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreeNodePropFind` / `dtreeNodePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreeNodePropFind` invocation is sufficient.

`dtreeNodePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreeNodePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreeNodePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreeNodePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreeNodePropAttach` function.

`dtreeNodePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreeNodePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreeNodePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreeNodePropValue	-	+	-	-
dtreeNodePropName	-	+	-	-
dtreeNodePropAlloc	+	+	-	+
dtreeNodePropFree	+	+	-	+
dtreeNodePropAttach	-	+	-	-
dtreeNodePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreeNodePropAdd	-	+	-	+
dtreeNodePathLeng	+	+	-	-
dtreeNodePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreeNodePropFind, dtreeNodePropFindNext, dtreeNodePropLength, dtreeNodePropValue, dtreeNodePropName, dtreeNodePropAlloc, dtreeNodePropFree, dtreeNodePropAttach, dtreeNodePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreeNodePropAdd, dtreeNodePathLeng, dtreeNodePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreeNodePropFind(DevNode node, char * name); DevProperty dtreeNodePropFindNext(DevProperty prop, char * name); unsigned int dtreeNodePropLength(DevProperty prop); void * dtreeNodePropValue(DevProperty prop); char * dtreeNodePropName(DevProperty prop); DevProperty dtreeNodePropAlloc(char * name, int length); void dtreeNodePropFree(DevProperty prop); void dtreeNodePropAttach(DevNode node, DevProperty prop); void dtreeNodePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreeNodePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreeNodePathLeng(DevNode node, int length); DevProperty dtreeNodePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreeNodePathLeng` returns the pathname length of a device node.

`dtreeNodePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreeNodePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreeNodePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreeNodePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreeNodePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreeNodePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreeNodePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreeNodePropFind</code>	-	+	-	-
<code>dtreeNodePropFindNext</code>	-	+	-	-
<code>dtreeNodePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreeNodePropFind, dtreeNodePropFindNext, dtreeNodePropLength, dtreeNodePropValue, dtreeNodePropName, dtreeNodePropAlloc, dtreeNodePropFree, dtreeNodePropAttach, dtreeNodePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreeNodePropAdd, dtreeNodePathLeng, dtreeNodePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreeNodePropFind(DevNode node, char * name); DevProperty dtreeNodePropFindNext(DevProperty prop, char * name); unsigned int dtreeNodePropLength(DevProperty prop); void * dtreeNodePropValue(DevProperty prop); char * dtreeNodePropName(DevProperty prop); DevProperty dtreeNodePropAlloc(char * name, int length); void dtreeNodePropFree(DevProperty prop); void dtreeNodePropAttach(DevNode node, DevProperty prop); void dtreeNodePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreeNodePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreeNodePathLeng(DevNode node, int length); DevProperty dtreeNodePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreeNodePathLeng` returns the pathname length of a device node.

`dtreeNodePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreeNodePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreeNodePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreeNodePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreeNodePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreeNodePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre> #include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value); </pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreeNodePathLeng` returns the pathname length of a device node.

`dtreeNodePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreeNodePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreeNodePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreeNodePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreeNodePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreeNodePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreeNodePropFind` / `dtreeNodePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreeNodePropFind` invocation is sufficient.

`dtreeNodePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreeNodePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreeNodePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreeNodePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreeNodePropAttach` function.

`dtreeNodePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreeNodePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreeNodePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreeNodePathLeng` returns the pathname length of a device node.

`dtreeNodePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreeNodePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreeNodePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreeNodePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreeNodePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreeNodePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreeNodePropFind` / `dtreeNodePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreeNodePropFind` invocation is sufficient.

`dtreeNodePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreeNodePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreeNodePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreeNodePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreeNodePropAttach` function.

`dtreeNodePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreeNodePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreeNodePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreeNodePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreeNodePropFind</code>	-	+	-	-
<code>dtreeNodePropFindNext</code>	-	+	-	-
<code>dtreeNodePropLength</code>	-	+	-	-

dtreeNodePropValue	-	+	-	-
dtreeNodePropName	-	+	-	-
dtreeNodePropAlloc	+	+	-	+
dtreeNodePropFree	+	+	-	+
dtreeNodePropAttach	-	+	-	-
dtreeNodePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreeNodePropAdd	-	+	-	+
dtreeNodePathLeng	+	+	-	-
dtreeNodePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre> #include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value); </pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreeNodePathLeng` returns the pathname length of a device node.

`dtreeNodePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreeNodePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreeNodePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreeNodePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreeNodePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreeNodePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION

Provides device tree operations.

EXTENDED DESCRIPTION

The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.

A device property is a *name* / *value* pair. The property *name* is a null terminated ASCII string. The property *value* is a sequence of bytes specified by the *length* / *address* pair.

Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).

These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the `booter`. For instance, the `booter` may include a pre-defined sequence of device tree function calls. Another use of the `booter` is to build the device tree from a hardware description provided by firmware.

In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.

Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the `booter` which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.

Device Tree Browsing

This section describes APIs related to device tree browsing. `DevNode` is an abstract type designating a device node object, and is opaque to the driver.

`dtreeNodeRoot` returns the root device node if the device tree is not empty, otherwise NULL is returned.

`dtreeNodeChild` returns the first child node from the list of children. NULL is returned when the list of children is empty.

`dtreeNodePeer` returns the next device node from the sibling list, if any, otherwise NULL is returned. The *node* argument specifies the current device node in the sibling list.

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreeNodePathLeng` returns the pathname length of a device node.

`dtreeNodePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreeNodePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreeNodePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreeNodePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreeNodePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreeNodePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreeNodePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreeNodePropFind</code>	-	+	-	-
<code>dtreeNodePropFindNext</code>	-	+	-	-
<code>dtreeNodePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre> #include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value); </pre>
FEATURES	DKI

DESCRIPTION

Provides device tree operations.

EXTENDED DESCRIPTION

The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.

A device property is a *name* / *value* pair. The property *name* is a null terminated ASCII string. The property *value* is a sequence of bytes specified by the *length* / *address* pair.

Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).

These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the `booter`. For instance, the `booter` may include a pre-defined sequence of device tree function calls. Another use of the `booter` is to build the device tree from a hardware description provided by firmware.

In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.

Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the `booter` which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.

Device Tree Browsing

This section describes APIs related to device tree browsing. `DevNode` is an abstract type designating a device node object, and is opaque to the driver.

`dtreeNodeRoot` returns the root device node if the device tree is not empty, otherwise NULL is returned.

`dtreeNodeChild` returns the first child node from the list of children. NULL is returned when the list of children is empty.

`dtreeNodePeer` returns the next device node from the sibling list, if any, otherwise NULL is returned. The *node* argument specifies the current device node in the sibling list.

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreeNodePathLeng` returns the pathname length of a device node.

`dtreeNodePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreeNodePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreeNodePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreeNodePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreeNodePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreeNodePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreeNodePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre> #include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value); </pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre> #include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value); </pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device noe. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre>#include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value);</pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	dtreeNodeRoot, dtreeNodeChild, dtreeNodePeer, dtreeNodeParent, dtreeNodeAlloc, dtreeNodeFree, dtreeNodeAttach, dtreeNodeDetach, dtreePropFind, dtreePropFindNext, dtreePropLength, dtreePropValue, dtreePropName, dtreePropAlloc, dtreePropFree, dtreePropAttach, dtreePropDetach, dtreeNodeAdd, dtreeNodeFind, dtreePropAdd, dtreePathLeng, dtreePathGet – device tree operations
SYNOPSIS	<pre> #include <dki/dki.h> DevNode dtreeNodeRoot(void); DevNode dtreeNodeChild(DevNode node); DevNode dtreeNodePeer(DevNode node); DevNode dtreeNodeParent(DevNode node); DevNode dtreeNodeAlloc(void); void dtreeNodeFree(DevNode); void dtreeNodeAttach(DevNode pnode, DevNode cnode); void dtreeNodeDetach(DevNode node); DevProperty dtreePropFind(DevNode node, char * name); DevProperty dtreePropFindNext(DevProperty prop, char * name); unsigned int dtreePropLength(DevProperty prop); void * dtreePropValue(DevProperty prop); char * dtreePropName(DevProperty prop); DevProperty dtreePropAlloc(char * name, int length); void dtreePropFree(DevProperty prop); void dtreePropAttach(DevNode node, DevProperty prop); void dtreePropDetach(DevProperty prop); DevNode dtreeNodeAdd(DevNode parent, char * name); DevNode dtreeNodeFind(DevNode parent, char * name); DevProperty dtreePropAdd(DevNode node, char * name, void * value, unsigned int length); DevProperty dtreePathLeng(DevNode node, int length); DevProperty dtreePathGet(DevNode node, char * buf, void * value); </pre>
FEATURES	DKI

DESCRIPTION	Provides device tree operations.
EXTENDED DESCRIPTION	<p>The device tree is a data structure which provides a description of the hardware topology and device properties. The hardware topology is specified in terms of parent/child relationships. Device properties associated with each device node in the tree are device specific.</p> <p>A device property is a <i>name</i> / <i>value</i> pair. The property <i>name</i> is a null terminated ASCII string. The property <i>value</i> is a sequence of bytes specified by the <i>length</i> / <i>address</i> pair.</p> <p>Note that the property value format is property specific and has to be standardized between the given producer and its consumers. For instance, among all device node properties there are some properties related to the bus resources required/allocated for the device (for instance, interrupt lines, I/O registers, DMA channels).</p> <p>These types of properties must be standardized in order to be understood by the bus driver as well as by any device drivers connected to the given bus. The device tree data structure may be built either statically or dynamically. In the static case, the device tree is populated by the <code>booter</code>. For instance, the <code>booter</code> may include a pre-defined sequence of device tree function calls. Another use of the <code>booter</code> is to build the device tree from a hardware description provided by firmware.</p> <p>In the dynamic case, the device tree is populated at system initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parents to children.</p> <p>Note that it is possible to combine both methods. In other words, an initial (non complete) device tree may be provided by the <code>booter</code> which will later be completed dynamically using an enumeration/probing mechanism. However it is implemented, the device tree structure can be modified (extended/truncated) dynamically at run time using hot-plug insertion/removal (for example, PCMCIA cards). The device tree API is described in detail below.</p>
Device Tree Browsing	<p>This section describes APIs related to device tree browsing. <code>DevNode</code> is an abstract type designating a device node object, and is opaque to the driver.</p> <p><code>dtreeNodeRoot</code> returns the root device node if the device tree is not empty, otherwise NULL is returned.</p> <p><code>dtreeNodeChild</code> returns the first child node from the list of children. NULL is returned when the list of children is empty.</p> <p><code>dtreeNodePeer</code> returns the next device node from the sibling list, if any, otherwise NULL is returned. The <i>node</i> argument specifies the current device node in the sibling list.</p>

`dtreeNodeParent` returns the parent device node, if any, otherwise NULL is returned. The *node* argument specifies the child device node.

`dtreePathLeng` returns the pathname length of a device node.

`dtreePathGet` returns in *buf* the absolute pathname of a device node. The trailing path, of the pathname, is the name of the node and is read in a node property. If this property does not exist, the trailing path of the returned pathname is set to '???'.

Device Tree Modification

This section describes APIs related to device tree topology modification.

`dtreeNodeAlloc` allocates a new device node object. A non zero `DevNode` cookie is returned in case of success, otherwise NULL is returned. The allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

`dtreeNodeFree` releases the memory allocated by the node object and all property objects attached to the node. The *node* argument specifies the node object being released.

`dtreeNodeAttach` adds the device node specified by the *cnode* argument in the child list of the parent node specified by the *pnode* argument.

`dtreeNodeDetach` detaches the node object specified by the *node* argument from its parent (if any). When *node* specifies the root node, the device tree is emptied.

Device Node Properties

This section describes APIs related to the device node properties.

`DevProperty` is an abstract type designating a device property object, and is opaque to the driver.

`dtreePropFind` searches the first property within the properties list of the device node specified by the *node* argument. If the *name* argument is not NULL, `dtreePropFind` returns the first property whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFind` returns the first property from the list regardless of its name.

In case of success, a non zero `DevProperty` cookie is returned, otherwise NULL is returned. Once the first property is found, the `dtreePropFindNext` routine may be used in order to find a subsequent property in the list.

`dtreePropFindNext` searches the next property within the properties list. The current position within the properties list is specified by the *prop* argument. If the *name* argument is not NULL, `dtreePropFindNext` returns the next property (with respect to the current position) whose name matches the *name* string. If the *name* argument is NULL, `dtreePropFindNext` returns the next property from the list (with respect to the current position) regardless of its name.

If successful, a non zero `DevProperty` cookie is returned, otherwise `NULL` is returned. The `dtreePropFind` / `dtreePropFindNext` pair are typically used to iterate through either all device node properties, or a subset of device node properties whose names match a given name. In a case where a user knows that there is only one property with a given name attached to the device node, a single `dtreePropFind` invocation is sufficient.

`dtreePropLength` returns the property value length (in bytes). The property object is specified by the *prop* argument.

`dtreePropValue` returns a pointer to the first byte of the property value. The property object is specified by the *prop* argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

`dtreePropName` returns a pointer to an ASCII string which designates the property name. The property object is specified by the *prop* argument.

`dtreePropAlloc` allocates a new device property object. A non zero `DevProperty` cookie is returned if successful, otherwise `NULL` is returned. The *name* argument specifies the property name. The *length* argument specifies the length of the property value. The property value is undefined. The property object allocated is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (that is, added to the node properties list) using the `dtreePropAttach` function.

`dtreePropFree` releases the memory allocated by the property object. The *prop* parameter specifies the property object being released. Note that the property object must not be attached to any device node.

`dtreePropAttach` adds the property object specified by the *prop* argument to the node properties list. The *node* argument specifies the device node to which the property will be attached.

`dtreePropDetach` detaches the property object from the device node to which it is attached (if any). The *prop* argument specifies the property object.

Note that the device tree structure must only be accessed in the DKI thread context. The DKI thread provides access synchronization to the device tree. In many cases, driver routines which typically examine or modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. This rule also applies to driver clients; a driver client must always switch to the DKI thread context when accessing the device tree structure.

**Device Tree
High-Level Services**

This section describes high-level APIs built upon other basic services. This API implements certain useful services for building and searching the device tree, in order to avoid implementing this code in all device tree users.

`dtreeNodeAdd` allocates a new device node object, and adds it to the child list of the *parent* device node specified..

The *name* argument specifies the name of the new device node allocated.. This means that a *node* property is allocated and attached to the new node with the value specified by *name* . If successful, the newly allocated node object is returned, otherwise, a NULL pointer is returned.

`dtreeNodeFind` looks for a named node in the child list of a specified device node. The *parent* argument specifies the device node within which the child list should be searched. The *name* argument specifies the value which must match the *node* property value of the node being searched.. If successful, the matching node object is returned, otherwise a NULL pointer is returned.

`dtreePropAdd` allocates a new property, sets its value and attaches it to a given device node. The *node* argument specifies the node to which the new property must be attached. The *name* argument specifies the name of the new property to allocate. The *length* argument specifies the memory size in bytes to be allocated for the new property value. The *value* argument specifies the value to be set for the newly allocated property. If successful, the newly allocated property object is returned, otherwise, a NULL pointer is returned.

**Allowed Calling
Contexts**

The following table specifies the contexts in which a caller may invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dtreeNodeRoot</code>	-	+	-	-
<code>dtreeNodeChild</code>	-	+	-	-
<code>dtreeNodePeer</code>	-	+	-	-
<code>dtreeNodeParent</code>	-	+	-	-
<code>dtreeNodeAlloc</code>	+	+	-	+
<code>dtreeNodeFree</code>	+	+	-	+
<code>dtreeNodeAttach</code>	-	+	-	-
<code>dtreeNodeDetach</code>	-	+	-	-
<code>dtreePropFind</code>	-	+	-	-
<code>dtreePropFindNext</code>	-	+	-	-
<code>dtreePropLength</code>	-	+	-	-

dtreePropValue	-	+	-	-
dtreePropName	-	+	-	-
dtreePropAlloc	+	+	-	+
dtreePropFree	+	+	-	+
dtreePropAttach	-	+	-	-
dtreePropDetach	-	+	-	-
dtreeNodeAdd	-	+	-	+
dtreeNodeFind	-	+	-	-
dtreePropAdd	-	+	-	+
dtreePathLeng	+	+	-	-
dtreePathGet	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	loadSwapEieio_16, storeSwapEieio_16, swapEieio_16, loadSwapEieio_32, storeSwapEieio_32, swapEieio_32, eieio - i/o services
FEATURES	DKI
DESCRIPTION	<p>Provides access to DKI i/o routines, optimized to facilitate byte swapping and synchronization of input/output.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ loadSwapEieio_16_powerpc(9DKI) ■ storeSwapEieio_16_powerpc(9DKI) ■ swapEieio_16_powerpc(9DKI) ■ loadSwapEieio_32_powerpc(9DKI) ■ storeSwapEieio_32_powerpc(9DKI) ■ swapEieio_32_powerpc(9DKI) ■ eieio(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	loadSwapEieio_16_powerpc, loadSwapSync_16_powerpc, storeSwapEieio_16_powerpc, storeSwapSync_16_powerpc, swapEieio_16_powerpc, loadSwapEieio_32_powerpc, loadSwapSync_32_powerpc, storeSwapEieio_32_powerpc, storeSwapSync_32_powerpc, swapEieio_32_powerpc, eieio_powerpc, ioSync_powerpc – PowerPC specific i/o services
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwapEieio_16(uint16_f* addr); uint16_f loadSwapSync_16(uint16_f* addr); void storeSwapEieio_16(uint16_f* addr, uint16_f value); void storeSwapSync_16(uint16_f* addr, uint16_f value); void swapEieio_16(uint16_f* addr); uint32_f loadSwapEieio_32(uint32_f* addr); uint32_f loadSwapSync_32(uint32_f* addr); void storeSwapEieio_32(uint32_f* addr, uint32_f value); void storeSwapSync_32(uint32_f* addr, uint32_f value); void swapEieio_32 (uint32_f* addr); void ioSync(void); void eieio(void);</pre>
FEATURES	DKI
DESCRIPTION	Provides access to PowerPC DKI I/O routines, optimized to facilitate byte swapping and synchronization of input/output.
EXTENDED DESCRIPTION	<p>The PowerPC DKI provides specific I/O routines to handle byte swapping and synchronization of I/O. These services are based on PowerPC specific instructions to store/load with byte swapping, and to enforce in-order execution of I/O.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the PowerPC bus and the host bus, or to enforce sequential execution of I/O. PowerPC specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <pre>_16 for 16-bit data _32 for 32-bit data</pre>

`loadSwapEieio_xx` loads data from a given address and returns the corresponding byte swapped value. A PowerPC `eieio` instruction is issued after the data is read, to enforce execution of subsequent I/O sequentially with respect to the current load. The *addr* argument specifies the address to read from.

`loadSwapSync_xx` behaves like `loadSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`storeSwapEieio_xx` stores, at a given address, the *value* byte swapped. A PowerPC `eieio` instruction is issued after the data is written, to enforce execution of subsequent I/O sequentially with respect to the current store. The *addr* argument specifies the address to write to.

`storeSwapSync_xx` behaves like `storeSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`swapEieio_xx` swaps in place the bytes of the data stored at a given address. A PowerPC `eieio` instruction is issued between reading the data and re-writing the byte-swapped data, to enforce execution in order. The *addr* argument specifies the address of the data to byte-swap.

`ioSync` enforces context synchronization, completion of all instructions, and completion of all exceptions before returning.

`eieio` enforces execution in order of I/O and is mapped to the corresponding PowerPC assembler instruction. This routine is intended for use in managing shared data structures, in doing memory-mapped I/O, and in preventing load/store combining operations in main memory. Refer to PowerPC programming manuals for more details.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwapEieio_xx</code>	+	+	+	-
<code>loadSwapSync_xx</code>	+	+	+	-
<code>storeSwapEieio_xx</code>	+	+	+	-
<code>storeSwapSync_xx</code>	+	+	+	-
<code>swapEieio_xx</code>	+	+	+	-
<code>ioSync</code>	+	+	+	-
<code>eieio</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap_16(9DKI)` , `svDkiThreadCall(9DKI)`

NAME	DISABLE_PREEMPT, ENABLE_PREEMPT – thread preemption disabling; thread preemption enabling															
SYNOPSIS	#include <dki/dki.h> DISABLE_PREEMPT(); ENABLE_PREEMPT();															
FEATURES	DKI															
DESCRIPTION	DKI provides a means for a driver to disable/enable the preemption of the current thread. These services can be used by a driver to prevent the current thread being preempted while interrupts are masked at bus/device level. <hr/> Note - These services are implemented as macros. <hr/>															
EXTENDED DESCRIPTION																
Thread Preemption Disabling	The <code>DISABLE_PREEMPT()</code> macro disables preemption of the thread which is currently executing. This macro increments a per-processor preemption mask count. When this count is not zero, the scheduler is locked. This occurs as when there is a preemption request, the scheduler simply raises a pending preemption flag deferring the real thread preemption until the preemption mask count drops to zero.															
Thread Preemption Enabling	The <code>ENABLE_PREEMPT()</code> macro enables preemption of the thread which is currently executing and was previously disabled by <code>DISABLE_PREEMPT()</code> as outlined above. This macro decrements the preemption mask count and, if it drops to zero, checks whether the pending preemption flag is raised and thus if the current thread should be preempted. <hr/> Note - As <code>DISABLE_PREEMPT()</code> / <code>ENABLE_PREEMPT()</code> rely on the preemption mask count, a driver may issue nested calls to these services. <hr/>															
Allowed Calling Contexts	The following table specifies the allowed calling contexts for each service: <table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Services</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Base level</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">DKI thread</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Interrupt</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Blocking</th> </tr> </thead> <tbody> <tr> <td style="border-bottom: 1px solid black;">DISABLE_PREEMPT</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">-</td> </tr> <tr> <td style="border-bottom: 1px solid black;">ENABLE_PREEMPT</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">-</td> </tr> </tbody> </table>	Services	Base level	DKI thread	Interrupt	Blocking	DISABLE_PREEMPT	+	+	+	-	ENABLE_PREEMPT	+	+	+	-
Services	Base level	DKI thread	Interrupt	Blocking												
DISABLE_PREEMPT	+	+	+	-												
ENABLE_PREEMPT	+	+	+	-												
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:															

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME | icacheInval, icacheLineInval, icacheBlockInval, dcacheFlush, dcacheLineFlush, dcacheBlockFlush – cache management

FEATURES | DKI

DESCRIPTION | The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.

Refer to architecture specific man pages:

- icacheInval_usparc(9DKI)
- icacheLineInval_usparc(9DKI)
- icacheBlockInval_usparc(9DKI)
- dcacheFlush_usparc(9DKI)
- dcacheLineFlush_usparc(9DKI)
- dcacheBlockFlush_usparc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	icacheInval_usparc, icacheLineInval_usparc, icacheBlockInval_usparc, dcacheFlush_usparc, dcacheLineFlush_usparc, dcacheBlockFlush_usparc – UltraSPARC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void icacheInval(void); void icacheLineInval(VmAddr addr); void icacheBlockInval(VmAddr addr, VmSize size); void dcacheFlush(void); void dcacheLineFlush(VmAddr addr); void dcacheBlockFlush(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.
EXTENDED DESCRIPTION	<p>UltraSPARC family processors have split instruction and data caches which are directly mapped and virtually indexed. The instruction cache is two-way associative, directly mapped. It is composed of two banks with an LRU replacement mechanism. The data cache is one-way associative, directly mapped and operates in the write-through mode.</p> <p>The data/instruction caches configuration is available as the "i-cache" and "d-cache" properties attached to the "cpu". The cpu node may be found in the device tree as a child node of the root node. The cpu node name is "cpu".</p> <p>The cache configuration property value is the SparcPropCache structure shown below.</p> <pre>typedef struct { uint32_f csize; /* cache size */ uint32_f bsize; /* cache bank size */ uint32_f lsize; /* cache line size */ uint32_f nbanks; /* number of banks (csize = bsize * nbanks) */ uint32_f type; /* cache type */ } SparcPropCache;</pre> <p>The <i>csize</i> field specifies the cache size in bytes.</p> <p>The <i>bsize</i> field specifies the cache bank size in bytes.</p> <p>The <i>lsize</i> field specifies the cache line size in bytes.</p> <p>The <i>nbanks</i> field specifies the number of banks.</p>

The *type* field specifies the cache properties. It is composed of the following bit-fields and flags:

CACHE_TYPE_IDX	Bit-field specifying the cache indexing mode: CACHE_TYPE_IDX_PHYS — physically indexed CACHE_TYPE_IDX_VIRT — virtually indexed
CACHE_TYPE_TAG	Bit-field specifying the cache tagging mode: CACHE_TYPE_TAG_PHYS — physically tagged CACHE_TYPE_TAG_VIRT — virtually tagged
CACHE_TYPE_MODE	Bit-field specifying the data cache mode: CACHE_TYPE_MODE_WT — write-through CACHE_TYPE_MODE_CB — copy-back
CACHE_TYPE_CFS	Flag specifying that the entire cache flush/invalidation is supported.
CACHE_TYPE_LFS	Flag specifying that line cache flush/invalidation is supported.
CACHE_TYPE_MATCH	Flag specifying that the tag match criteria is used for the line flush/invalidation

Note that, typically, a host bus driver does not need to examine the caches' properties because they are already taken into account by the cache management service routines described below.

`icacheInval` invalidates the entire CPU Instruction Cache.

`icacheLineInval` invalidates a given cache line within the CPU Instruction Cache. The cache line being invalidated is specified by a virtual address.

`icacheBlockInval` invalidates a given range within the CPU Instruction Cache. The range being invalidated is specified by the virtual start address and range size.

`dcacheFlush` flushes and invalidates the entire CPU Data Cache.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

**ALLOWED
CALLING
CONTEXTS**

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
icacheInval	+	+	+	-
icacheLineInval	+	+	+	-
icacheBlockInval	+	+	+	-
dcacheFlush	+	+	+	-
dcacheLineFlush	+	+	+	-
dcacheBlockFlush	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME | icacheInval, icacheLineInval, icacheBlockInval, dcacheFlush, dcacheLineFlush, dcacheBlockFlush – cache management

FEATURES | DKI

DESCRIPTION | The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.

Refer to architecture specific man pages:

- icacheInval_usparc(9DKI)
- icacheLineInval_usparc(9DKI)
- icacheBlockInval_usparc(9DKI)
- dcacheFlush_usparc(9DKI)
- dcacheLineFlush_usparc(9DKI)
- dcacheBlockFlush_usparc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	icacheInval_usparc, icacheLineInval_usparc, icacheBlockInval_usparc, dcacheFlush_usparc, dcacheLineFlush_usparc, dcacheBlockFlush_usparc – UltraSPARC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void icacheInval(void); void icacheLineInval(VmAddr addr); void icacheBlockInval(VmAddr addr, VmSize size); void dcacheFlush(void); void dcacheLineFlush(VmAddr addr); void dcacheBlockFlush(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.
EXTENDED DESCRIPTION	<p>UltraSPARC family processors have split instruction and data caches which are directly mapped and virtually indexed. The instruction cache is two-way associative, directly mapped. It is composed of two banks with an LRU replacement mechanism. The data cache is one-way associative, directly mapped and operates in the write-through mode.</p> <p>The data/instruction caches configuration is available as the "i-cache" and "d-cache" properties attached to the "cpu". The cpu node may be found in the device tree as a child node of the root node. The cpu node name is "cpu".</p> <p>The cache configuration property value is the SparcPropCache structure shown below.</p> <pre>typedef struct { uint32_f csize; /* cache size */ uint32_f bsize; /* cache bank size */ uint32_f lsize; /* cache line size */ uint32_f nbanks; /* number of banks (csize = bsize * nbanks) */ uint32_f type; /* cache type */ } SparcPropCache;</pre> <p>The <i>csize</i> field specifies the cache size in bytes.</p> <p>The <i>bsize</i> field specifies the cache bank size in bytes.</p> <p>The <i>lsize</i> field specifies the cache line size in bytes.</p> <p>The <i>nbanks</i> field specifies the number of banks.</p>

The *type* field specifies the cache properties. It is composed of the following bit-fields and flags:

CACHE_TYPE_IDX	Bit-field specifying the cache indexing mode: CACHE_TYPE_IDX_PHYS — physically indexed CACHE_TYPE_IDX_VIRT — virtually indexed
CACHE_TYPE_TAG	Bit-field specifying the cache tagging mode: CACHE_TYPE_TAG_PHYS — physically tagged CACHE_TYPE_TAG_VIRT — virtually tagged
CACHE_TYPE_MODE	Bit-field specifying the data cache mode: CACHE_TYPE_MODE_WT — write-through CACHE_TYPE_MODE_CB — copy-back
CACHE_TYPE_CFS	Flag specifying that the entire cache flush/invalidation is supported.
CACHE_TYPE_LFS	Flag specifying that line cache flush/invalidation is supported.
CACHE_TYPE_MATCH	Flag specifying that the tag match criteria is used for the line flush/invalidation

Note that, typically, a host bus driver does not need to examine the caches' properties because they are already taken into account by the cache management service routines described below.

`icacheInval` invalidates the entire CPU Instruction Cache.

`icacheLineInval` invalidates a given cache line within the CPU Instruction Cache. The cache line being invalidated is specified by a virtual address.

`icacheBlockInval` invalidates a given range within the CPU Instruction Cache. The range being invalidated is specified by the virtual start address and range size.

`dcacheFlush` flushes and invalidates the entire CPU Data Cache.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
icacheInval	+	+	+	-
icacheLineInval	+	+	+	-
icacheBlockInval	+	+	+	-
dcacheFlush	+	+	+	-
dcacheLineFlush	+	+	+	-
dcacheBlockFlush	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME | icacheInval, icacheLineInval, icacheBlockInval, dcacheFlush, dcacheLineFlush, dcacheBlockFlush – cache management

FEATURES | DKI

DESCRIPTION | The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.

Refer to architecture specific man pages:

- icacheInval_usparc(9DKI)
- icacheLineInval_usparc(9DKI)
- icacheBlockInval_usparc(9DKI)
- dcacheFlush_usparc(9DKI)
- dcacheLineFlush_usparc(9DKI)
- dcacheBlockFlush_usparc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	icacheInval_usparc, icacheLineInval_usparc, icacheBlockInval_usparc, dcacheFlush_usparc, dcacheLineFlush_usparc, dcacheBlockFlush_usparc – UltraSPARC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void icacheInval(void); void icacheLineInval(VmAddr addr); void icacheBlockInval(VmAddr addr, VmSize size); void dcacheFlush(void); void dcacheLineFlush(VmAddr addr); void dcacheBlockFlush(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.
EXTENDED DESCRIPTION	<p>UltraSPARC family processors have split instruction and data caches which are directly mapped and virtually indexed. The instruction cache is two-way associative, directly mapped. It is composed of two banks with an LRU replacement mechanism. The data cache is one-way associative, directly mapped and operates in the write-through mode.</p> <p>The data/instruction caches configuration is available as the "i-cache" and "d-cache" properties attached to the "cpu". The cpu node may be found in the device tree as a child node of the root node. The cpu node name is "cpu".</p> <p>The cache configuration property value is the SparcPropCache structure shown below.</p> <pre>typedef struct { uint32_f csize; /* cache size */ uint32_f bsize; /* cache bank size */ uint32_f lsize; /* cache line size */ uint32_f nbanks; /* number of banks (csize = bsize * nbanks) */ uint32_f type; /* cache type */ } SparcPropCache;</pre> <p>The <i>csize</i> field specifies the cache size in bytes.</p> <p>The <i>bsize</i> field specifies the cache bank size in bytes.</p> <p>The <i>lsize</i> field specifies the cache line size in bytes.</p> <p>The <i>nbanks</i> field specifies the number of banks.</p>

The *type* field specifies the cache properties. It is composed of the following bit-fields and flags:

CACHE_TYPE_IDX	Bit-field specifying the cache indexing mode: CACHE_TYPE_IDX_PHYS — physically indexed CACHE_TYPE_IDX_VIRT — virtually indexed
CACHE_TYPE_TAG	Bit-field specifying the cache tagging mode: CACHE_TYPE_TAG_PHYS — physically tagged CACHE_TYPE_TAG_VIRT — virtually tagged
CACHE_TYPE_MODE	Bit-field specifying the data cache mode: CACHE_TYPE_MODE_WT — write-through CACHE_TYPE_MODE_CB — copy-back
CACHE_TYPE_CFS	Flag specifying that the entire cache flush/invalidation is supported.
CACHE_TYPE_LFS	Flag specifying that line cache flush/invalidation is supported.
CACHE_TYPE_MATCH	Flag specifying that the tag match criteria is used for the line flush/invalidation

Note that, typically, a host bus driver does not need to examine the caches' properties because they are already taken into account by the cache management service routines described below.

`icacheInval` invalidates the entire CPU Instruction Cache.

`icacheLineInval` invalidates a given cache line within the CPU Instruction Cache. The cache line being invalidated is specified by a virtual address.

`icacheBlockInval` invalidates a given range within the CPU Instruction Cache. The range being invalidated is specified by the virtual start address and range size.

`dcacheFlush` flushes and invalidates the entire CPU Data Cache.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

`dcacheLineFlush` flushes and invalidates a given cache line within the CPU Data Cache. The cache line being flushed and invalidated is specified by a virtual address.

**ALLOWED
CALLING
CONTEXTS**

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
icacheInval	+	+	+	-
icacheLineInval	+	+	+	-
icacheBlockInval	+	+	+	-
dcacheFlush	+	+	+	-
dcacheLineFlush	+	+	+	-
dcacheBlockFlush	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME | imsIntrMask_f, imsIntrUnmask_f – global interrupts masking

SYNOPSIS | #include <dki/dki.h>
 void **imsIntrMask_f** (void);
 void **imsIntrUnmask_f** (void);

FEATURES | DKI

DESCRIPTION | Provides global interrupts masking operations.

EXTENDED DESCRIPTION | Some of the Interrupt Management Service (IMS) routines are part of the DKI in order to provide drivers with global interrupts masking services. These services can be useful for a driver to protect a critical section from interrupts, if necessary.

imsIntrMask_f masks all maskable interrupts at processor level, and increments the imsIntrMaskCount_f kernel variable.

imsIntrUnmask_f decrements the imsIntrMaskCount_f kernel variable, and unmask interrupts at processor level, if imsIntrMaskCount_f becomes equal to zero. Note that, as they rely on a kernel interrupt mask count, a driver may issue nested calls to these services.

Allowed Calling Contexts | The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
imsIntrMask_f	+	+	+	-
imsIntrUnmask_f	+	+	+	-

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO | svDkiThreadCall(9DKI)

NAME	imsIntrMask_f, imsIntrUnmask_f – global interrupts masking															
SYNOPSIS	<pre>#include <dki/dki.h> void imsIntrMask_f (void); void imsIntrUnmask_f (void);</pre>															
FEATURES	DKI															
DESCRIPTION	Provides global interrupts masking operations.															
EXTENDED DESCRIPTION	<p>Some of the Interrupt Management Service (IMS) routines are part of the DKI in order to provide drivers with global interrupts masking services. These services can be useful for a driver to protect a critical section from interrupts, if necessary.</p> <p>imsIntrMask_f masks all maskable interrupts at processor level, and increments the imsIntrMaskCount_f kernel variable.</p> <p>imsIntrUnmask_f decrements the imsIntrMaskCount_f kernel variable, and unmask interrupts at processor level, if imsIntrMaskCount_f becomes equal to zero. Note that, as they rely on a kernel interrupt mask count, a driver may issue nested calls to these services.</p>															
Allowed Calling Contexts	<p>The following table specifies the contexts in which a caller is allowed to invoke each service.</p> <table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="text-align: left;">Services</th> <th>Base level</th> <th>DKI thread</th> <th>Interrupt</th> <th>Blocking</th> </tr> </thead> <tbody> <tr> <td>imsIntrMask_f</td> <td>+</td> <td>+</td> <td>+</td> <td>-</td> </tr> <tr> <td>imsIntrUnmask_f</td> <td>+</td> <td>+</td> <td>+</td> <td>-</td> </tr> </tbody> </table>	Services	Base level	DKI thread	Interrupt	Blocking	imsIntrMask_f	+	+	+	-	imsIntrUnmask_f	+	+	+	-
Services	Base level	DKI thread	Interrupt	Blocking												
imsIntrMask_f	+	+	+	-												
imsIntrUnmask_f	+	+	+	-												
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="text-align: left;">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	svDkiThreadCall(9DKI)															

NAME dataCacheBlockFlush, dataCacheInvalidate, dataCacheBlockInvalidate, instCacheInvalidate, instCacheBlockInvalidate – cache management

FEATURES DKI

DESCRIPTION See the architecture specific man pages:

- dataCacheBlockFlush_powerpc(9DKI)
- dataCacheInvalidate_powerpc(9DKI)
- dataCacheBlockInvalidate_powerpc(9DKI)
- instCacheInvalidate_powerpc(9DKI)
- instCacheBlockInvalidate_powerpc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	dataCacheBlockFlush_powerpc, dataCacheInvalidate_powerpc, dataCacheBlockInvalidate_powerpc, instCacheInvalidate_powerpc, instCacheBlockInvalidate_powerpc – PowerPC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void dataCacheFlush(void); void dataCacheBlockFlush(VmAddr addr, VmSize size); void dataCacheInvalidate(void); void dataCacheBlockInvalidate(VmAddr addr, VmSize size); void instCacheInvalidate(void); void instCacheBlockInvalidate(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	Provides PowerPC cache management services.
EXTENDED DESCRIPTION	<p>The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.</p> <p>Typically, PowerPC family processors have separate instruction and data caches which are virtually indexed and physically tagged. However, the PowerPC architecture does not specify the type or existence of a cache. This allows for various different cache types (unified, or no cache at all). In any case, cache management services should behave consistently. They should do nothing if there is no cache, and data or instruction invalidation routines should be equivalent for a unified cache.</p> <hr/> <p>Note - The data cache may operate in either write-through or copy-back mode, on a per line basis, depending on the cached memory attributes. The data/instruction cache size and data/instruction line size are processor implementation specific.</p> <hr/> <p>The data/instruction cache configuration is available from the PPC_PROP_CACHE property attached to the NODE_CPU node. (The cpu node may be found in the device tree as a child node of the root node. The cpu node name is NODE_CPU.)</p> <p>The cache configuration property value is the PpcPropCache structure.</p> <p>The <i>blockNumber</i> field specifies the number of cache blocks in each data or instruction cache.</p> <p>The <i>blockSize</i> field specifies the cache block size in bytes.</p> <p>The <i>blockSizeShift</i> field specifies the number of bits to shift right/left to divide/multiply by the cache block size (cache block size is always a power of 2).</p>

`dataCacheFlush` globally flushes and invalidates the data cache.

`dataCacheBlockFlush` flushes and invalidates a given range of addresses within the CPU data cache. The range being flushed is specified by the virtual start address (within the current MMU context) and the range size.

`dataCacheInvalidate` globally invalidates the data cache. Note that all blocks in the data cache are marked as invalid without writing back any modified lines to memory. This function does nothing if the data cache is disabled.

`dataCacheBlockInvalidate` invalidates a given range of addresses within the CPU data cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that invalidated blocks are not written back to memory. Note also that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

`instCacheInvalidate` globally invalidates the instruction cache; that is, all blocks in the instruction cache are marked as invalid. This function does nothing if the instruction cache is disabled.

`instCacheBlockInvalidate` invalidates a given range of addresses within the CPU instruction cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dataCacheFlush</code>	+	+	+	-
<code>dataCacheBlockFlush</code>	+	+	+	-
<code>dataCacheInvalidate</code>	+	+	+	-
<code>dataCacheBlockInvalidate</code>	+	+	+	-
<code>instCacheInvalidate</code>	+	+	+	-
<code>instCacheBlockInvalidate</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME dataCacheBlockFlush, dataCacheInvalidate, dataCacheBlockInvalidate, instCacheInvalidate, instCacheBlockInvalidate – cache management

FEATURES DKI

DESCRIPTION See the architecture specific man pages:

- dataCacheBlockFlush_powerpc(9DKI)
- dataCacheInvalidate_powerpc(9DKI)
- dataCacheBlockInvalidate_powerpc(9DKI)
- instCacheInvalidate_powerpc(9DKI)
- instCacheBlockInvalidate_powerpc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	dataCacheBlockFlush_powerpc, dataCacheInvalidate_powerpc, dataCacheBlockInvalidate_powerpc, instCacheInvalidate_powerpc, instCacheBlockInvalidate_powerpc – PowerPC cache management
SYNOPSIS	<pre>#include <dki/f_dki.h> void dataCacheFlush(void); void dataCacheBlockFlush(VmAddr addr, VmSize size); void dataCacheInvalidate(void); void dataCacheBlockInvalidate(VmAddr addr, VmSize size); void instCacheInvalidate(void); void instCacheBlockInvalidate(VmAddr addr, VmSize size);</pre>
FEATURES	DKI
DESCRIPTION	Provides PowerPC cache management services.
EXTENDED DESCRIPTION	<p>The microkernel provides cache management services, mainly to allow host bus drivers to manage memory coherency for DMA purposes.</p> <p>Typically, PowerPC family processors have separate instruction and data caches which are virtually indexed and physically tagged. However, the PowerPC architecture does not specify the type or existence of a cache. This allows for various different cache types (unified, or no cache at all). In any case, cache management services should behave consistently. They should do nothing if there is no cache, and data or instruction invalidation routines should be equivalent for a unified cache.</p> <hr/> <p>Note - The data cache may operate in either write-through or copy-back mode, on a per line basis, depending on the cached memory attributes. The data/instruction cache size and data/instruction line size are processor implementation specific.</p> <hr/> <p>The data/instruction cache configuration is available from the PPC_PROP_CACHE property attached to the NODE_CPU node. (The cpu node may be found in the device tree as a child node of the root node. The cpu node name is NODE_CPU.)</p> <p>The cache configuration property value is the PpcPropCache structure.</p> <p>The <i>blockNumber</i> field specifies the number of cache blocks in each data or instruction cache.</p> <p>The <i>blockSize</i> field specifies the cache block size in bytes.</p> <p>The <i>blockSizeShift</i> field specifies the number of bits to shift right/left to divide/multiply by the cache block size (cache block size is always a power of 2).</p>

`dataCacheFlush` globally flushes and invalidates the data cache.

`dataCacheBlockFlush` flushes and invalidates a given range of addresses within the CPU data cache. The range being flushed is specified by the virtual start address (within the current MMU context) and the range size.

`dataCacheInvalidate` globally invalidates the data cache. Note that all blocks in the data cache are marked as invalid without writing back any modified lines to memory. This function does nothing if the data cache is disabled.

`dataCacheBlockInvalidate` invalidates a given range of addresses within the CPU data cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that invalidated blocks are not written back to memory. Note also that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

`instCacheInvalidate` globally invalidates the instruction cache; that is, all blocks in the instruction cache are marked as invalid. This function does nothing if the instruction cache is disabled.

`instCacheBlockInvalidate` invalidates a given range of addresses within the CPU instruction cache. The range being invalidated is specified by the virtual start address (within the current MMU context) and the range size.

Note that this service should only be used to invalidate a small memory range; that is, a memory range smaller than the cache size. Otherwise, global invalidation should be used.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>dataCacheFlush</code>	+	+	+	-
<code>dataCacheBlockFlush</code>	+	+	+	-
<code>dataCacheInvalidate</code>	+	+	+	-
<code>dataCacheBlockInvalidate</code>	+	+	+	-
<code>instCacheInvalidate</code>	+	+	+	-
<code>instCacheBlockInvalidate</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME | ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services

FEATURES | DKI

DESCRIPTION | The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

See the architecture specific man pages:

- ioLoad8_x86(9DKI)
- ioStore8_x86(9DKI)
- ioRead8_x86(9DKI)
- ioWrite8_x86(9DKI)
- ioLoad16_x86(9DKI)
- ioStore16_x86(9DKI)
- ioRead16_x86(9DKI)
- ioWrite16_x86(9DKI)
- ioLoad32_x86(9DKI)
- ioStore32_x86(9DKI)
- ioRead32_x86(9DKI)
- ioWrite32_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services
SYNOPSIS	<pre>#include <dki/dki.h> uint8_f ioLoad8(uint32_f base, uint32_f offset); void ioStore8(uint32_f base, uint32_f offset, uint8_f value); void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); uint16_f ioLoad16(uint32_f base, uint32_f offset); void ioStore16(uint32_f base, uint32_f offset, uint16_f value); void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); uint32_f ioLoad32(uint32_f base, uint32_f offset); void ioStore32(uint32_f base, uint32_f offset, uint32_f value); void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count); void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);</pre>
FEATURES	DKI
EXTENDED DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>Specific I/O services are defined below as sets of routines where the xx suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:</p> <ul style="list-style-type: none"> ■ 8 for 8-bit data ■ 16 for 16-bit data ■ 32 for 32-bit data <p>ioLoadxx</p> <p>The <code>ioLoadxx</code> loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The <i>base</i> argument specifies the base address. The <i>offset</i> argument specifies the offset from the base address.</p> <p>ioStorexx</p>

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services

FEATURES DKI

DESCRIPTION The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

See the architecture specific man pages:

- ioLoad8_x86(9DKI)
- ioStore8_x86(9DKI)
- ioRead8_x86(9DKI)
- ioWrite8_x86(9DKI)
- ioLoad16_x86(9DKI)
- ioStore16_x86(9DKI)
- ioRead16_x86(9DKI)
- ioWrite16_x86(9DKI)
- ioLoad32_x86(9DKI)
- ioStore32_x86(9DKI)
- ioRead32_x86(9DKI)
- ioWrite32_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME | ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services

SYNOPSIS

```
#include <dki/dki.h>
uint8_f ioLoad8(uint32_f base, uint32_f offset);

void ioStore8(uint32_f base, uint32_f offset, uint8_f value);

void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

uint16_f ioLoad16(uint32_f base, uint32_f offset);

void ioStore16(uint32_f base, uint32_f offset, uint16_f value);

void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

uint32_f ioLoad32(uint32_f base, uint32_f offset);

void ioStore32(uint32_f base, uint32_f offset, uint32_f value);

void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);

void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);
```

FEATURES

DKI

EXTENDED DESCRIPTION

The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

Specific I/O services are defined below as sets of routines where the **xx** suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:

- 8 for 8-bit data
- 16 for 16-bit data
- 32 for 32-bit data

ioLoadxx

The `ioLoadxx` loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address.

ioStorexx

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services

FEATURES | DK1

DESCRIPTION | The DK1 provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

See the architecture specific man pages:

- ioLoad8_x86(9DKI)
- ioStore8_x86(9DKI)
- ioRead8_x86(9DKI)
- ioWrite8_x86(9DKI)
- ioLoad16_x86(9DKI)
- ioStore16_x86(9DKI)
- ioRead16_x86(9DKI)
- ioWrite16_x86(9DKI)
- ioLoad32_x86(9DKI)
- ioStore32_x86(9DKI)
- ioRead32_x86(9DKI)
- ioWrite32_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services
SYNOPSIS	<pre>#include <dki/dki.h> uint8_f ioLoad8(uint32_f base, uint32_f offset); void ioStore8(uint32_f base, uint32_f offset, uint8_f value); void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); uint16_f ioLoad16(uint32_f base, uint32_f offset); void ioStore16(uint32_f base, uint32_f offset, uint16_f value); void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); uint32_f ioLoad32(uint32_f base, uint32_f offset); void ioStore32(uint32_f base, uint32_f offset, uint32_f value); void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count); void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);</pre>
FEATURES	DKI
EXTENDED DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>Specific I/O services are defined below as sets of routines where the xx suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:</p> <ul style="list-style-type: none"> ■ 8 for 8-bit data ■ 16 for 16-bit data ■ 32 for 32-bit data <p>ioLoadxx</p> <p>The <code>ioLoadxx</code> loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The <i>base</i> argument specifies the base address. The <i>offset</i> argument specifies the offset from the base address.</p> <p>ioStorexx</p>

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services

FEATURES DKI

DESCRIPTION The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

See the architecture specific man pages:

- ioLoad8_x86(9DKI)
- ioStore8_x86(9DKI)
- ioRead8_x86(9DKI)
- ioWrite8_x86(9DKI)
- ioLoad16_x86(9DKI)
- ioStore16_x86(9DKI)
- ioRead16_x86(9DKI)
- ioWrite16_x86(9DKI)
- ioLoad32_x86(9DKI)
- ioStore32_x86(9DKI)
- ioRead32_x86(9DKI)
- ioWrite32_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME | ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services

SYNOPSIS

```
#include <dki/dki.h>
uint8_f ioLoad8(uint32_f base, uint32_f offset);

void ioStore8(uint32_f base, uint32_f offset, uint8_f value);

void ioRead8(uint32_f base, uint32_f offset, uint8_f* addr, uint32_f count);

void ioWrite8(uint32_f base, uint32_f offset, uint8_f* addr, uint32_f count);

uint16_f ioLoad16(uint32_f base, uint32_f offset);

void ioStore16(uint32_f base, uint32_f offset, uint16_f value);

void ioRead16(uint32_f base, uint32_f offset, uint16_f* addr, uint32_f count);

void ioWrite16(uint32_f base, uint32_f offset, uint16_f* addr, uint32_f count);

uint32_f ioLoad32(uint32_f base, uint32_f offset);

void ioStore32(uint32_f base, uint32_f offset, uint32_f value);

void ioRead32(uint32_f base, uint32_f offset, uint32_f* addr, uint32_f count);

void ioWrite32(uint32_f base, uint32_f offset, uint32_f* addr, uint32_f count);
```

FEATURES

DKI

EXTENDED DESCRIPTION

The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

Specific I/O services are defined below as sets of routines where the **xx** suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:

- 8 for 8-bit data
- 16 for 16-bit data
- 32 for 32-bit data

ioLoadxx

The *ioLoadxx* loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address.

ioStorexx

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services

FEATURES | DKI

DESCRIPTION | The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

See the architecture specific man pages:

- ioLoad8_x86(9DKI)
- ioStore8_x86(9DKI)
- ioRead8_x86(9DKI)
- ioWrite8_x86(9DKI)
- ioLoad16_x86(9DKI)
- ioStore16_x86(9DKI)
- ioRead16_x86(9DKI)
- ioWrite16_x86(9DKI)
- ioLoad32_x86(9DKI)
- ioStore32_x86(9DKI)
- ioRead32_x86(9DKI)
- ioWrite32_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services
SYNOPSIS	<pre>#include <dki/dki.h> uint8_f ioLoad8(uint32_f base, uint32_f offset); void ioStore8(uint32_f base, uint32_f offset, uint8_f value); void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); uint16_f ioLoad16(uint32_f base, uint32_f offset); void ioStore16(uint32_f base, uint32_f offset, uint16_f value); void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); uint32_f ioLoad32(uint32_f base, uint32_f offset); void ioStore32(uint32_f base, uint32_f offset, uint32_f value); void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count); void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);</pre>
FEATURES	DKI
EXTENDED DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>Specific I/O services are defined below as sets of routines where the xx suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:</p> <ul style="list-style-type: none"> ■ 8 for 8-bit data ■ 16 for 16-bit data ■ 32 for 32-bit data <p>ioLoadxx</p> <p>The <code>ioLoadxx</code> loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The <i>base</i> argument specifies the base address. The <i>offset</i> argument specifies the offset from the base address.</p> <p>ioStorexx</p>

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services
FEATURES	DKI
DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ ioLoad8_x86(9DKI) ■ ioStore8_x86(9DKI) ■ ioRead8_x86(9DKI) ■ ioWrite8_x86(9DKI) ■ ioLoad16_x86(9DKI) ■ ioStore16_x86(9DKI) ■ ioRead16_x86(9DKI) ■ ioWrite16_x86(9DKI) ■ ioLoad32_x86(9DKI) ■ ioStore32_x86(9DKI) ■ ioRead32_x86(9DKI) ■ ioWrite32_x86(9DKI)

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME | ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services

SYNOPSIS

```
#include <dki/dki.h>
uint8_f ioLoad8(uint32_f base, uint32_f offset);

void ioStore8(uint32_f base, uint32_f offset, uint8_f value);

void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

uint16_f ioLoad16(uint32_f base, uint32_f offset);

void ioStore16(uint32_f base, uint32_f offset, uint16_f value);

void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

uint32_f ioLoad32(uint32_f base, uint32_f offset);

void ioStore32(uint32_f base, uint32_f offset, uint32_f value);

void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);

void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);
```

FEATURES

DKI

EXTENDED DESCRIPTION

The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

Specific I/O services are defined below as sets of routines where the **xx** suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:

- 8 for 8-bit data
- 16 for 16-bit data
- 32 for 32-bit data

ioLoadxx

The *ioLoadxx* loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address.

ioStorexx

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services

FEATURES | DKI

DESCRIPTION | The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.
 See the architecture specific man pages:

- ioLoad8_x86(9DKI)
- ioStore8_x86(9DKI)
- ioRead8_x86(9DKI)
- ioWrite8_x86(9DKI)
- ioLoad16_x86(9DKI)
- ioStore16_x86(9DKI)
- ioRead16_x86(9DKI)
- ioWrite16_x86(9DKI)
- ioLoad32_x86(9DKI)
- ioStore32_x86(9DKI)
- ioRead32_x86(9DKI)
- ioWrite32_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services
SYNOPSIS	<pre>#include <dki/dki.h> uint8_f ioLoad8(uint32_f base, uint32_f offset); void ioStore8(uint32_f base, uint32_f offset, uint8_f value); void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); uint16_f ioLoad16(uint32_f base, uint32_f offset); void ioStore16(uint32_f base, uint32_f offset, uint16_f value); void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); uint32_f ioLoad32(uint32_f base, uint32_f offset); void ioStore32(uint32_f base, uint32_f offset, uint32_f value); void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count); void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);</pre>
FEATURES	DKI
EXTENDED DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>Specific I/O services are defined below as sets of routines where the xx suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:</p> <ul style="list-style-type: none"> ■ 8 for 8-bit data ■ 16 for 16-bit data ■ 32 for 32-bit data <p>ioLoadxx</p> <p>The <code>ioLoadxx</code> loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The <i>base</i> argument specifies the base address. The <i>offset</i> argument specifies the offset from the base address.</p> <p>ioStorexx</p>

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services
FEATURES	DKI
DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ ioLoad8_x86(9DKI) ■ ioStore8_x86(9DKI) ■ ioRead8_x86(9DKI) ■ ioWrite8_x86(9DKI) ■ ioLoad16_x86(9DKI) ■ ioStore16_x86(9DKI) ■ ioRead16_x86(9DKI) ■ ioWrite16_x86(9DKI) ■ ioLoad32_x86(9DKI) ■ ioStore32_x86(9DKI) ■ ioRead32_x86(9DKI) ■ ioWrite32_x86(9DKI)

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME | ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services

SYNOPSIS

```
#include <dki/dki.h>
uint8_f ioLoad8(uint32_f base, uint32_f offset);

void ioStore8(uint32_f base, uint32_f offset, uint8_f value);

void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

uint16_f ioLoad16(uint32_f base, uint32_f offset);

void ioStore16(uint32_f base, uint32_f offset, uint16_f value);

void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

uint32_f ioLoad32(uint32_f base, uint32_f offset);

void ioStore32(uint32_f base, uint32_f offset, uint32_f value);

void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);

void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);
```

FEATURES

DKI

EXTENDED DESCRIPTION

The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

Specific I/O services are defined below as sets of routines where the **xx** suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:

- 8 for 8-bit data
- 16 for 16-bit data
- 32 for 32-bit data

ioLoadxx

The *ioLoadxx* loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address.

ioStorexx

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services

FEATURES | DKI

DESCRIPTION | The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

See the architecture specific man pages:

- ioLoad8_x86(9DKI)
- ioStore8_x86(9DKI)
- ioRead8_x86(9DKI)
- ioWrite8_x86(9DKI)
- ioLoad16_x86(9DKI)
- ioStore16_x86(9DKI)
- ioRead16_x86(9DKI)
- ioWrite16_x86(9DKI)
- ioLoad32_x86(9DKI)
- ioStore32_x86(9DKI)
- ioRead32_x86(9DKI)
- ioWrite32_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services
SYNOPSIS	<pre>#include <dki/dki.h> uint8_f ioLoad8(uint32_f base, uint32_f offset); void ioStore8(uint32_f base, uint32_f offset, uint8_f value); void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); uint16_f ioLoad16(uint32_f base, uint32_f offset); void ioStore16(uint32_f base, uint32_f offset, uint16_f value); void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); uint32_f ioLoad32(uint32_f base, uint32_f offset); void ioStore32(uint32_f base, uint32_f offset, uint32_f value); void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count); void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);</pre>
FEATURES	DKI
EXTENDED DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>Specific I/O services are defined below as sets of routines where the xx suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:</p> <ul style="list-style-type: none"> ■ 8 for 8-bit data ■ 16 for 16-bit data ■ 32 for 32-bit data <p>ioLoadxx</p> <p>The <code>ioLoadxx</code> loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The <i>base</i> argument specifies the base address. The <i>offset</i> argument specifies the offset from the base address.</p> <p>ioStorexx</p>

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services
FEATURES	DKI
DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ ioLoad8_x86(9DKI) ■ ioStore8_x86(9DKI) ■ ioRead8_x86(9DKI) ■ ioWrite8_x86(9DKI) ■ ioLoad16_x86(9DKI) ■ ioStore16_x86(9DKI) ■ ioRead16_x86(9DKI) ■ ioWrite16_x86(9DKI) ■ ioLoad32_x86(9DKI) ■ ioStore32_x86(9DKI) ■ ioRead32_x86(9DKI) ■ ioWrite32_x86(9DKI)

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME | ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services

SYNOPSIS

```
#include <dki/dki.h>
uint8_f ioLoad8(uint32_f base, uint32_f offset);

void ioStore8(uint32_f base, uint32_f offset, uint8_f value);

void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

uint16_f ioLoad16(uint32_f base, uint32_f offset);

void ioStore16(uint32_f base, uint32_f offset, uint16_f value);

void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

uint32_f ioLoad32(uint32_f base, uint32_f offset);

void ioStore32(uint32_f base, uint32_f offset, uint32_f value);

void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);

void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);
```

FEATURES

DKI

EXTENDED DESCRIPTION

The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

Specific I/O services are defined below as sets of routines where the **xx** suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:

- 8 for 8-bit data
- 16 for 16-bit data
- 32 for 32-bit data

ioLoadxx

The *ioLoadxx* loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address.

ioStorexx

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services

FEATURES | DKI

DESCRIPTION | The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.
 See the architecture specific man pages:

- ioLoad8_x86(9DKI)
- ioStore8_x86(9DKI)
- ioRead8_x86(9DKI)
- ioWrite8_x86(9DKI)
- ioLoad16_x86(9DKI)
- ioStore16_x86(9DKI)
- ioRead16_x86(9DKI)
- ioWrite16_x86(9DKI)
- ioLoad32_x86(9DKI)
- ioStore32_x86(9DKI)
- ioRead32_x86(9DKI)
- ioWrite32_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services
SYNOPSIS	<pre>#include <dki/dki.h> uint8_f ioLoad8(uint32_f base, uint32_f offset); void ioStore8(uint32_f base, uint32_f offset, uint8_f value); void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count); uint16_f ioLoad16(uint32_f base, uint32_f offset); void ioStore16(uint32_f base, uint32_f offset, uint16_f value); void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count); uint32_f ioLoad32(uint32_f base, uint32_f offset); void ioStore32(uint32_f base, uint32_f offset, uint32_f value); void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count); void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);</pre>
FEATURES	DKI
EXTENDED DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>Specific I/O services are defined below as sets of routines where the xx suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:</p> <ul style="list-style-type: none"> ■ 8 for 8-bit data ■ 16 for 16-bit data ■ 32 for 32-bit data <p>ioLoadxx</p> <p>The <code>ioLoadxx</code> loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The <i>base</i> argument specifies the base address. The <i>offset</i> argument specifies the offset from the base address.</p> <p>ioStorexx</p>

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	ioLoad8, ioStore8, ioRead8, ioWrite8, ioLoad16, ioStore16, ioRead16, ioWrite16, ioLoad32, ioStore32, ioRead32, ioWrite32 – I/O services
FEATURES	DKI
DESCRIPTION	<p>The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ ioLoad8_x86(9DKI) ■ ioStore8_x86(9DKI) ■ ioRead8_x86(9DKI) ■ ioWrite8_x86(9DKI) ■ ioLoad16_x86(9DKI) ■ ioStore16_x86(9DKI) ■ ioRead16_x86(9DKI) ■ ioWrite16_x86(9DKI) ■ ioLoad32_x86(9DKI) ■ ioStore32_x86(9DKI) ■ ioRead32_x86(9DKI) ■ ioWrite32_x86(9DKI)

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME | ioLoad8_x86, ioStore8_x86, ioRead8_x86, ioWrite8_x86, ioLoad16_x86, ioStore16_x86, ioRead16_x86, ioWrite16_x86, ioLoad32_x86, ioStore32_x86, ioRead32_x86, ioWrite32_x86 – Intel x86 specific I/O services

SYNOPSIS

```
#include <dki/dki.h>
uint8_f ioLoad8(uint32_f base, uint32_f offset);

void ioStore8(uint32_f base, uint32_f offset, uint8_f value);

void ioRead8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

void ioWrite8(uint32_f base, uint32_f offset, uint8_f * addr, uint32_f count);

uint16_f ioLoad16(uint32_f base, uint32_f offset);

void ioStore16(uint32_f base, uint32_f offset, uint16_f value);

void ioRead16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

void ioWrite16(uint32_f base, uint32_f offset, uint16_f * addr, uint32_f count);

uint32_f ioLoad32(uint32_f base, uint32_f offset);

void ioStore32(uint32_f base, uint32_f offset, uint32_f value);

void ioRead32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);

void ioWrite32(uint32_f base, uint32_f offset, uint32_f * addr, uint32_f count);
```

FEATURES

DKI

EXTENDED DESCRIPTION

The DKI provides specific I/O routines which can be used by a host bus driver to implement bus I/O operations.

Specific I/O services are defined below as sets of routines where the **xx** suffix indicates the bit length of the data to which the service applies. This suffix may take one of the following values:

- 8 for 8-bit data
- 16 for 16-bit data
- 32 for 32-bit data

ioLoadxx

The *ioLoadxx* loads data from a given I/O address and returns the read value. The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address.

ioStorexx

The `ioStorexx` stores a specified value at a given I/O address. The I/O address written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *value* argument specifies the value to store (write).

`ioReadxx`

The `ioReadxx` reads data from a given I/O address *count* times and stores values at the memory location pointed to by *addr*.

The read location is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *addr* argument specifies the location where read values will be stored. The *count* argument specifies the desired number of reads.

`ioWritexx`

The `ioWritexx` reads *count* data from the memory location at *addr* and writes them at the I/O address specified by *base + offset*.

The read location is specified by the parameter *addr*. The I/O location written is composed of a base address and an offset from this base. The *base* argument specifies the base address. The *offset* argument specifies the offset from the base address. The *count* argument specifies the desired number of writes to perform. The table below specifies the contexts in which a caller is allowed to invoke each service:

Specific I/O Allowed Calling Contexts

Services	Base level	DKI thread	Interrupt	Blocking
<code>ioLoadxx</code>	+	+	+	-
<code>ioStorexx</code>	+	+	+	-
<code>ioReadxx</code>	+	+	+	-
<code>ioWritexx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | loadSwap_16, storeSwap_16, swap_16, loadSwap_32, storeSwap_32, swap_32, loadSwap_64, storeSwap_64, swap_64 – specific i/o services

SYNOPSIS

```
#include <dki/dki.h>
uint16_f loadSwap_16(uint16_f * addr);

void storeSwap_16(uint16_f * addr, uint16_f value);

void swap_16(uint16_f * addr);

uint32_f loadSwap_32(uint32_f * addr);

void storeSwap_32(uint32_f * addr, uint32_f value);

void swap_32(uint32_f * addr);

uint64_f loadSwap_64(uint64_f * addr);

void storeSwap_64(uint64_f * addr, uint64_f value);

void swap_64(uint64_f * addr);
```

FEATURES | DKI

DESCRIPTION | This function provides DKI routines to handle different byte ordering between the processor bus and the host bus.

EXTENDED DESCRIPTION | Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus. Specific I/O services are defined below as sets of routines where the *_xx* suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- _16* for 16-bit data
- _32* for 32-bit data
- _64* for 64-bit data

loadSwap_xx loads data from a given address and returns the corresponding byte swapped value. The *addr* argument specifies the address to read from.

storeSwap_xx stores into a given address the *value* byte swapped. The *addr* argument specifies the address to write to.

swap_xx swap in place the bytes of the data stored at a given address. The *addr* argument specifies the address of the data to byte-swap.

Allowed Calling Contexts | The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
loadSwap_xx	+	+	+	-

storeSwap_xx	+	+	+	-
swap_xx	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | loadSwap_16, storeSwap_16, swap_16, loadSwap_32, storeSwap_32, swap_32, loadSwap_64, storeSwap_64, swap_64 – specific i/o services

SYNOPSIS

```
#include <dki/dki.h>
uint16_f loadSwap_16(uint16_f * addr);

void storeSwap_16(uint16_f * addr, uint16_f value);

void swap_16(uint16_f * addr);

uint32_f loadSwap_32(uint32_f * addr);

void storeSwap_32(uint32_f * addr, uint32_f value);

void swap_32(uint32_f * addr);

uint64_f loadSwap_64(uint64_f * addr);

void storeSwap_64(uint64_f * addr, uint64_f value);

void swap_64(uint64_f * addr);
```

FEATURES | DKI

DESCRIPTION | This function provides DKI routines to handle different byte ordering between the processor bus and the host bus.

EXTENDED DESCRIPTION | Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus. Specific I/O services are defined below as sets of routines where the *_xx* suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- _16* for 16-bit data
- _32* for 32-bit data
- _64* for 64-bit data

loadSwap_xx loads data from a given address and returns the corresponding byte swapped value. The *addr* argument specifies the address to read from.

storeSwap_xx stores into a given address the *value* byte swapped. The *addr* argument specifies the address to write to.

swap_xx swap in place the bytes of the data stored at a given address. The *addr* argument specifies the address of the data to byte-swap.

Allowed Calling Contexts | The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
loadSwap_xx	+	+	+	-

storeSwap_xx	+	+	+	-
swap_xx	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | loadSwap_16, storeSwap_16, swap_16, loadSwap_32, storeSwap_32, swap_32, loadSwap_64, storeSwap_64, swap_64 – specific i/o services

SYNOPSIS

```
#include <dki/dki.h>
uint16_f loadSwap_16(uint16_f * addr);

void storeSwap_16(uint16_f * addr, uint16_f value);

void swap_16(uint16_f * addr);

uint32_f loadSwap_32(uint32_f * addr);

void storeSwap_32(uint32_f * addr, uint32_f value);

void swap_32(uint32_f * addr);

uint64_f loadSwap_64(uint64_f * addr);

void storeSwap_64(uint64_f * addr, uint64_f value);

void swap_64(uint64_f * addr);
```

FEATURES | DKI

DESCRIPTION | This function provides DKI routines to handle different byte ordering between the processor bus and the host bus.

EXTENDED DESCRIPTION | Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus. Specific I/O services are defined below as sets of routines where the *_xx* suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- _16* for 16-bit data
- _32* for 32-bit data
- _64* for 64-bit data

loadSwap_xx loads data from a given address and returns the corresponding byte swapped value. The *addr* argument specifies the address to read from.

storeSwap_xx stores into a given address the *value* byte swapped. The *addr* argument specifies the address to write to.

swap_xx swap in place the bytes of the data stored at a given address. The *addr* argument specifies the address of the data to byte-swap.

Allowed Calling Contexts | The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
loadSwap_xx	+	+	+	-

storeSwap_xx	+	+	+	-
swap_xx	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | loadSwapEieio_16, storeSwapEieio_16, swapEieio_16, loadSwapEieio_32, storeSwapEieio_32, swapEieio_32, eieio - i/o services

FEATURES | DKI

DESCRIPTION | Provides access to DKI i/o routines, optimized to facilitate byte swapping and synchronization of input/output.
 See the architecture specific man pages:

- loadSwapEieio_16_powerpc(9DKI)
- storeSwapEieio_16_powerpc(9DKI)
- swapEieio_16_powerpc(9DKI)
- loadSwapEieio_32_powerpc(9DKI)
- storeSwapEieio_32_powerpc(9DKI)
- swapEieio_32_powerpc(9DKI)
- eieio(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	loadSwapEieio_16_powerpc, loadSwapSync_16_powerpc, storeSwapEieio_16_powerpc, storeSwapSync_16_powerpc, swapEieio_16_powerpc, loadSwapEieio_32_powerpc, loadSwapSync_32_powerpc, storeSwapEieio_32_powerpc, storeSwapSync_32_powerpc, swapEieio_32_powerpc, eieio_powerpc, ioSync_powerpc – PowerPC specific i/o services
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwapEieio_16(uint16_f* addr); uint16_f loadSwapSync_16(uint16_f* addr); void storeSwapEieio_16(uint16_f* addr, uint16_f value); void storeSwapSync_16(uint16_f* addr, uint16_f value); void swapEieio_16(uint16_f* addr); uint32_f loadSwapEieio_32(uint32_f* addr); uint32_f loadSwapSync_32(uint32_f* addr); void storeSwapEieio_32(uint32_f* addr, uint32_f value); void storeSwapSync_32(uint32_f* addr, uint32_f value); void swapEieio_32 (uint32_f* addr); void ioSync(void); void eieio(void);</pre>
FEATURES	DKI
DESCRIPTION	Provides access to PowerPC DKI I/O routines, optimized to facilitate byte swapping and synchronization of input/output.
EXTENDED DESCRIPTION	<p>The PowerPC DKI provides specific I/O routines to handle byte swapping and synchronization of I/O. These services are based on PowerPC specific instructions to store/load with byte swapping, and to enforce in-order execution of I/O.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the PowerPC bus and the host bus, or to enforce sequential execution of I/O. PowerPC specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <pre>_16 for 16-bit data _32 for 32-bit data</pre>

`loadSwapEieio_xx` loads data from a given address and returns the corresponding byte swapped value. A PowerPC `eieio` instruction is issued after the data is read, to enforce execution of subsequent I/O sequentially with respect to the current load. The *addr* argument specifies the address to read from.

`loadSwapSync_xx` behaves like `loadSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`storeSwapEieio_xx` stores, at a given address, the *value* byte swapped. A PowerPC `eieio` instruction is issued after the data is written, to enforce execution of subsequent I/O sequentially with respect to the current store. The *addr* argument specifies the address to write to.

`storeSwapSync_xx` behaves like `storeSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`swapEieio_xx` swaps in place the bytes of the data stored at a given address. A PowerPC `eieio` instruction is issued between reading the data and re-writing the byte-swapped data, to enforce execution in order. The *addr* argument specifies the address of the data to byte-swap.

`ioSync` enforces context synchronization, completion of all instructions, and completion of all exceptions before returning.

`eieio` enforces execution in order of I/O and is mapped to the corresponding PowerPC assembler instruction. This routine is intended for use in managing shared data structures, in doing memory-mapped I/O, and in preventing load/store combining operations in main memory. Refer to PowerPC programming manuals for more details.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwapEieio_xx</code>	+	+	+	-
<code>loadSwapSync_xx</code>	+	+	+	-
<code>storeSwapEieio_xx</code>	+	+	+	-
<code>storeSwapSync_xx</code>	+	+	+	-
<code>swapEieio_xx</code>	+	+	+	-
<code>ioSync</code>	+	+	+	-
<code>eieio</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap_16(9DKI)` , `svDkiThreadCall(9DKI)`

NAME | loadSwapEieio_16, storeSwapEieio_16, swapEieio_16, loadSwapEieio_32, storeSwapEieio_32, swapEieio_32, eieio - i/o services

FEATURES | DKI

DESCRIPTION | Provides access to DKI i/o routines, optimized to facilitate byte swapping and synchronization of input/output.
 See the architecture specific man pages:

- loadSwapEieio_16_powerpc(9DKI)
- storeSwapEieio_16_powerpc(9DKI)
- swapEieio_16_powerpc(9DKI)
- loadSwapEieio_32_powerpc(9DKI)
- storeSwapEieio_32_powerpc(9DKI)
- swapEieio_32_powerpc(9DKI)
- eieio(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	loadSwapEieio_16_powerpc, loadSwapSync_16_powerpc, storeSwapEieio_16_powerpc, storeSwapSync_16_powerpc, swapEieio_16_powerpc, loadSwapEieio_32_powerpc, loadSwapSync_32_powerpc, storeSwapEieio_32_powerpc, storeSwapSync_32_powerpc, swapEieio_32_powerpc, eieio_powerpc, ioSync_powerpc – PowerPC specific i/o services
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwapEieio_16(uint16_f* addr); uint16_f loadSwapSync_16(uint16_f* addr); void storeSwapEieio_16(uint16_f* addr, uint16_f value); void storeSwapSync_16(uint16_f* addr, uint16_f value); void swapEieio_16(uint16_f* addr); uint32_f loadSwapEieio_32(uint32_f* addr); uint32_f loadSwapSync_32(uint32_f* addr); void storeSwapEieio_32(uint32_f* addr, uint32_f value); void storeSwapSync_32(uint32_f* addr, uint32_f value); void swapEieio_32 (uint32_f* addr); void ioSync(void); void eieio(void);</pre>
FEATURES	DKI
DESCRIPTION	Provides access to PowerPC DKI I/O routines, optimized to facilitate byte swapping and synchronization of input/output.
EXTENDED DESCRIPTION	<p>The PowerPC DKI provides specific I/O routines to handle byte swapping and synchronization of I/O. These services are based on PowerPC specific instructions to store/load with byte swapping, and to enforce in-order execution of I/O.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the PowerPC bus and the host bus, or to enforce sequential execution of I/O. PowerPC specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <pre>_16 for 16-bit data _32 for 32-bit data</pre>

`loadSwapEieio_xx` loads data from a given address and returns the corresponding byte swapped value. A PowerPC `eieio` instruction is issued after the data is read, to enforce execution of subsequent I/O sequentially with respect to the current load. The *addr* argument specifies the address to read from.

`loadSwapSync_xx` behaves like `loadSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`storeSwapEieio_xx` stores, at a given address, the *value* byte swapped. A PowerPC `eieio` instruction is issued after the data is written, to enforce execution of subsequent I/O sequentially with respect to the current store. The *addr* argument specifies the address to write to.

`storeSwapSync_xx` behaves like `storeSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`swapEieio_xx` swaps in place the bytes of the data stored at a given address. A PowerPC `eieio` instruction is issued between reading the data and re-writing the byte-swapped data, to enforce execution in order. The *addr* argument specifies the address of the data to byte-swap.

`ioSync` enforces context synchronization, completion of all instructions, and completion of all exceptions before returning.

`eieio` enforces execution in order of I/O and is mapped to the corresponding PowerPC assembler instruction. This routine is intended for use in managing shared data structures, in doing memory-mapped I/O, and in preventing load/store combining operations in main memory. Refer to PowerPC programming manuals for more details.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwapEieio_xx</code>	+	+	+	-
<code>loadSwapSync_xx</code>	+	+	+	-
<code>storeSwapEieio_xx</code>	+	+	+	-
<code>storeSwapSync_xx</code>	+	+	+	-
<code>swapEieio_xx</code>	+	+	+	-
<code>ioSync</code>	+	+	+	-
<code>eieio</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap_16(9DKI)` , `svDkiThreadCall(9DKI)`

NAME | load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services

SYNOPSIS

```
#include <dki/f_dki.h>
uint8_f load_sync_8(uint8_f * addr);

void store_sync_8(uint8_f * addr, uint8_f value);

uint16_f loadSwap_sync_16(uint16_f * addr);

void storeSwap_sync_16(uint16_f * addr, uint16_f value);

uint16_f load_sync_16(uint16_f * addr);

void store_sync_16(uint16_f * addr, uint16_f value);

uint32_f loadSwap_sync_32(uint32_f * addr);

void storeSwap_sync_32(uint32_f * addr, uint32_f value);

uint32_f load_sync_32(uint32_f * addr);

void store_sync_32(uint32_f * addr, uint32_f value);

uint64_f loadSwap_sync_64(uint64_f * addr);

void storeSwap_sync_64(uint64_f * addr, uint64_f value);

uint64_f load_sync_64(uint64_f * addr);

void store_sync_64(uint64_f * addr, uint64_f value);
```

FEATURES | DKI

DESCRIPTION

The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.

Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.

In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.

Specific I/O services are defined below as sets of routines where the `_xx` suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the *value* byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the *value*. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME	load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services
SYNOPSIS	<pre>#include <dki/f_dki.h> uint8_f load_sync_8(uint8_f * addr); void store_sync_8(uint8_f * addr, uint8_f value); uint16_f loadSwap_sync_16(uint16_f * addr); void storeSwap_sync_16(uint16_f * addr, uint16_f value); uint16_f load_sync_16(uint16_f * addr); void store_sync_16(uint16_f * addr, uint16_f value); uint32_f loadSwap_sync_32(uint32_f * addr); void storeSwap_sync_32(uint32_f * addr, uint32_f value); uint32_f load_sync_32(uint32_f * addr); void store_sync_32(uint32_f * addr, uint32_f value); uint64_f loadSwap_sync_64(uint64_f * addr); void storeSwap_sync_64(uint64_f * addr, uint64_f value); uint64_f load_sync_64(uint64_f * addr); void store_sync_64(uint64_f * addr, uint64_f value);</pre>
FEATURES	DKI
DESCRIPTION	<p>The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.</p> <p>In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.</p> <p>Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p>

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME | load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services

SYNOPSIS

```
#include <dki/f_dki.h>
uint8_f load_sync_8(uint8_f * addr);

void store_sync_8(uint8_f * addr, uint8_f value);

uint16_f loadSwap_sync_16(uint16_f * addr);

void storeSwap_sync_16(uint16_f * addr, uint16_f value);

uint16_f load_sync_16(uint16_f * addr);

void store_sync_16(uint16_f * addr, uint16_f value);

uint32_f loadSwap_sync_32(uint32_f * addr);

void storeSwap_sync_32(uint32_f * addr, uint32_f value);

uint32_f load_sync_32(uint32_f * addr);

void store_sync_32(uint32_f * addr, uint32_f value);

uint64_f loadSwap_sync_64(uint64_f * addr);

void storeSwap_sync_64(uint64_f * addr, uint64_f value);

uint64_f load_sync_64(uint64_f * addr);

void store_sync_64(uint64_f * addr, uint64_f value);
```

FEATURES

DKI

DESCRIPTION

The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.

Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.

In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.

Specific I/O services are defined below as sets of routines where the `_xx` suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME	load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services
SYNOPSIS	<pre>#include <dki/f_dki.h> uint8_f load_sync_8(uint8_f * addr); void store_sync_8(uint8_f * addr, uint8_f value); uint16_f loadSwap_sync_16(uint16_f * addr); void storeSwap_sync_16(uint16_f * addr, uint16_f value); uint16_f load_sync_16(uint16_f * addr); void store_sync_16(uint16_f * addr, uint16_f value); uint32_f loadSwap_sync_32(uint32_f * addr); void storeSwap_sync_32(uint32_f * addr, uint32_f value); uint32_f load_sync_32(uint32_f * addr); void store_sync_32(uint32_f * addr, uint32_f value); uint64_f loadSwap_sync_64(uint64_f * addr); void storeSwap_sync_64(uint64_f * addr, uint64_f value); uint64_f load_sync_64(uint64_f * addr); void store_sync_64(uint64_f * addr, uint64_f value);</pre>
FEATURES	DKI
DESCRIPTION	<p>The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.</p> <p>In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.</p> <p>Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p>

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME | load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services

SYNOPSIS

```
#include <dki/f_dki.h>
uint8_f load_sync_8(uint8_f * addr);

void store_sync_8(uint8_f * addr, uint8_f value);

uint16_f loadSwap_sync_16(uint16_f * addr);

void storeSwap_sync_16(uint16_f * addr, uint16_f value);

uint16_f load_sync_16(uint16_f * addr);

void store_sync_16(uint16_f * addr, uint16_f value);

uint32_f loadSwap_sync_32(uint32_f * addr);

void storeSwap_sync_32(uint32_f * addr, uint32_f value);

uint32_f load_sync_32(uint32_f * addr);

void store_sync_32(uint32_f * addr, uint32_f value);

uint64_f loadSwap_sync_64(uint64_f * addr);

void storeSwap_sync_64(uint64_f * addr, uint64_f value);

uint64_f load_sync_64(uint64_f * addr);

void store_sync_64(uint64_f * addr, uint64_f value);
```

FEATURES

DKI

DESCRIPTION

The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.

Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.

In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.

Specific I/O services are defined below as sets of routines where the `_xx` suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME	load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services
SYNOPSIS	<pre>#include <dki/f_dki.h> uint8_f load_sync_8(uint8_f * addr); void store_sync_8(uint8_f * addr, uint8_f value); uint16_f loadSwap_sync_16(uint16_f * addr); void storeSwap_sync_16(uint16_f * addr, uint16_f value); uint16_f load_sync_16(uint16_f * addr); void store_sync_16(uint16_f * addr, uint16_f value); uint32_f loadSwap_sync_32(uint32_f * addr); void storeSwap_sync_32(uint32_f * addr, uint32_f value); uint32_f load_sync_32(uint32_f * addr); void store_sync_32(uint32_f * addr, uint32_f value); uint64_f loadSwap_sync_64(uint64_f * addr); void storeSwap_sync_64(uint64_f * addr, uint64_f value); uint64_f load_sync_64(uint64_f * addr); void store_sync_64(uint64_f * addr, uint64_f value);</pre>
FEATURES	DKI
DESCRIPTION	<p>The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.</p> <p>In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.</p> <p>Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p>

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME | load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services

SYNOPSIS

```
#include <dki/f_dki.h>
uint8_f load_sync_8(uint8_f * addr);

void store_sync_8(uint8_f * addr, uint8_f value);

uint16_f loadSwap_sync_16(uint16_f * addr);

void storeSwap_sync_16(uint16_f * addr, uint16_f value);

uint16_f load_sync_16(uint16_f * addr);

void store_sync_16(uint16_f * addr, uint16_f value);

uint32_f loadSwap_sync_32(uint32_f * addr);

void storeSwap_sync_32(uint32_f * addr, uint32_f value);

uint32_f load_sync_32(uint32_f * addr);

void store_sync_32(uint32_f * addr, uint32_f value);

uint64_f loadSwap_sync_64(uint64_f * addr);

void storeSwap_sync_64(uint64_f * addr, uint64_f value);

uint64_f load_sync_64(uint64_f * addr);

void store_sync_64(uint64_f * addr, uint64_f value);
```

FEATURES | DKI

DESCRIPTION

The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.

Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.

In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.

Specific I/O services are defined below as sets of routines where the `_xx` suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc(9DKI) ,
svAsyncExcepDetach_usparc(9DKI)

NAME	loadSwap_16, storeSwap_16, swap_16, loadSwap_32, storeSwap_32, swap_32, loadSwap_64, storeSwap_64, swap_64 – specific i/o services										
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwap_16(uint16_f * addr); void storeSwap_16(uint16_f * addr, uint16_f value); void swap_16(uint16_f * addr); uint32_f loadSwap_32(uint32_f * addr); void storeSwap_32(uint32_f * addr, uint32_f value); void swap_32(uint32_f * addr); uint64_f loadSwap_64(uint64_f * addr); void storeSwap_64(uint64_f * addr, uint64_f value); void swap_64(uint64_f * addr);</pre>										
FEATURES	DKI										
DESCRIPTION	This function provides DKI routines to handle different byte ordering between the processor bus and the host bus.										
EXTENDED DESCRIPTION	<p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus. Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <p><code>_16</code> for 16-bit data</p> <p><code>_32</code> for 32-bit data</p> <p><code>_64</code> for 64-bit data</p> <p><code>loadSwap_xx</code> loads data from a given address and returns the corresponding byte swapped value. The <code>addr</code> argument specifies the address to read from.</p> <p><code>storeSwap_xx</code> stores into a given address the <code>value</code> byte swapped. The <code>addr</code> argument specifies the address to write to.</p> <p><code>swap_xx</code> swap in place the bytes of the data stored at a given address. The <code>addr</code> argument specifies the address of the data to byte-swap.</p>										
Allowed Calling Contexts	<p>The following table specifies the contexts in which a caller is allowed to invoke each service:</p> <table border="1"> <thead> <tr> <th>Services</th> <th>Base level</th> <th>DKI thread</th> <th>Interrupt</th> <th>Blocking</th> </tr> </thead> <tbody> <tr> <td>loadSwap_xx</td> <td style="text-align: center;">+</td> <td style="text-align: center;">+</td> <td style="text-align: center;">+</td> <td style="text-align: center;">-</td> </tr> </tbody> </table>	Services	Base level	DKI thread	Interrupt	Blocking	loadSwap_xx	+	+	+	-
Services	Base level	DKI thread	Interrupt	Blocking							
loadSwap_xx	+	+	+	-							

storeSwap_xx	+	+	+	-
swap_xx	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	loadSwap_16, storeSwap_16, swap_16, loadSwap_32, storeSwap_32, swap_32, loadSwap_64, storeSwap_64, swap_64 – specific i/o services										
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwap_16(uint16_f * addr); void storeSwap_16(uint16_f * addr, uint16_f value); void swap_16(uint16_f * addr); uint32_f loadSwap_32(uint32_f * addr); void storeSwap_32(uint32_f * addr, uint32_f value); void swap_32(uint32_f * addr); uint64_f loadSwap_64(uint64_f * addr); void storeSwap_64(uint64_f * addr, uint64_f value); void swap_64(uint64_f * addr);</pre>										
FEATURES	DKI										
DESCRIPTION	This function provides DKI routines to handle different byte ordering between the processor bus and the host bus.										
EXTENDED DESCRIPTION	<p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus. Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <p><code>_16</code> for 16-bit data</p> <p><code>_32</code> for 32-bit data</p> <p><code>_64</code> for 64-bit data</p> <p><code>loadSwap_xx</code> loads data from a given address and returns the corresponding byte swapped value. The <code>addr</code> argument specifies the address to read from.</p> <p><code>storeSwap_xx</code> stores into a given address the <code>value</code> byte swapped. The <code>addr</code> argument specifies the address to write to.</p> <p><code>swap_xx</code> swap in place the bytes of the data stored at a given address. The <code>addr</code> argument specifies the address of the data to byte-swap.</p>										
Allowed Calling Contexts	<p>The following table specifies the contexts in which a caller is allowed to invoke each service:</p> <table border="1"> <thead> <tr> <th>Services</th> <th>Base level</th> <th>DKI thread</th> <th>Interrupt</th> <th>Blocking</th> </tr> </thead> <tbody> <tr> <td>loadSwap_xx</td> <td style="text-align: center;">+</td> <td style="text-align: center;">+</td> <td style="text-align: center;">+</td> <td style="text-align: center;">-</td> </tr> </tbody> </table>	Services	Base level	DKI thread	Interrupt	Blocking	loadSwap_xx	+	+	+	-
Services	Base level	DKI thread	Interrupt	Blocking							
loadSwap_xx	+	+	+	-							

storeSwap_xx	+	+	+	-
swap_xx	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	loadSwap_16, storeSwap_16, swap_16, loadSwap_32, storeSwap_32, swap_32, loadSwap_64, storeSwap_64, swap_64 – specific i/o services										
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwap_16(uint16_f * addr); void storeSwap_16(uint16_f * addr, uint16_f value); void swap_16(uint16_f * addr); uint32_f loadSwap_32(uint32_f * addr); void storeSwap_32(uint32_f * addr, uint32_f value); void swap_32(uint32_f * addr); uint64_f loadSwap_64(uint64_f * addr); void storeSwap_64(uint64_f * addr, uint64_f value); void swap_64(uint64_f * addr);</pre>										
FEATURES	DKI										
DESCRIPTION	This function provides DKI routines to handle different byte ordering between the processor bus and the host bus.										
EXTENDED DESCRIPTION	<p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus. Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <p><code>_16</code> for 16-bit data</p> <p><code>_32</code> for 32-bit data</p> <p><code>_64</code> for 64-bit data</p> <p><code>loadSwap_xx</code> loads data from a given address and returns the corresponding byte swapped value. The <code>addr</code> argument specifies the address to read from.</p> <p><code>storeSwap_xx</code> stores into a given address the <code>value</code> byte swapped. The <code>addr</code> argument specifies the address to write to.</p> <p><code>swap_xx</code> swap in place the bytes of the data stored at a given address. The <code>addr</code> argument specifies the address of the data to byte-swap.</p>										
Allowed Calling Contexts	<p>The following table specifies the contexts in which a caller is allowed to invoke each service:</p> <table border="1"> <thead> <tr> <th>Services</th> <th>Base level</th> <th>DKI thread</th> <th>Interrupt</th> <th>Blocking</th> </tr> </thead> <tbody> <tr> <td>loadSwap_xx</td> <td style="text-align: center;">+</td> <td style="text-align: center;">+</td> <td style="text-align: center;">+</td> <td style="text-align: center;">-</td> </tr> </tbody> </table>	Services	Base level	DKI thread	Interrupt	Blocking	loadSwap_xx	+	+	+	-
Services	Base level	DKI thread	Interrupt	Blocking							
loadSwap_xx	+	+	+	-							

storeSwap_xx	+	+	+	-
swap_xx	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME | loadSwapEieio_16, storeSwapEieio_16, swapEieio_16, loadSwapEieio_32, storeSwapEieio_32, swapEieio_32, eieio – i/o services

FEATURES | DKI

DESCRIPTION | Provides access to DKI i/o routines, optimized to facilitate byte swapping and synchronization of input/output.

See the architecture specific man pages:

- loadSwapEieio_16_powerpc(9DKI)
- storeSwapEieio_16_powerpc(9DKI)
- swapEieio_16_powerpc(9DKI)
- loadSwapEieio_32_powerpc(9DKI)
- storeSwapEieio_32_powerpc(9DKI)
- swapEieio_32_powerpc(9DKI)
- eieio(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	loadSwapEieio_16_powerpc, loadSwapSync_16_powerpc, storeSwapEieio_16_powerpc, storeSwapSync_16_powerpc, swapEieio_16_powerpc, loadSwapEieio_32_powerpc, loadSwapSync_32_powerpc, storeSwapEieio_32_powerpc, storeSwapSync_32_powerpc, swapEieio_32_powerpc, eieio_powerpc, ioSync_powerpc – PowerPC specific i/o services
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwapEieio_16(uint16_f * addr); unit16_f loadSwapSync_16(uint16_f * addr); void storeSwapEieio_16(uint16_f * addr, uint16_f value); void storeSwapSync_16(uint16_f * addr, uint16_f value); void swapEieio_16(uint16_f * addr); uint32_f loadSwapEieio_32(uint32_f * addr); unit32_f loadSwapSync_32(uint32_f * addr); void storeSwapEieio_32(uint32_f * addr, uint32_f value); void storeSwapSync_32(uint32_f * addr, uint32_f value); void swapEieio_32 (uint32_f * addr); void ioSync(void); void eieio(void);</pre>
FEATURES	DKI
DESCRIPTION	Provides access to PowerPC DKI I/O routines, optimized to facilitate byte swapping and synchronization of input/output.
EXTENDED DESCRIPTION	<p>The PowerPC DKI provides specific I/O routines to handle byte swapping and synchronization of I/O. These services are based on PowerPC specific instructions to store/load with byte swapping, and to enforce in-order execution of I/O.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the PowerPC bus and the host bus, or to enforce sequential execution of I/O. PowerPC specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <pre>_16 for 16-bit data _32 for 32-bit data</pre>

`loadSwapEieio_xx` loads data from a given address and returns the corresponding byte swapped value. A PowerPC `eieio` instruction is issued after the data is read, to enforce execution of subsequent I/O sequentially with respect to the current load. The *addr* argument specifies the address to read from.

`loadSwapSync_xx` behaves like `loadSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`storeSwapEieio_xx` stores, at a given address, the *value* byte swapped. A PowerPC `eieio` instruction is issued after the data is written, to enforce execution of subsequent I/O sequentially with respect to the current store. The *addr* argument specifies the address to write to.

`storeSwapSync_xx` behaves like `storeSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`swapEieio_xx` swaps in place the bytes of the data stored at a given address. A PowerPC `eieio` instruction is issued between reading the data and re-writing the byte-swapped data, to enforce execution in order. The *addr* argument specifies the address of the data to byte-swap.

`ioSync` enforces context synchronization, completion of all instructions, and completion of all exceptions before returning.

`eieio` enforces execution in order of I/O and is mapped to the corresponding PowerPC assembler instruction. This routine is intended for use in managing shared data structures, in doing memory-mapped I/O, and in preventing load/store combining operations in main memory. Refer to PowerPC programming manuals for more details.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwapEieio_xx</code>	+	+	+	-
<code>loadSwapSync_xx</code>	+	+	+	-
<code>storeSwapEieio_xx</code>	+	+	+	-
<code>storeSwapSync_xx</code>	+	+	+	-
<code>swapEieio_xx</code>	+	+	+	-
<code>ioSync</code>	+	+	+	-
<code>eieio</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap_16(9DKI)` , `svDkiThreadCall(9DKI)`

NAME | loadSwapEieio_16, storeSwapEieio_16, swapEieio_16, loadSwapEieio_32, storeSwapEieio_32, swapEieio_32, eieio - i/o services

FEATURES | DKI

DESCRIPTION | Provides access to DKI i/o routines, optimized to facilitate byte swapping and synchronization of input/output.
 See the architecture specific man pages:

- loadSwapEieio_16_powerpc(9DKI)
- storeSwapEieio_16_powerpc(9DKI)
- swapEieio_16_powerpc(9DKI)
- loadSwapEieio_32_powerpc(9DKI)
- storeSwapEieio_32_powerpc(9DKI)
- swapEieio_32_powerpc(9DKI)
- eieio(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	loadSwapEieio_16_powerpc, loadSwapSync_16_powerpc, storeSwapEieio_16_powerpc, storeSwapSync_16_powerpc, swapEieio_16_powerpc, loadSwapEieio_32_powerpc, loadSwapSync_32_powerpc, storeSwapEieio_32_powerpc, storeSwapSync_32_powerpc, swapEieio_32_powerpc, eieio_powerpc, ioSync_powerpc – PowerPC specific i/o services
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwapEieio_16(uint16_f * addr); unit16_f loadSwapSync_16(uint16_f * addr); void storeSwapEieio_16(uint16_f * addr, uint16_f value); void storeSwapSync_16(uint16_f * addr, uint16_f value); void swapEieio_16(uint16_f * addr); uint32_f loadSwapEieio_32(uint32_f * addr); unit32_f loadSwapSync_32(uint32_f * addr); void storeSwapEieio_32(uint32_f * addr, uint32_f value); void storeSwapSync_32(uint32_f * addr, uint32_f value); void swapEieio_32 (uint32_f * addr); void ioSync(void); void eieio(void);</pre>
FEATURES	DKI
DESCRIPTION	Provides access to PowerPC DKI I/O routines, optimized to facilitate byte swapping and synchronization of input/output.
EXTENDED DESCRIPTION	<p>The PowerPC DKI provides specific I/O routines to handle byte swapping and synchronization of I/O. These services are based on PowerPC specific instructions to store/load with byte swapping, and to enforce in-order execution of I/O.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the PowerPC bus and the host bus, or to enforce sequential execution of I/O. PowerPC specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <pre>_16 for 16-bit data _32 for 32-bit data</pre>

`loadSwapEieio_xx` loads data from a given address and returns the corresponding byte swapped value. A PowerPC `eieio` instruction is issued after the data is read, to enforce execution of subsequent I/O sequentially with respect to the current load. The *addr* argument specifies the address to read from.

`loadSwapSync_xx` behaves like `loadSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`storeSwapEieio_xx` stores, at a given address, the *value* byte swapped. A PowerPC `eieio` instruction is issued after the data is written, to enforce execution of subsequent I/O sequentially with respect to the current store. The *addr* argument specifies the address to write to.

`storeSwapSync_xx` behaves like `storeSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`swapEieio_xx` swaps in place the bytes of the data stored at a given address. A PowerPC `eieio` instruction is issued between reading the data and re-writing the byte-swapped data, to enforce execution in order. The *addr* argument specifies the address of the data to byte-swap.

`ioSync` enforces context synchronization, completion of all instructions, and completion of all exceptions before returning.

`eieio` enforces execution in order of I/O and is mapped to the corresponding PowerPC assembler instruction. This routine is intended for use in managing shared data structures, in doing memory-mapped I/O, and in preventing load/store combining operations in main memory. Refer to PowerPC programming manuals for more details.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwapEieio_xx</code>	+	+	+	-
<code>loadSwapSync_xx</code>	+	+	+	-
<code>storeSwapEieio_xx</code>	+	+	+	-
<code>storeSwapSync_xx</code>	+	+	+	-
<code>swapEieio_xx</code>	+	+	+	-
<code>ioSync</code>	+	+	+	-
<code>eieio</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap_16(9DKI)` , `svDkiThreadCall(9DKI)`

NAME	load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services
SYNOPSIS	<pre>#include <dki/f_dki.h> uint8_f load_sync_8(uint8_f * addr); void store_sync_8(uint8_f * addr, uint8_f value); uint16_f loadSwap_sync_16(uint16_f * addr); void storeSwap_sync_16(uint16_f * addr, uint16_f value); uint16_f load_sync_16(uint16_f * addr); void store_sync_16(uint16_f * addr, uint16_f value); uint32_f loadSwap_sync_32(uint32_f * addr); void storeSwap_sync_32(uint32_f * addr, uint32_f value); uint32_f load_sync_32(uint32_f * addr); void store_sync_32(uint32_f * addr, uint32_f value); uint64_f loadSwap_sync_64(uint64_f * addr); void storeSwap_sync_64(uint64_f * addr, uint64_f value); uint64_f load_sync_64(uint64_f * addr); void store_sync_64(uint64_f * addr, uint64_f value);</pre>
FEATURES	DKI
DESCRIPTION	<p>The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.</p> <p>In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.</p> <p>Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p>

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME | load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services

SYNOPSIS

```
#include <dki/f_dki.h>
uint8_f load_sync_8(uint8_f * addr);

void store_sync_8(uint8_f * addr, uint8_f value);

uint16_f loadSwap_sync_16(uint16_f * addr);

void storeSwap_sync_16(uint16_f * addr, uint16_f value);

uint16_f load_sync_16(uint16_f * addr);

void store_sync_16(uint16_f * addr, uint16_f value);

uint32_f loadSwap_sync_32(uint32_f * addr);

void storeSwap_sync_32(uint32_f * addr, uint32_f value);

uint32_f load_sync_32(uint32_f * addr);

void store_sync_32(uint32_f * addr, uint32_f value);

uint64_f loadSwap_sync_64(uint64_f * addr);

void storeSwap_sync_64(uint64_f * addr, uint64_f value);

uint64_f load_sync_64(uint64_f * addr);

void store_sync_64(uint64_f * addr, uint64_f value);
```

FEATURES | DKI

DESCRIPTION

The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.

Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.

In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.

Specific I/O services are defined below as sets of routines where the `_xx` suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME	load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services
SYNOPSIS	<pre>#include <dki/f_dki.h> uint8_f load_sync_8(uint8_f * addr); void store_sync_8(uint8_f * addr, uint8_f value); uint16_f loadSwap_sync_16(uint16_f * addr); void storeSwap_sync_16(uint16_f * addr, uint16_f value); uint16_f load_sync_16(uint16_f * addr); void store_sync_16(uint16_f * addr, uint16_f value); uint32_f loadSwap_sync_32(uint32_f * addr); void storeSwap_sync_32(uint32_f * addr, uint32_f value); uint32_f load_sync_32(uint32_f * addr); void store_sync_32(uint32_f * addr, uint32_f value); uint64_f loadSwap_sync_64(uint64_f * addr); void storeSwap_sync_64(uint64_f * addr, uint64_f value); uint64_f load_sync_64(uint64_f * addr); void store_sync_64(uint64_f * addr, uint64_f value);</pre>
FEATURES	DKI
DESCRIPTION	<p>The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.</p> <p>In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.</p> <p>Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p>

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME | load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services

SYNOPSIS

```
#include <dki/f_dki.h>
uint8_f load_sync_8(uint8_f * addr);

void store_sync_8(uint8_f * addr, uint8_f value);

uint16_f loadSwap_sync_16(uint16_f * addr);

void storeSwap_sync_16(uint16_f * addr, uint16_f value);

uint16_f load_sync_16(uint16_f * addr);

void store_sync_16(uint16_f * addr, uint16_f value);

uint32_f loadSwap_sync_32(uint32_f * addr);

void storeSwap_sync_32(uint32_f * addr, uint32_f value);

uint32_f load_sync_32(uint32_f * addr);

void store_sync_32(uint32_f * addr, uint32_f value);

uint64_f loadSwap_sync_64(uint64_f * addr);

void storeSwap_sync_64(uint64_f * addr, uint64_f value);

uint64_f load_sync_64(uint64_f * addr);

void store_sync_64(uint64_f * addr, uint64_f value);
```

FEATURES | DKI

DESCRIPTION | The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.

Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.

In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.

Specific I/O services are defined below as sets of routines where the `_xx` suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME	load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services
SYNOPSIS	<pre>#include <dki/f_dki.h> uint8_f load_sync_8(uint8_f * addr); void store_sync_8(uint8_f * addr, uint8_f value); uint16_f loadSwap_sync_16(uint16_f * addr); void storeSwap_sync_16(uint16_f * addr, uint16_f value); uint16_f load_sync_16(uint16_f * addr); void store_sync_16(uint16_f * addr, uint16_f value); uint32_f loadSwap_sync_32(uint32_f * addr); void storeSwap_sync_32(uint32_f * addr, uint32_f value); uint32_f load_sync_32(uint32_f * addr); void store_sync_32(uint32_f * addr, uint32_f value); uint64_f loadSwap_sync_64(uint64_f * addr); void storeSwap_sync_64(uint64_f * addr, uint64_f value); uint64_f load_sync_64(uint64_f * addr); void store_sync_64(uint64_f * addr, uint64_f value);</pre>
FEATURES	DKI
DESCRIPTION	<p>The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.</p> <p>In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.</p> <p>Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p>

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME | load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services

SYNOPSIS

```
#include <dki/f_dki.h>
uint8_f load_sync_8(uint8_f * addr);

void store_sync_8(uint8_f * addr, uint8_f value);

uint16_f loadSwap_sync_16(uint16_f * addr);

void storeSwap_sync_16(uint16_f * addr, uint16_f value);

uint16_f load_sync_16(uint16_f * addr);

void store_sync_16(uint16_f * addr, uint16_f value);

uint32_f loadSwap_sync_32(uint32_f * addr);

void storeSwap_sync_32(uint32_f * addr, uint32_f value);

uint32_f load_sync_32(uint32_f * addr);

void store_sync_32(uint32_f * addr, uint32_f value);

uint64_f loadSwap_sync_64(uint64_f * addr);

void storeSwap_sync_64(uint64_f * addr, uint64_f value);

uint64_f load_sync_64(uint64_f * addr);

void store_sync_64(uint64_f * addr, uint64_f value);
```

FEATURES | DKI

DESCRIPTION

The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.

Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.

In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.

Specific I/O services are defined below as sets of routines where the `_xx` suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME	load_sync_8_usparc, store_sync_8_usparc, loadSwap_sync_16_usparc, storeSwap_sync_16_usparc, load_sync_16_usparc, store_sync_16_usparc, loadSwap_sync_32_usparc, storeSwap_sync_32_usparc, load_sync_32_usparc, store_sync_32_usparc, loadSwap_sync_64_usparc, storeSwap_sync_64_usparc, load_sync_64_usparc, store_sync_64_usparc – UltraSparc specific i/o services
SYNOPSIS	<pre>#include <dki/f_dki.h> uint8_f load_sync_8(uint8_f * addr); void store_sync_8(uint8_f * addr, uint8_f value); uint16_f loadSwap_sync_16(uint16_f * addr); void storeSwap_sync_16(uint16_f * addr, uint16_f value); uint16_f load_sync_16(uint16_f * addr); void store_sync_16(uint16_f * addr, uint16_f value); uint32_f loadSwap_sync_32(uint32_f * addr); void storeSwap_sync_32(uint32_f * addr, uint32_f value); uint32_f load_sync_32(uint32_f * addr); void store_sync_32(uint32_f * addr, uint32_f value); uint64_f loadSwap_sync_64(uint64_f * addr); void storeSwap_sync_64(uint64_f * addr, uint64_f value); uint64_f load_sync_64(uint64_f * addr); void store_sync_64(uint64_f * addr, uint64_f value);</pre>
FEATURES	DKI
DESCRIPTION	<p>The UltraSPARC DKI provides specific I/O routines optimized to handle byte swapping. These services are analogous to the generic ones except the membar #Sync instruction is issued each time the data is read/written from/to the memory. Note that this allows recovery from an asynchronous exception caused by these types of service routines.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus.</p> <p>In addition, the UltraSPARC DKI provides specific synchronized I/O routines. These routines do not perform the byte swapping but they are safe and allow recovery from an asynchronous exception caused by this type of service routine.</p> <p>Specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p>

- `_8` for 8-bit data
- `_16` for 16-bit data
- `_32` for 32-bit data
- `_64` for 64-bit data

Note that the `_8` suffix is only applied to the `load_sync_xx` and `store_sync_xx` routines.

loadSwap_sync_xx

`loadSwap_sync_xx` loads data from a given address and returns the corresponding byte swapped value. The `membar #Sync` instruction is issued once the data has been loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value` byte swapped. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

load_sync_xx

`load_sync_xx` loads data from a given address. The `membar #Sync` instruction is issued once the data is loaded.

The `addr` argument specifies the address from which to read.

storeSwap_sync_xx

`storeSwap_sync_xx` stores into a given address the `value`. The `membar #Sync` instruction is issued once the data is stored.

The `addr` argument specifies the address to which to write.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwap_sync_xx</code>	+	+	+	-
<code>storeSwap_sync_xx</code>	+	+	+	-
<code>load_sync_xx</code>	+	+	+	-
<code>store_sync_xx</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svAsyncExcepAttach_usparc>(9DKI) ,
svAsyncExcepDetach_usparc>(9DKI)

NAME svAsyncExcepAttach, svAsyncExcepDetach – asynchronous exceptions management

FEATURES DKI

DESCRIPTION Provides asynchronous exception management services.
 See the architecture specific man pages:

- svAsyncExcepAttach_usparc(9DKI)
- svAsyncExcepDetach_usparc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svAsyncExcepAttach_usparc, svAsyncExcepDetach_usparc – UltraSPARC asynchronous exceptions management								
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svAsyncExcepAttach(CpuExcepHandler <i>excepHandler</i>, void * <i>excepCookie</i>, CpuExcepId * <i>excepId</i>); void svAsyncExcepDetach(CpuExcepId * <i>excepId</i>);</pre>								
FEATURES	DKI								
DESCRIPTION svAsyncExcepAttach	<p>Provides UltraSPARC asynchronous exceptions management services. svAsyncExcepAttach attaches a given exception handler to the UltraSPARC asynchronous exceptions.</p> <p>The <i>excepHandler</i> argument specifies the handler to call back when an asynchronous exception occurs.</p> <p>The <i>excepCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On succes, K_OK is returned and an identifier for the attached exception handler is also returned in <i>excepId</i>. This identifier must be used in subsequent invocations of svAsyncExcepDetach.</p> <p>On failure, an error code is returned as follows:</p> <table border="0"> <tr> <td>K_ENOMEM</td> <td>The system is out of memory.</td> </tr> </table> <pre>typedef CpuExcepStatus (*CpuExcepHandler) (void* cookie, uint64_f afsr, uint64_f afar);</pre> <p>The asynchronous exception handler is invoked by the microkernel when one of the follwing bits is set in the asynchronous fault status register (AFSR):</p> <table border="0"> <tr> <td>AFSR_TO</td> <td>Time-out from system bus.</td> </tr> <tr> <td>AFSR_BERR</td> <td>Bus error from system bus.</td> </tr> <tr> <td>AFSR_CP</td> <td>DMA external cache parity error.</td> </tr> </table> <p>The asynchronous exception handler is called with masked CPU interrupts and disabled asynchronous exceptions. In addition, the CPU instruction and data caches are disabled. A cookie specified in svAsyncExcepAttach is passed back to the exception handler as the first argument. The asynchronous fault status and address registers are given to the handler.</p>	K_ENOMEM	The system is out of memory.	AFSR_TO	Time-out from system bus.	AFSR_BERR	Bus error from system bus.	AFSR_CP	DMA external cache parity error.
K_ENOMEM	The system is out of memory.								
AFSR_TO	Time-out from system bus.								
AFSR_BERR	Bus error from system bus.								
AFSR_CP	DMA external cache parity error.								

The exception handler should return either `CPU_EXCEP_CLAIMED` or `CPU_EXCEP_UNCLAIMED`. The `CPU_EXCEP_CLAIMED` value is returned by the handler in cases where the handler detects that the exception is due to an I/O or DMA access on the underlying bus, and the exception is recoverable. For instance, the host bus driver exception handler simply notifies an appropriate child driver (via the error handler invocation) if the exception is due to an I/O access initiated by this driver. In this case, the exception is recoverable and the host bus driver asks the microkernel to continue execution returning `CPU_EXCEP_CLAIMED`.

Note that multiple handlers may be attached to the asynchronous exceptions. When an exception occurs, the microkernel invokes all handlers sequentially iterating through the handlers list. Once `CPU_EXCEP_CLAIMED` is returned by a handler, the iteration is aborted and the microkernel tries to continue the execution. It re-enables the CPU instruction and data caches and asynchronous exceptions, sets `NPC` to `PC+4` and returns from the exception. Note that `NPC` is incorrect when an asynchronous exception occurs, and therefore, it should be corrected prior to continuing the execution. As a consequence, accesses to the bus I/O space should be synchronized by issuing a `membar #Sync` instruction after a `load/store` one. Also, the `membar` instruction should not be put in a delayed slot. Otherwise, asynchronous exceptions related to the I/O accesses are not granted to be recoverable. Note that the specific I/O service routines provided by the UltraSPARC DKI allow drivers to recover from an asynchronous exception caused by this type of service routine.

The microkernel considers the exception as unrecoverable and panics, if all exception handlers return `CPU_EXCEP_UNCLAIMED`.

svAsyncExcepDetach

`svAsyncExcepDetach` detaches the exception handler previously connected by `svAsyncExcepAttach`.

The *exceptId* argument identifies the attached exception handler, previously returned by `svAsyncExcepAttach`.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each exceptions management service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svAsyncExcepAttach</code>	+	+	-	+
<code>svAsyncExcepDetach</code>	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

load_sync_8_usparc(9DKI) , store_sync_8_usparc(9DKI) ,
loadSwap_sync_16_usparc(9DKI) , storeSwap_sync_16_usparc(9DKI)
 , load_sync_16_usparc(9DKI) , store_sync_16_usparc(9DKI) ,
loadSwap_sync_32_usparc(9DKI) , storeSwap_sync_32_usparc(9DKI)
 , load_sync_32_usparc(9DKI) , store_sync_32_usparc(9DKI) ,
loadSwap_sync_64_usparc(9DKI) , storeSwap_sync_64_usparc(9DKI) ,
load_sync_64_usparc(9DKI) , store_sync_64_usparc(9DKI) ,

NAME svAsyncExcepAttach, svAsyncExcepDetach – asynchronous exceptions management

FEATURES DKI

DESCRIPTION Provides asynchronous exception management services.
 See the architecture specific man pages:

- svAsyncExcepAttach_usparc(9DKI)
- svAsyncExcepDetach_usparc(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svAsyncExcepAttach_usparc, svAsyncExcepDetach_usparc – UltraSPARC asynchronous exceptions management								
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svAsyncExcepAttach(CpuExcepHandler <i>excepHandler</i>, void * <i>excepCookie</i>, CpuExcepId * <i>excepId</i>); void svAsyncExcepDetach(CpuExcepId * <i>excepId</i>);</pre>								
FEATURES	DKI								
DESCRIPTION svAsyncExcepAttach	<p>Provides UltraSPARC asynchronous exceptions management services. svAsyncExcepAttach attaches a given exception handler to the UltraSPARC asynchronous exceptions.</p> <p>The <i>excepHandler</i> argument specifies the handler to call back when an asynchronous exception occurs.</p> <p>The <i>excepCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On succes, K_OK is returned and an identifier for the attached exception handler is also returned in <i>excepId</i>. This identifier must be used in subsequent invocations of svAsyncExcepDetach.</p> <p>On failure, an error code is returned as follows:</p> <table border="0"> <tr> <td>K_ENOMEM</td> <td>The system is out of memory.</td> </tr> </table> <pre>typedef CpuExcepStatus (*CpuExcepHandler) (void* cookie, uint64_f afsr, uint64_f afar);</pre> <p>The asynchronous exception handler is invoked by the microkernel when one of the follwing bits is set in the asynchronous fault status register (AFSR):</p> <table border="0"> <tr> <td>AFSR_TO</td> <td>Time-out from system bus.</td> </tr> <tr> <td>AFSR_BERR</td> <td>Bus error from system bus.</td> </tr> <tr> <td>AFSR_CP</td> <td>DMA external cache parity error.</td> </tr> </table> <p>The asynchronous exception handler is called with masked CPU interrupts and disabled asynchronous exceptions. In addition, the CPU instruction and data caches are disabled. A cookie specified in svAsyncExcepAttach is passed back to the exception handler as the first argument. The asynchronous fault status and address registers are given to the handler.</p>	K_ENOMEM	The system is out of memory.	AFSR_TO	Time-out from system bus.	AFSR_BERR	Bus error from system bus.	AFSR_CP	DMA external cache parity error.
K_ENOMEM	The system is out of memory.								
AFSR_TO	Time-out from system bus.								
AFSR_BERR	Bus error from system bus.								
AFSR_CP	DMA external cache parity error.								

The exception handler should return either `CPU_EXCEP_CLAIMED` or `CPU_EXCEP_UNCLAIMED`. The `CPU_EXCEP_CLAIMED` value is returned by the handler in cases where the handler detects that the exception is due to an I/O or DMA access on the underlying bus, and the exception is recoverable. For instance, the host bus driver exception handler simply notifies an appropriate child driver (via the error handler invocation) if the exception is due to an I/O access initiated by this driver. In this case, the exception is recoverable and the host bus driver asks the microkernel to continue execution returning `CPU_EXCEP_CLAIMED`.

Note that multiple handlers may be attached to the asynchronous exceptions. When an exception occurs, the microkernel invokes all handlers sequentially iterating through the handlers list. Once `CPU_EXCEP_CLAIMED` is returned by a handler, the iteration is aborted and the microkernel tries to continue the execution. It re-enables the CPU instruction and data caches and asynchronous exceptions, sets `NPC` to `PC+4` and returns from the exception. Note that `NPC` is incorrect when an asynchronous exception occurs, and therefore, it should be corrected prior to continuing the execution. As a consequence, accesses to the bus I/O space should be synchronized by issuing a `membar #Sync` instruction after a `load/store` one. Also, the `membar` instruction should not be put in a delayed slot. Otherwise, asynchronous exceptions related to the I/O accesses are not granted to be recoverable. Note that the specific I/O service routines provided by the UltraSPARC DKI allow drivers to recover from an asynchronous exception caused by this type of service routine.

The microkernel considers the exception as unrecoverable and panics, if all exception handlers return `CPU_EXCEP_UNCLAIMED`.

svAsyncExcepDetach

`svAsyncExcepDetach` detaches the exception handler previously connected by `svAsyncExcepAttach`.

The *excepId* argument identifies the attached exception handler, previously returned by `svAsyncExcepAttach`.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each exceptions management service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svAsyncExcepAttach</code>	+	+	-	+
<code>svAsyncExcepDetach</code>	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

load_sync_8_usparc(9DKI) , store_sync_8_usparc(9DKI) ,
loadSwap_sync_16_usparc(9DKI) , storeSwap_sync_16_usparc(9DKI)
 , load_sync_16_usparc(9DKI) , store_sync_16_usparc(9DKI) ,
loadSwap_sync_32_usparc(9DKI) , storeSwap_sync_32_usparc(9DKI)
 , load_sync_32_usparc(9DKI) , store_sync_32_usparc(9DKI) ,
loadSwap_sync_64_usparc(9DKI) , storeSwap_sync_64_usparc(9DKI) ,
load_sync_64_usparc(9DKI) , store_sync_64_usparc(9DKI) ,

NAME	svDeviceRegister, svDeviceAlloc, svDeviceFree, svDeviceUnregister, svDeviceEvent, svDeviceLookup, svDeviceEntry, svDeviceRelease – device registry operations
SYNOPSIS	<pre>#include <dki/dki.h> void svDeviceRegister(DevRegId dev_id); DevRegId svDeviceAlloc(DevRegEntry * entry, unsigned int version, Bool shared, DevRelHandler handler); KnError svDeviceUnregister(DevRegId dev_id); void svDeviceEvent(DevRegId dev_id, DevEvent event, void* arg); void svDeviceFree(DevRegId dev_id); KnError svDeviceLookup(char * dev_class, unsigned int dev_version, unsigned int dev_unit, DevEventHandler cli_handler, void * cli_cookie, DevClientId * cli_id); DevRegEntry * svDeviceEntry(DevClientId cli_id); void svDeviceRelease(DevClientId cli_id); typedef struct { char* dev_class; void* dev_ops; void* dev_id; DevNode dev_node; } DevRegEntry;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to microkernel services implementing the device registry.
EXTENDED DESCRIPTION	<p>The device registry microkernel module implements a data base of driver instances servicing devices currently supported by the system. The device registry data base is populated by drivers which perform self-registration (using svDeviceRegister) at device initialization time. The device registry data base is accessed by driver clients in order to obtain a pointer to the driver instance servicing a given (logical) device. The device registry API is described below in detail.</p> <p>Note that only the svDeviceLookup , svDeviceRelease and svDeviceEntry microkernel calls should be used by driver clients. The rest of the API is dedicated to device drivers.</p>

`svDeviceAlloc` allocates a device registry entry for a given device driver instance. Note that the entry is allocated in an invalid state. This means that the entry is not included in the registry (it is not visible to clients via `svDeviceLookup`).

The entry becomes valid (visible to clients) when a subsequent `svDeviceRegister` is invoked. On the other hand, the driver is allowed to call `svDeviceEvent` on this type of entry. Shut down events signaled on an invalid entry are memorized by the device registry and they are processed when the entry becomes valid (that is, when `svDeviceRegister` is called).

The *entry* argument points to the `DevRegEntry` structure which designates the device driver instance. The `DevRegEntry` structure is described below:

<code>dev_class</code>	Points to a string specifying the device class name.
<code>dev_ops</code>	Points to a structure of driver service routines.
<code>dev_id</code>	Points to a handle which is usually passed back to the driver (as first argument) each time a driver service routine is invoked.
<code>dev_node</code>	Specifies the device node in the tree which is serviced by the driver. Note that a driver client must switch to the DKI thread context when accessing the device tree structure (see section "Device Node Properties").

The *version* argument specifies the driver interface version implemented by the device driver instance.

The *dev_class* field and *version* argument specify APIs implemented by the driver. This type of API is provided to driver clients as a structure of indirect functions implementing the API service routines. For instance, the "uart" device API is specified by the `UartDevOps` structure, the "timer" device API is specified by the `TimerDevOps` structure. A pointer to this type of structure is obtained by a device driver client from the *dev_ops* field of the device registry entry.

The *shared* argument specifies whether the device driver instance may be shared between multiple clients. In other words, it indicates whether the registered instance may be looked up multiple times.

The *handler* argument specifies the driver handler which is invoked by the device registry module as an acknowledgement to a shut-down event (see `svDeviceEvent`). *handler* is called when the last reference to the device registry

entry goes away and a shut-down event has been previously signalled on the entry. *entry* is passed back to the driver as the argument of *handler*.

In case of success, a non zero *DevRegId* is returned, otherwise `NULL` is returned.

Device registry assigns a logical unit number to the physical device. The logical unit number is unique within *dev_class*. The device registry handles a logical unit counter per class. All counters are initialized to zero. Each time a new device entry's is allocated, the current counter value is assigned to the entry logical unit number and the counter is incremented. Thus, the logical device order within a class corresponds to the allocation order.

The logical unit number is mainly used by device driver clients to iterate through a given device class looking for a certain device. Using a valid logical unit number, a client is able to access the corresponding device registry entry, in particular, the *dev_node* field which points to the device node. The physical device corresponding to a given logical unit number may then be detected by the device node properties and the node position within the device tree.

Note - The device driver may allocate (and then register) multiple entries in order to declare different classes (APIs) for the same physical device. For instance, a driver may implement orthogonal sets of operations like power management which are additional to the main device functionality. Another example is a device driver servicing a multi-function device. This type of driver would register multiple entries: one entry per device function. From the client's point of view, this type of device would be visible as multiple (mono-function) independent devices.

`svDeviceRegister` adds a given device registry entry to the registry. The entry must be previously allocated by `svDeviceAlloc`. The entry becomes valid (and therefore visible for clients) only if there is no shut-down event signalled on the entry. Otherwise, the entry remains invalid and the device registry module invokes a handler previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The *handler* routine is called in the DK1 thread context.

Note - When *handler* is called the entry is no longer registered. Thus, it is useless to invoke `svDeviceUnregister` once *handler* has been called. On the other hand, in order to free memory resources, the driver should release the device registry entry invoking `svDeviceFree` once the entry is no longer used by the driver. Typically, the entry is no longer used by the driver when the connection to the parent bus/nexus driver is closed, and the parent driver is unable to invoke a call-back handler signalling a bus/nexus event.

`svDeviceUnregister` removes the device entry from the device registry (if the entry is valid). The device entry being removed is specified by the `dev_id` argument. `svDeviceUnregister` returns the following results when a valid entry is specified:

<code>K_OK</code>	The device entry has been successfully removed from the registry.
<code>K_EBUSY</code>	The device entry has not been removed from the registry because it is locked by a driver client. In other words, the <code>svDeviceLookup/Release</code> pairs are not balanced.

`svDeviceUnregister` returns `K_EBUSY` when an invalid entry is specified. The driver must remove all related device entries from the registry prior to being unloaded from the system. Typically, `svDeviceUnregister` is issued by the device driver (for each device driver instance) when the driver code is requested to be unloaded via `drv_unload`. The driver unload fails when the driver code is still being used by the system.

`svDeviceEvent` notifies the device registry module that a shutdown event has occurred (for example, that a hot-pluggable device has been removed). The shutdown event basically means that the driver instance should no longer be used by driver clients and the device entry must be removed from the registry (that is, the driver instance will disappear). The shutdown event is immediately propagated to all driver clients in order to stop the device usage and to release the device entry as soon as possible. The device is specified by the `dev_id` argument.

`svDeviceEvent` propagates the shutdown event invoking the `event_handler` routines specified by the driver clients in `svDeviceLookup`. The `event` and `arg` arguments are opaque for the device registry. They are passed (as arguments) to the client event handlers. `event` specifies the reason for the shut-down as follows:

`DEV_EVENT_SHUTDOWN` Normal device shut down

`DEV_EVENT_REMOVAL` Hot-plug (surprise) device removal

`arg` is `event` specific.

All driver clients are requested to release the device entry (as soon as possible) invoking the `svDeviceRelease` routine. Note that, prior to the shutdown event propagation, `svDeviceEvent` removes the device entry from the registry in order to prevent the entry from being found (and locked) by new driver clients. Once all driver client handlers are invoked, `svDeviceEvent` returns to the driver. Note that the real device shutdown has to be deferred until the `handler` routine invocation.

Once the device entry is released by the last driver client, the device registry module invokes *handler* previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The handler routine is called in the DKI thread context.

Typically, `svDeviceEvent` is used by a device driver servicing a hot-pluggable device. `svDeviceEvent` is invoked by the device driver when the driver is notified (by its parent driver) that the device has been removed from the bus.

Note - In the case of hot-plug removal, the device driver must still be operational until the device entry is released (ignoring all requests to the driver, for example). The driver is allowed to call `svDeviceEvent` on an invalid (unregistered) entry. In this type of case, *handler* will be invoked only when a registration attempt is made, that is, `svDeviceRegister` will be invoked. In addition, the registration will fail and the entry will remain invalid.

`svDeviceFree` releases a given device registry entry previously allocated by `svDeviceAlloc`. The device entry being released is specified by the *dev_id* argument. The entry being released must be invalid (unregistered). Thus, if the entry was registered by `svDeviceRegister`, the driver is allowed to release it if either the entry is successfully unregistered by `svDeviceUnregister` or *handler* (previously specified by `svDeviceAlloc`) is invoked.

`svDeviceLookup` searches the device entry in the registry matching the specified device class and logical unit.

<code>dev_class</code>	Specifies the device class.
<code>dev_version</code>	Specifies the minimum device driver interface version required.
<code>dev_unit</code>	Specifies the logical device unit in the class.
<code>cli_handler</code>	Specifies the event handler which is called when a device event is signalled.
<code>cli_cookie</code>	Specifies the first argument of <i>cli_handler</i> .
<code>cli_id</code>	Is an output argument identifying the client token on the matching device entry. This <i>cli_id</i> is passed back as an argument to <code>svDeviceEntry</code> to get the associated device entry, or to <code>svDeviceRelease</code> to release the device driver instance.

`svDeviceLookup` returns the following results:

K_OK	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was found in the registry.
K_EBUSY	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was not found in the registry. Or, the device is being used by another client and device sharing is not allowed.
K_EUNKNOWN	There is no device entry in the registry matching <i>dev_class</i> .
K_UNDEF	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Or, there is a device entry in the <i>dev_class</i> with a logical unit greater than <i>dev_unit</i> . In other words, <i>dev_unit</i> falls into a hole in the <i>dev_class</i> unit ranges. The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair found in the registry is implementing an older version of the interface than the one required.
K_ETOOMUCH	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Moreover <i>dev_unit</i> is greater than all existing logical units in the <i>dev_class</i> .

In case of success, the corresponding device entry is locked in the registry until a subsequent `svDeviceRelease` . Note that the device registry lock may or may not be exclusive, depending on the value of the *shared* argument specified in `svDeviceAlloc` . In other words, the device may be shared between multiple driver clients, if *shared* was `TRUE` at registration time. The device driver instance must not disappear while its device entry is locked in the registry.

In the case of a catastrophic device shutdown (for example, hot-plug removal) the device entry lock owners are notified (through the *cli_handler* routine invocation) that the device entry must be released as soon as possible.

The `svDeviceEntry` routine returns the device entry associated with a given client identifier. The *cli_id* argument specifies the client identifier previously returned by `svDeviceLookup` .

The `svDeviceRelease` routine releases the lock on the given device entry. The device entry is specified by the *cli_id* argument. Obviously, a device driver client should no longer access the device driver instance and the device node once the device entry has been released.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
svDeviceAlloc	+	+	-	+
svDeviceFree	+	+	-	+
svDeviceRegister	+	+	-	+
svDeviceUnregister	+	+	-	+
svDeviceEvent	+	+	+	-
svDeviceLookup	+	+	-	+
svDeviceRelease	+	+	-	+
svDeviceEntry	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

dtreeNodeRoot(9DKI) , svDkiThreadCall(9DKI)

NAME	svDeviceRegister, svDeviceAlloc, svDeviceFree, svDeviceUnregister, svDeviceEvent, svDeviceLookup, svDeviceEntry, svDeviceRelease – device registry operations
SYNOPSIS	<pre>#include <dki/dki.h> void svDeviceRegister(DevRegId dev_id); DevRegId svDeviceAlloc(DevRegEntry * entry, unsigned int version, Bool shared, DevRelHandler handler); KnError svDeviceUnregister(DevRegId dev_id); void svDeviceEvent(DevRegId dev_id, DevEvent event, void* arg); void svDeviceFree(DevRegId dev_id); KnError svDeviceLookup(char * dev_class, unsigned int dev_version, unsigned int dev_unit, DevEventHandler cli_handler, void * cli_cookie, DevClientId * cli_id); DevRegEntry * svDeviceEntry(DevClientId cli_id); void svDeviceRelease(DevClientId cli_id); typedef struct { char* dev_class; void* dev_ops; void* dev_id; DevNode dev_node; } DevRegEntry;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to microkernel services implementing the device registry.
EXTENDED DESCRIPTION	<p>The device registry microkernel module implements a data base of driver instances servicing devices currently supported by the system. The device registry data base is populated by drivers which perform self-registration (using svDeviceRegister) at device initialization time. The device registry data base is accessed by driver clients in order to obtain a pointer to the driver instance servicing a given (logical) device. The device registry API is described below in detail.</p> <p>Note that only the svDeviceLookup, svDeviceRelease and svDeviceEntry microkernel calls should be used by driver clients. The rest of the API is dedicated to device drivers.</p>

`svDeviceAlloc` allocates a device registry entry for a given device driver instance. Note that the entry is allocated in an invalid state. This means that the entry is not included in the registry (it is not visible to clients via `svDeviceLookup`).

The entry becomes valid (visible to clients) when a subsequent `svDeviceRegister` is invoked. On the other hand, the driver is allowed to call `svDeviceEvent` on this type of entry. Shut down events signaled on an invalid entry are memorized by the device registry and they are processed when the entry becomes valid (that is, when `svDeviceRegister` is called).

The *entry* argument points to the `DevRegEntry` structure which designates the device driver instance. The `DevRegEntry` structure is described below:

<code>dev_class</code>	Points to a string specifying the device class name.
<code>dev_ops</code>	Points to a structure of driver service routines.
<code>dev_id</code>	Points to a handle which is usually passed back to the driver (as first argument) each time a driver service routine is invoked.
<code>dev_node</code>	Specifies the device node in the tree which is serviced by the driver. Note that a driver client must switch to the DKI thread context when accessing the device tree structure (see section "Device Node Properties").

The *version* argument specifies the driver interface version implemented by the device driver instance.

The *dev_class* field and *version* argument specify APIs implemented by the driver. This type of API is provided to driver clients as a structure of indirect functions implementing the API service routines. For instance, the "uart" device API is specified by the `UartDevOps` structure, the "timer" device API is specified by the `TimerDevOps` structure. A pointer to this type of structure is obtained by a device driver client from the *dev_ops* field of the device registry entry.

The *shared* argument specifies whether the device driver instance may be shared between multiple clients. In other words, it indicates whether the registered instance may be looked up multiple times.

The *handler* argument specifies the driver handler which is invoked by the device registry module as an acknowledgement to a shut-down event (see `svDeviceEvent`). *handler* is called when the last reference to the device registry

entry goes away and a shut-down event has been previously signalled on the entry. *entry* is passed back to the driver as the argument of *handler* .

In case of success, a non zero *DevRegId* is returned, otherwise `NULL` is returned.

Device registry assigns a logical unit number to the physical device. The logical unit number is unique within *dev_class* . The device registry handles a logical unit counter per class. All counters are initialized to zero. Each time a new device entry's is allocated, the current counter value is assigned to the entry logical unit number and the counter is incremented. Thus, the logical device order within a class corresponds to the allocation order.

The logical unit number is mainly used by device driver clients to iterate through a given device class looking for a certain device. Using a valid logical unit number, a client is able to access the corresponding device registry entry, in particular, the *dev_node* field which points to the device node. The physical device corresponding to a given logical unit number may then be detected by the device node properties and the node position within the device tree.

Note - The device driver may allocate (and then register) multiple entries in order to declare different classes (APIs) for the same physical device. For instance, a driver may implement orthogonal sets of operations like power management which are additional to the main device functionality. Another example is a device driver servicing a multi-function device. This type of driver would register multiple entries: one entry per device function. From the client's point of view, this type of device would be visible as multiple (mono-function) independent devices.

`svDeviceRegister` adds a given device registry entry to the registry. The entry must be previously allocated by `svDeviceAlloc` . The entry becomes valid (and therefore visible for clients) only if there is no shut-down event signalled on the entry. Otherwise, the entry remains invalid and the device registry module invokes a handler previously specified via `svDeviceAlloc` . A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler* . The *handler* routine is called in the DK1 thread context.

Note - When *handler* is called the entry is no longer registered. Thus, it is useless to invoke `svDeviceUnregister` once *handler* has been called. On the other hand, in order to free memory resources, the driver should release the device registry entry invoking `svDeviceFree` once the entry is no longer used by the driver. Typically, the entry is no longer used by the driver when the connection to the parent bus/nexus driver is closed, and the parent driver is unable to invoke a call-back handler signalling a bus/nexus event.

`svDeviceUnregister` removes the device entry from the device registry (if the entry is valid). The device entry being removed is specified by the `dev_id` argument. `svDeviceUnregister` returns the following results when a valid entry is specified:

<code>K_OK</code>	The device entry has been successfully removed from the registry.
<code>K_EBUSY</code>	The device entry has not been removed from the registry because it is locked by a driver client. In other words, the <code>svDeviceLookup/Release</code> pairs are not balanced.

`svDeviceUnregister` returns `K_EBUSY` when an invalid entry is specified. The driver must remove all related device entries from the registry prior to being unloaded from the system. Typically, `svDeviceUnregister` is issued by the device driver (for each device driver instance) when the driver code is requested to be unloaded via `drv_unload`. The driver unload fails when the driver code is still being used by the system.

`svDeviceEvent` notifies the device registry module that a shutdown event has occurred (for example, that a hot-pluggable device has been removed). The shutdown event basically means that the driver instance should no longer be used by driver clients and the device entry must be removed from the registry (that is, the driver instance will disappear). The shutdown event is immediately propagated to all driver clients in order to stop the device usage and to release the device entry as soon as possible. The device is specified by the `dev_id` argument.

`svDeviceEvent` propagates the shutdown event invoking the `event_handler` routines specified by the driver clients in `svDeviceLookup`. The `event` and `arg` arguments are opaque for the device registry. They are passed (as arguments) to the client event handlers. `event` specifies the reason for the shut-down as follows:

`DEV_EVENT_SHUTDOWN` Normal device shut down

`DEV_EVENT_REMOVAL` Hot-plug (surprise) device removal

`arg` is `event` specific.

All driver clients are requested to release the device entry (as soon as possible) invoking the `svDeviceRelease` routine. Note that, prior to the shutdown event propagation, `svDeviceEvent` removes the device entry from the registry in order to prevent the entry from being found (and locked) by new driver clients. Once all driver client handlers are invoked, `svDeviceEvent` returns to the driver. Note that the real device shutdown has to be deferred until the `handler` routine invocation.

Once the device entry is released by the last driver client, the device registry module invokes *handler* previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The handler routine is called in the DKI thread context.

Typically, `svDeviceEvent` is used by a device driver servicing a hot-pluggable device. `svDeviceEvent` is invoked by the device driver when the driver is notified (by its parent driver) that the device has been removed from the bus.

Note - In the case of hot-plug removal, the device driver must still be operational until the device entry is released (ignoring all requests to the driver, for example). The driver is allowed to call `svDeviceEvent` on an invalid (unregistered) entry. In this type of case, *handler* will be invoked only when a registration attempt is made, that is, `svDeviceRegister` will be invoked. In addition, the registration will fail and the entry will remain invalid.

`svDeviceFree` releases a given device registry entry previously allocated by `svDeviceAlloc`. The device entry being released is specified by the *dev_id* argument. The entry being released must be invalid (unregistered). Thus, if the entry was registered by `svDeviceRegister`, the driver is allowed to release it if either the entry is successfully unregistered by `svDeviceUnregister` or *handler* (previously specified by `svDeviceAlloc`) is invoked.

`svDeviceLookup` searches the device entry in the registry matching the specified device class and logical unit.

<code>dev_class</code>	Specifies the device class.
<code>dev_version</code>	Specifies the minimum device driver interface version required.
<code>dev_unit</code>	Specifies the logical device unit in the class.
<code>cli_handler</code>	Specifies the event handler which is called when a device event is signalled.
<code>cli_cookie</code>	Specifies the first argument of <i>cli_handler</i> .
<code>cli_id</code>	Is an output argument identifying the client token on the matching device entry. This <i>cli_id</i> is passed back as an argument to <code>svDeviceEntry</code> to get the associated device entry, or to <code>svDeviceRelease</code> to release the device driver instance.

`svDeviceLookup` returns the following results:

K_OK	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was found in the registry.
K_EBUSY	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was not found in the registry. Or, the device is being used by another client and device sharing is not allowed.
K_EUNKNOWN	There is no device entry in the registry matching <i>dev_class</i> .
K_UNDEF	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Or, there is a device entry in the <i>dev_class</i> with a logical unit greater than <i>dev_unit</i> . In other words, <i>dev_unit</i> falls into a hole in the <i>dev_class</i> unit ranges. The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair found in the registry is implementing an older version of the interface than the one required.
K_ETOOMUCH	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Moreover <i>dev_unit</i> is greater than all existing logical units in the <i>dev_class</i> .

In case of success, the corresponding device entry is locked in the registry until a subsequent `svDeviceRelease` . Note that the device registry lock may or may not be exclusive, depending on the value of the *shared* argument specified in `svDeviceAlloc` . In other words, the device may be shared between multiple driver clients, if *shared* was `TRUE` at registration time. The device driver instance must not disappear while its device entry is locked in the registry.

In the case of a catastrophic device shutdown (for example, hot-plug removal) the device entry lock owners are notified (through the *cli_handler* routine invocation) that the device entry must be released as soon as possible.

The `svDeviceEntry` routine returns the device entry associated with a given client identifier. The *cli_id* argument specifies the client identifier previously returned by `svDeviceLookup` .

The `svDeviceRelease` routine releases the lock on the given device entry. The device entry is specified by the *cli_id* argument. Obviously, a device driver client should no longer access the device driver instance and the device node once the device entry has been released.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
svDeviceAlloc	+	+	-	+
svDeviceFree	+	+	-	+
svDeviceRegister	+	+	-	+
svDeviceUnregister	+	+	-	+
svDeviceEvent	+	+	+	-
svDeviceLookup	+	+	-	+
svDeviceRelease	+	+	-	+
svDeviceEntry	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`dtreeNodeRoot(9DKI)` , `svDkiThreadCall(9DKI)`

NAME	svDeviceRegister, svDeviceAlloc, svDeviceFree, svDeviceUnregister, svDeviceEvent, svDeviceLookup, svDeviceEntry, svDeviceRelease – device registry operations
SYNOPSIS	<pre>#include <dki/dki.h> void svDeviceRegister(DevRegId dev_id); DevRegId svDeviceAlloc(DevRegEntry * entry, unsigned int version, Bool shared, DevRelHandler handler); KnError svDeviceUnregister(DevRegId dev_id); void svDeviceEvent(DevRegId dev_id, DevEvent event, void* arg); void svDeviceFree(DevRegId dev_id); KnError svDeviceLookup(char * dev_class, unsigned int dev_version, unsigned int dev_unit, DevEventHandler cli_handler, void * cli_cookie, DevClientId * cli_id); DevRegEntry * svDeviceEntry(DevClientId cli_id); void svDeviceRelease(DevClientId cli_id); typedef struct { char* dev_class; void* dev_ops; void* dev_id; DevNode dev_node; } DevRegEntry;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to microkernel services implementing the device registry.
EXTENDED DESCRIPTION	<p>The device registry microkernel module implements a data base of driver instances servicing devices currently supported by the system. The device registry data base is populated by drivers which perform self-registration (using svDeviceRegister) at device initialization time. The device registry data base is accessed by driver clients in order to obtain a pointer to the driver instance servicing a given (logical) device. The device registry API is described below in detail.</p> <p>Note that only the svDeviceLookup , svDeviceRelease and svDeviceEntry microkernel calls should be used by driver clients. The rest of the API is dedicated to device drivers.</p>

`svDeviceAlloc` allocates a device registry entry for a given device driver instance. Note that the entry is allocated in an invalid state. This means that the entry is not included in the registry (it is not visible to clients via `svDeviceLookup`).

The entry becomes valid (visible to clients) when a subsequent `svDeviceRegister` is invoked. On the other hand, the driver is allowed to call `svDeviceEvent` on this type of entry. Shut down events signaled on an invalid entry are memorized by the device registry and they are processed when the entry becomes valid (that is, when `svDeviceRegister` is called).

The *entry* argument points to the `DevRegEntry` structure which designates the device driver instance. The `DevRegEntry` structure is described below:

<code>dev_class</code>	Points to a string specifying the device class name.
<code>dev_ops</code>	Points to a structure of driver service routines.
<code>dev_id</code>	Points to a handle which is usually passed back to the driver (as first argument) each time a driver service routine is invoked.
<code>dev_node</code>	Specifies the device node in the tree which is serviced by the driver. Note that a driver client must switch to the DKI thread context when accessing the device tree structure (see section "Device Node Properties").

The *version* argument specifies the driver interface version implemented by the device driver instance.

The *dev_class* field and *version* argument specify APIs implemented by the driver. This type of API is provided to driver clients as a structure of indirect functions implementing the API service routines. For instance, the "uart" device API is specified by the `UartDevOps` structure, the "timer" device API is specified by the `TimerDevOps` structure. A pointer to this type of structure is obtained by a device driver client from the *dev_ops* field of the device registry entry.

The *shared* argument specifies whether the device driver instance may be shared between multiple clients. In other words, it indicates whether the registered instance may be looked up multiple times.

The *handler* argument specifies the driver handler which is invoked by the device registry module as an acknowledgement to a shut-down event (see `svDeviceEvent`). *handler* is called when the last reference to the device registry

entry goes away and a shut-down event has been previously signalled on the entry. *entry* is passed back to the driver as the argument of *handler*.

In case of success, a non zero *DevRegId* is returned, otherwise `NULL` is returned.

Device registry assigns a logical unit number to the physical device. The logical unit number is unique within *dev_class*. The device registry handles a logical unit counter per class. All counters are initialized to zero. Each time a new device entry's is allocated, the current counter value is assigned to the entry logical unit number and the counter is incremented. Thus, the logical device order within a class corresponds to the allocation order.

The logical unit number is mainly used by device driver clients to iterate through a given device class looking for a certain device. Using a valid logical unit number, a client is able to access the corresponding device registry entry, in particular, the *dev_node* field which points to the device node. The physical device corresponding to a given logical unit number may then be detected by the device node properties and the node position within the device tree.

Note - The device driver may allocate (and then register) multiple entries in order to declare different classes (APIs) for the same physical device. For instance, a driver may implement orthogonal sets of operations like power management which are additional to the main device functionality. Another example is a device driver servicing a multi-function device. This type of driver would register multiple entries: one entry per device function. From the client's point of view, this type of device would be visible as multiple (mono-function) independent devices.

`svDeviceRegister` adds a given device registry entry to the registry. The entry must be previously allocated by `svDeviceAlloc`. The entry becomes valid (and therefore visible for clients) only if there is no shut-down event signalled on the entry. Otherwise, the entry remains invalid and the device registry module invokes a handler previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The *handler* routine is called in the DKI thread context.

Note - When *handler* is called the entry is no longer registered. Thus, it is useless to invoke `svDeviceUnregister` once *handler* has been called. On the other hand, in order to free memory resources, the driver should release the device registry entry invoking `svDeviceFree` once the entry is no longer used by the driver. Typically, the entry is no longer used by the driver when the connection to the parent bus/nexus driver is closed, and the parent driver is unable to invoke a call-back handler signalling a bus/nexus event.

`svDeviceUnregister` removes the device entry from the device registry (if the entry is valid). The device entry being removed is specified by the `dev_id` argument. `svDeviceUnregister` returns the following results when a valid entry is specified:

<code>K_OK</code>	The device entry has been successfully removed from the registry.
<code>K_EBUSY</code>	The device entry has not been removed from the registry because it is locked by a driver client. In other words, the <code>svDeviceLookup/Release</code> pairs are not balanced.

`svDeviceUnregister` returns `K_EBUSY` when an invalid entry is specified. The driver must remove all related device entries from the registry prior to being unloaded from the system. Typically, `svDeviceUnregister` is issued by the device driver (for each device driver instance) when the driver code is requested to be unloaded via `drv_unload`. The driver unload fails when the driver code is still being used by the system.

`svDeviceEvent` notifies the device registry module that a shutdown event has occurred (for example, that a hot-pluggable device has been removed). The shutdown event basically means that the driver instance should no longer be used by driver clients and the device entry must be removed from the registry (that is, the driver instance will disappear). The shutdown event is immediately propagated to all driver clients in order to stop the device usage and to release the device entry as soon as possible. The device is specified by the `dev_id` argument.

`svDeviceEvent` propagates the shutdown event invoking the `event_handler` routines specified by the driver clients in `svDeviceLookup`. The `event` and `arg` arguments are opaque for the device registry. They are passed (as arguments) to the client event handlers. `event` specifies the reason for the shut-down as follows:

`DEV_EVENT_SHUTDOWN` Normal device shut down

`DEV_EVENT_REMOVAL` Hot-plug (surprise) device removal

`arg` is `event` specific.

All driver clients are requested to release the device entry (as soon as possible) invoking the `svDeviceRelease` routine. Note that, prior to the shutdown event propagation, `svDeviceEvent` removes the device entry from the registry in order to prevent the entry from being found (and locked) by new driver clients. Once all driver client handlers are invoked, `svDeviceEvent` returns to the driver. Note that the real device shutdown has to be deferred until the `handler` routine invocation.

Once the device entry is released by the last driver client, the device registry module invokes *handler* previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The handler routine is called in the DKI thread context.

Typically, `svDeviceEvent` is used by a device driver servicing a hot-pluggable device. `svDeviceEvent` is invoked by the device driver when the driver is notified (by its parent driver) that the device has been removed from the bus.

Note - In the case of hot-plug removal, the device driver must still be operational until the device entry is released (ignoring all requests to the driver, for example). The driver is allowed to call `svDeviceEvent` on an invalid (unregistered) entry. In this type of case, *handler* will be invoked only when a registration attempt is made, that is, `svDeviceRegister` will be invoked. In addition, the registration will fail and the entry will remain invalid.

`svDeviceFree` releases a given device registry entry previously allocated by `svDeviceAlloc`. The device entry being released is specified by the *dev_id* argument. The entry being released must be invalid (unregistered). Thus, if the entry was registered by `svDeviceRegister`, the driver is allowed to release it if either the entry is successfully unregistered by `svDeviceUnregister` or *handler* (previously specified by `svDeviceAlloc`) is invoked.

`svDeviceLookup` searches the device entry in the registry matching the specified device class and logical unit.

<code>dev_class</code>	Specifies the device class.
<code>dev_version</code>	Specifies the minimum device driver interface version required.
<code>dev_unit</code>	Specifies the logical device unit in the class.
<code>cli_handler</code>	Specifies the event handler which is called when a device event is signalled.
<code>cli_cookie</code>	Specifies the first argument of <i>cli_handler</i> .
<code>cli_id</code>	Is an output argument identifying the client token on the matching device entry. This <i>cli_id</i> is passed back as an argument to <code>svDeviceEntry</code> to get the associated device entry, or to <code>svDeviceRelease</code> to release the device driver instance.

`svDeviceLookup` returns the following results:

K_OK	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was found in the registry.
K_EBUSY	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was not found in the registry. Or, the device is being used by another client and device sharing is not allowed.
K_EUNKNOWN	There is no device entry in the registry matching <i>dev_class</i> .
K_UNDEF	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Or, there is a device entry in the <i>dev_class</i> with a logical unit greater than <i>dev_unit</i> . In other words, <i>dev_unit</i> falls into a hole in the <i>dev_class</i> unit ranges. The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair found in the registry is implementing an older version of the interface than the one required.
K_ETOOMUCH	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Moreover <i>dev_unit</i> is greater than all existing logical units in the <i>dev_class</i> .

In case of success, the corresponding device entry is locked in the registry until a subsequent `svDeviceRelease` . Note that the device registry lock may or may not be exclusive, depending on the value of the *shared* argument specified in `svDeviceAlloc` . In other words, the device may be shared between multiple driver clients, if *shared* was `TRUE` at registration time. The device driver instance must not disappear while its device entry is locked in the registry.

In the case of a catastrophic device shutdown (for example, hot-plug removal) the device entry lock owners are notified (through the *cli_handler* routine invocation) that the device entry must be released as soon as possible.

The `svDeviceEntry` routine returns the device entry associated with a given client identifier. The *cli_id* argument specifies the client identifier previously returned by `svDeviceLookup` .

The `svDeviceRelease` routine releases the lock on the given device entry. The device entry is specified by the *cli_id* argument. Obviously, a device driver client should no longer access the device driver instance and the device node once the device entry has been released.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
svDeviceAlloc	+	+	-	+
svDeviceFree	+	+	-	+
svDeviceRegister	+	+	-	+
svDeviceUnregister	+	+	-	+
svDeviceEvent	+	+	+	-
svDeviceLookup	+	+	-	+
svDeviceRelease	+	+	-	+
svDeviceEntry	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

dtreeNodeRoot(9DKI) , svDkiThreadCall(9DKI)

NAME	svDeviceRegister, svDeviceAlloc, svDeviceFree, svDeviceUnregister, svDeviceEvent, svDeviceLookup, svDeviceEntry, svDeviceRelease – device registry operations
SYNOPSIS	<pre>#include <dki/dki.h> void svDeviceRegister(DevRegId dev_id); DevRegId svDeviceAlloc(DevRegEntry * entry, unsigned int version, Bool shared, DevRelHandler handler); KnError svDeviceUnregister(DevRegId dev_id); void svDeviceEvent(DevRegId dev_id, DevEvent event, void* arg); void svDeviceFree(DevRegId dev_id); KnError svDeviceLookup(char * dev_class, unsigned int dev_version, unsigned int dev_unit, DevEventHandler cli_handler, void * cli_cookie, DevClientId * cli_id); DevRegEntry * svDeviceEntry(DevClientId cli_id); void svDeviceRelease(DevClientId cli_id); typedef struct { char* dev_class; void* dev_ops; void* dev_id; DevNode dev_node; } DevRegEntry;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to microkernel services implementing the device registry.
EXTENDED DESCRIPTION	<p>The device registry microkernel module implements a data base of driver instances servicing devices currently supported by the system. The device registry data base is populated by drivers which perform self-registration (using svDeviceRegister) at device initialization time. The device registry data base is accessed by driver clients in order to obtain a pointer to the driver instance servicing a given (logical) device. The device registry API is described below in detail.</p> <p>Note that only the svDeviceLookup, svDeviceRelease and svDeviceEntry microkernel calls should be used by driver clients. The rest of the API is dedicated to device drivers.</p>

`svDeviceAlloc` allocates a device registry entry for a given device driver instance. Note that the entry is allocated in an invalid state. This means that the entry is not included in the registry (it is not visible to clients via `svDeviceLookup`).

The entry becomes valid (visible to clients) when a subsequent `svDeviceRegister` is invoked. On the other hand, the driver is allowed to call `svDeviceEvent` on this type of entry. Shut down events signaled on an invalid entry are memorized by the device registry and they are processed when the entry becomes valid (that is, when `svDeviceRegister` is called).

The *entry* argument points to the `DevRegEntry` structure which designates the device driver instance. The `DevRegEntry` structure is described below:

<code>dev_class</code>	Points to a string specifying the device class name.
<code>dev_ops</code>	Points to a structure of driver service routines.
<code>dev_id</code>	Points to a handle which is usually passed back to the driver (as first argument) each time a driver service routine is invoked.
<code>dev_node</code>	Specifies the device node in the tree which is serviced by the driver. Note that a driver client must switch to the DKI thread context when accessing the device tree structure (see section "Device Node Properties").

The *version* argument specifies the driver interface version implemented by the device driver instance.

The *dev_class* field and *version* argument specify APIs implemented by the driver. This type of API is provided to driver clients as a structure of indirect functions implementing the API service routines. For instance, the "uart" device API is specified by the `UartDevOps` structure, the "timer" device API is specified by the `TimerDevOps` structure. A pointer to this type of structure is obtained by a device driver client from the *dev_ops* field of the device registry entry.

The *shared* argument specifies whether the device driver instance may be shared between multiple clients. In other words, it indicates whether the registered instance may be looked up multiple times.

The *handler* argument specifies the driver handler which is invoked by the device registry module as an acknowledgement to a shut-down event (see `svDeviceEvent`). *handler* is called when the last reference to the device registry

entry goes away and a shut-down event has been previously signalled on the entry. *entry* is passed back to the driver as the argument of *handler*.

In case of success, a non zero *DevRegId* is returned, otherwise `NULL` is returned.

Device registry assigns a logical unit number to the physical device. The logical unit number is unique within *dev_class*. The device registry handles a logical unit counter per class. All counters are initialized to zero. Each time a new device entry's is allocated, the current counter value is assigned to the entry logical unit number and the counter is incremented. Thus, the logical device order within a class corresponds to the allocation order.

The logical unit number is mainly used by device driver clients to iterate through a given device class looking for a certain device. Using a valid logical unit number, a client is able to access the corresponding device registry entry, in particular, the *dev_node* field which points to the device node. The physical device corresponding to a given logical unit number may then be detected by the device node properties and the node position within the device tree.

Note - The device driver may allocate (and then register) multiple entries in order to declare different classes (APIs) for the same physical device. For instance, a driver may implement orthogonal sets of operations like power management which are additional to the main device functionality. Another example is a device driver servicing a multi-function device. This type of driver would register multiple entries: one entry per device function. From the client's point of view, this type of device would be visible as multiple (mono-function) independent devices.

`svDeviceRegister` adds a given device registry entry to the registry. The entry must be previously allocated by `svDeviceAlloc`. The entry becomes valid (and therefore visible for clients) only if there is no shut-down event signalled on the entry. Otherwise, the entry remains invalid and the device registry module invokes a handler previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The *handler* routine is called in the DK1 thread context.

Note - When *handler* is called the entry is no longer registered. Thus, it is useless to invoke `svDeviceUnregister` once *handler* has been called. On the other hand, in order to free memory resources, the driver should release the device registry entry invoking `svDeviceFree` once the entry is no longer used by the driver. Typically, the entry is no longer used by the driver when the connection to the parent bus/nexus driver is closed, and the parent driver is unable to invoke a call-back handler signalling a bus/nexus event.

Once the device entry is released by the last driver client, the device registry module invokes *handler* previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The handler routine is called in the DKI thread context.

Typically, `svDeviceEvent` is used by a device driver servicing a hot-pluggable device. `svDeviceEvent` is invoked by the device driver when the driver is notified (by its parent driver) that the device has been removed from the bus.

Note - In the case of hot-plug removal, the device driver must still be operational until the device entry is released (ignoring all requests to the driver, for example). The driver is allowed to call `svDeviceEvent` on an invalid (unregistered) entry. In this type of case, *handler* will be invoked only when a registration attempt is made, that is, `svDeviceRegister` will be invoked. In addition, the registration will fail and the entry will remain invalid.

`svDeviceFree` releases a given device registry entry previously allocated by `svDeviceAlloc`. The device entry being released is specified by the *dev_id* argument. The entry being released must be invalid (unregistered). Thus, if the entry was registered by `svDeviceRegister`, the driver is allowed to release it if either the entry is successfully unregistered by `svDeviceUnregister` or *handler* (previously specified by `svDeviceAlloc`) is invoked.

`svDeviceLookup` searches the device entry in the registry matching the specified device class and logical unit.

<code>dev_class</code>	Specifies the device class.
<code>dev_version</code>	Specifies the minimum device driver interface version required.
<code>dev_unit</code>	Specifies the logical device unit in the class.
<code>cli_handler</code>	Specifies the event handler which is called when a device event is signalled.
<code>cli_cookie</code>	Specifies the first argument of <i>cli_handler</i> .
<code>cli_id</code>	Is an output argument identifying the client token on the matching device entry. This <i>cli_id</i> is passed back as an argument to <code>svDeviceEntry</code> to get the associated device entry, or to <code>svDeviceRelease</code> to release the device driver instance.

`svDeviceLookup` returns the following results:

K_OK	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was found in the registry.
K_EBUSY	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was not found in the registry. Or, the device is being used by another client and device sharing is not allowed.
K_EUNKNOWN	There is no device entry in the registry matching <i>dev_class</i> .
K_UNDEF	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Or, there is a device entry in the <i>dev_class</i> with a logical unit greater than <i>dev_unit</i> . In other words, <i>dev_unit</i> falls into a hole in the <i>dev_class</i> unit ranges. The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair found in the registry is implementing an older version of the interface than the one required.
K_ETOOMUCH	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Moreover <i>dev_unit</i> is greater than all existing logical units in the <i>dev_class</i> .

In case of success, the corresponding device entry is locked in the registry until a subsequent `svDeviceRelease` . Note that the device registry lock may or may not be exclusive, depending on the value of the *shared* argument specified in `svDeviceAlloc` . In other words, the device may be shared between multiple driver clients, if *shared* was `TRUE` at registration time. The device driver instance must not disappear while its device entry is locked in the registry.

In the case of a catastrophic device shutdown (for example, hot-plug removal) the device entry lock owners are notified (through the *cli_handler* routine invocation) that the device entry must be released as soon as possible.

The `svDeviceEntry` routine returns the device entry associated with a given client identifier. The *cli_id* argument specifies the client identifier previously returned by `svDeviceLookup` .

The `svDeviceRelease` routine releases the lock on the given device entry. The device entry is specified by the *cli_id* argument. Obviously, a device driver client should no longer access the device driver instance and the device node once the device entry has been released.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
svDeviceAlloc	+	+	-	+
svDeviceFree	+	+	-	+
svDeviceRegister	+	+	-	+
svDeviceUnregister	+	+	-	+
svDeviceEvent	+	+	+	-
svDeviceLookup	+	+	-	+
svDeviceRelease	+	+	-	+
svDeviceEntry	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`dtreeNodeRoot(9DKI)` , `svDkiThreadCall(9DKI)`

NAME	svDeviceRegister, svDeviceAlloc, svDeviceFree, svDeviceUnregister, svDeviceEvent, svDeviceLookup, svDeviceEntry, svDeviceRelease – device registry operations
SYNOPSIS	<pre>#include <dki/dki.h> void svDeviceRegister(DevRegId dev_id); DevRegId svDeviceAlloc(DevRegEntry * entry, unsigned int version, Bool shared, DevRelHandler handler); KnError svDeviceUnregister(DevRegId dev_id); void svDeviceEvent(DevRegId dev_id, DevEvent event, void* arg); void svDeviceFree(DevRegId dev_id); KnError svDeviceLookup(char * dev_class, unsigned int dev_version, unsigned int dev_unit, DevEventHandler cli_handler, void * cli_cookie, DevClientId * cli_id); DevRegEntry * svDeviceEntry(DevClientId cli_id); void svDeviceRelease(DevClientId cli_id); typedef struct { char* dev_class; void* dev_ops; void* dev_id; DevNode dev_node; } DevRegEntry;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to microkernel services implementing the device registry.
EXTENDED DESCRIPTION	<p>The device registry microkernel module implements a data base of driver instances servicing devices currently supported by the system. The device registry data base is populated by drivers which perform self-registration (using svDeviceRegister) at device initialization time. The device registry data base is accessed by driver clients in order to obtain a pointer to the driver instance servicing a given (logical) device. The device registry API is described below in detail.</p> <p>Note that only the svDeviceLookup , svDeviceRelease and svDeviceEntry microkernel calls should be used by driver clients. The rest of the API is dedicated to device drivers.</p>

`svDeviceAlloc` allocates a device registry entry for a given device driver instance. Note that the entry is allocated in an invalid state. This means that the entry is not included in the registry (it is not visible to clients via `svDeviceLookup`).

The entry becomes valid (visible to clients) when a subsequent `svDeviceRegister` is invoked. On the other hand, the driver is allowed to call `svDeviceEvent` on this type of entry. Shut down events signaled on an invalid entry are memorized by the device registry and they are processed when the entry becomes valid (that is, when `svDeviceRegister` is called).

The *entry* argument points to the `DevRegEntry` structure which designates the device driver instance. The `DevRegEntry` structure is described below:

<code>dev_class</code>	Points to a string specifying the device class name.
<code>dev_ops</code>	Points to a structure of driver service routines.
<code>dev_id</code>	Points to a handle which is usually passed back to the driver (as first argument) each time a driver service routine is invoked.
<code>dev_node</code>	Specifies the device node in the tree which is serviced by the driver. Note that a driver client must switch to the DKI thread context when accessing the device tree structure (see section "Device Node Properties").

The *version* argument specifies the driver interface version implemented by the device driver instance.

The *dev_class* field and *version* argument specify APIs implemented by the driver. This type of API is provided to driver clients as a structure of indirect functions implementing the API service routines. For instance, the "uart" device API is specified by the `UartDevOps` structure, the "timer" device API is specified by the `TimerDevOps` structure. A pointer to this type of structure is obtained by a device driver client from the *dev_ops* field of the device registry entry.

The *shared* argument specifies whether the device driver instance may be shared between multiple clients. In other words, it indicates whether the registered instance may be looked up multiple times.

The *handler* argument specifies the driver handler which is invoked by the device registry module as an acknowledgement to a shut-down event (see `svDeviceEvent`). *handler* is called when the last reference to the device registry

entry goes away and a shut-down event has been previously signalled on the entry. *entry* is passed back to the driver as the argument of *handler*.

In case of success, a non zero *DevRegId* is returned, otherwise `NULL` is returned.

Device registry assigns a logical unit number to the physical device. The logical unit number is unique within *dev_class*. The device registry handles a logical unit counter per class. All counters are initialized to zero. Each time a new device entry's is allocated, the current counter value is assigned to the entry logical unit number and the counter is incremented. Thus, the logical device order within a class corresponds to the allocation order.

The logical unit number is mainly used by device driver clients to iterate through a given device class looking for a certain device. Using a valid logical unit number, a client is able to access the corresponding device registry entry, in particular, the *dev_node* field which points to the device node. The physical device corresponding to a given logical unit number may then be detected by the device node properties and the node position within the device tree.

Note - The device driver may allocate (and then register) multiple entries in order to declare different classes (APIs) for the same physical device. For instance, a driver may implement orthogonal sets of operations like power management which are additional to the main device functionality. Another example is a device driver servicing a multi-function device. This type of driver would register multiple entries: one entry per device function. From the client's point of view, this type of device would be visible as multiple (mono-function) independent devices.

`svDeviceRegister` adds a given device registry entry to the registry. The entry must be previously allocated by `svDeviceAlloc`. The entry becomes valid (and therefore visible for clients) only if there is no shut-down event signalled on the entry. Otherwise, the entry remains invalid and the device registry module invokes a handler previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The *handler* routine is called in the DKI thread context.

Note - When *handler* is called the entry is no longer registered. Thus, it is useless to invoke `svDeviceUnregister` once *handler* has been called. On the other hand, in order to free memory resources, the driver should release the device registry entry invoking `svDeviceFree` once the entry is no longer used by the driver. Typically, the entry is no longer used by the driver when the connection to the parent bus/nexus driver is closed, and the parent driver is unable to invoke a call-back handler signalling a bus/nexus event.

`svDeviceUnregister` removes the device entry from the device registry (if the entry is valid). The device entry being removed is specified by the `dev_id` argument. `svDeviceUnregister` returns the following results when a valid entry is specified:

<code>K_OK</code>	The device entry has been successfully removed from the registry.
<code>K_EBUSY</code>	The device entry has not been removed from the registry because it is locked by a driver client. In other words, the <code>svDeviceLookup/Release</code> pairs are not balanced.

`svDeviceUnregister` returns `K_EBUSY` when an invalid entry is specified. The driver must remove all related device entries from the registry prior to being unloaded from the system. Typically, `svDeviceUnregister` is issued by the device driver (for each device driver instance) when the driver code is requested to be unloaded via `drv_unload`. The driver unload fails when the driver code is still being used by the system.

`svDeviceEvent` notifies the device registry module that a shutdown event has occurred (for example, that a hot-pluggable device has been removed). The shutdown event basically means that the driver instance should no longer be used by driver clients and the device entry must be removed from the registry (that is, the driver instance will disappear). The shutdown event is immediately propagated to all driver clients in order to stop the device usage and to release the device entry as soon as possible. The device is specified by the `dev_id` argument.

`svDeviceEvent` propagates the shutdown event invoking the `event_handler` routines specified by the driver clients in `svDeviceLookup`. The `event` and `arg` arguments are opaque for the device registry. They are passed (as arguments) to the client event handlers. `event` specifies the reason for the shut-down as follows:

`DEV_EVENT_SHUTDOWN` Normal device shut down

`DEV_EVENT_REMOVAL` Hot-plug (surprise) device removal

`arg` is `event` specific.

All driver clients are requested to release the device entry (as soon as possible) invoking the `svDeviceRelease` routine. Note that, prior to the shutdown event propagation, `svDeviceEvent` removes the device entry from the registry in order to prevent the entry from being found (and locked) by new driver clients. Once all driver client handlers are invoked, `svDeviceEvent` returns to the driver. Note that the real device shutdown has to be deferred until the `handler` routine invocation.

Once the device entry is released by the last driver client, the device registry module invokes *handler* previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The handler routine is called in the DKI thread context.

Typically, `svDeviceEvent` is used by a device driver servicing a hot-pluggable device. `svDeviceEvent` is invoked by the device driver when the driver is notified (by its parent driver) that the device has been removed from the bus.

Note - In the case of hot-plug removal, the device driver must still be operational until the device entry is released (ignoring all requests to the driver, for example). The driver is allowed to call `svDeviceEvent` on an invalid (unregistered) entry. In this type of case, *handler* will be invoked only when a registration attempt is made, that is, `svDeviceRegister` will be invoked. In addition, the registration will fail and the entry will remain invalid.

`svDeviceFree` releases a given device registry entry previously allocated by `svDeviceAlloc`. The device entry being released is specified by the *dev_id* argument. The entry being released must be invalid (unregistered). Thus, if the entry was registered by `svDeviceRegister`, the driver is allowed to release it if either the entry is successfully unregistered by `svDeviceUnregister` or *handler* (previously specified by `svDeviceAlloc`) is invoked.

`svDeviceLookup` searches the device entry in the registry matching the specified device class and logical unit.

<code>dev_class</code>	Specifies the device class.
<code>dev_version</code>	Specifies the minimum device driver interface version required.
<code>dev_unit</code>	Specifies the logical device unit in the class.
<code>cli_handler</code>	Specifies the event handler which is called when a device event is signalled.
<code>cli_cookie</code>	Specifies the first argument of <i>cli_handler</i> .
<code>cli_id</code>	Is an output argument identifying the client token on the matching device entry. This <i>cli_id</i> is passed back as an argument to <code>svDeviceEntry</code> to get the associated device entry, or to <code>svDeviceRelease</code> to release the device driver instance.

`svDeviceLookup` returns the following results:

K_OK	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was found in the registry.
K_EBUSY	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was not found in the registry. Or, the device is being used by another client and device sharing is not allowed.
K_EUNKNOWN	There is no device entry in the registry matching <i>dev_class</i> .
K_UNDEF	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Or, there is a device entry in the <i>dev_class</i> with a logical unit greater than <i>dev_unit</i> . In other words, <i>dev_unit</i> falls into a hole in the <i>dev_class</i> unit ranges. The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair found in the registry is implementing an older version of the interface than the one required.
K_ETOOMUCH	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Moreover <i>dev_unit</i> is greater than all existing logical units in the <i>dev_class</i> .

In case of success, the corresponding device entry is locked in the registry until a subsequent `svDeviceRelease` . Note that the device registry lock may or may not be exclusive, depending on the value of the *shared* argument specified in `svDeviceAlloc` . In other words, the device may be shared between multiple driver clients, if *shared* was `TRUE` at registration time. The device driver instance must not disappear while its device entry is locked in the registry.

In the case of a catastrophic device shutdown (for example, hot-plug removal) the device entry lock owners are notified (through the *cli_handler* routine invocation) that the device entry must be released as soon as possible.

The `svDeviceEntry` routine returns the device entry associated with a given client identifier. The *cli_id* argument specifies the client identifier previously returned by `svDeviceLookup` .

The `svDeviceRelease` routine releases the lock on the given device entry. The device entry is specified by the *cli_id* argument. Obviously, a device driver client should no longer access the device driver instance and the device node once the device entry has been released.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
svDeviceAlloc	+	+	-	+
svDeviceFree	+	+	-	+
svDeviceRegister	+	+	-	+
svDeviceUnregister	+	+	-	+
svDeviceEvent	+	+	+	-
svDeviceLookup	+	+	-	+
svDeviceRelease	+	+	-	+
svDeviceEntry	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`dtreeNodeRoot(9DKI)` , `svDkiThreadCall(9DKI)`

NAME	svDeviceRegister, svDeviceAlloc, svDeviceFree, svDeviceUnregister, svDeviceEvent, svDeviceLookup, svDeviceEntry, svDeviceRelease – device registry operations
SYNOPSIS	<pre>#include <dki/dki.h> void svDeviceRegister(DevRegId dev_id); DevRegId svDeviceAlloc(DevRegEntry * entry, unsigned int version, Bool shared, DevRelHandler handler); KnError svDeviceUnregister(DevRegId dev_id); void svDeviceEvent(DevRegId dev_id, DevEvent event, void* arg); void svDeviceFree(DevRegId dev_id); KnError svDeviceLookup(char * dev_class, unsigned int dev_version, unsigned int dev_unit, DevEventHandler cli_handler, void * cli_cookie, DevClientId * cli_id); DevRegEntry * svDeviceEntry(DevClientId cli_id); void svDeviceRelease(DevClientId cli_id); typedef struct { char* dev_class; void* dev_ops; void* dev_id; DevNode dev_node; } DevRegEntry;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to microkernel services implementing the device registry.
EXTENDED DESCRIPTION	<p>The device registry microkernel module implements a data base of driver instances servicing devices currently supported by the system. The device registry data base is populated by drivers which perform self-registration (using svDeviceRegister) at device initialization time. The device registry data base is accessed by driver clients in order to obtain a pointer to the driver instance servicing a given (logical) device. The device registry API is described below in detail.</p> <p>Note that only the svDeviceLookup, svDeviceRelease and svDeviceEntry microkernel calls should be used by driver clients. The rest of the API is dedicated to device drivers.</p>

`svDeviceAlloc` allocates a device registry entry for a given device driver instance. Note that the entry is allocated in an invalid state. This means that the entry is not included in the registry (it is not visible to clients via `svDeviceLookup`).

The entry becomes valid (visible to clients) when a subsequent `svDeviceRegister` is invoked. On the other hand, the driver is allowed to call `svDeviceEvent` on this type of entry. Shut down events signaled on an invalid entry are memorized by the device registry and they are processed when the entry becomes valid (that is, when `svDeviceRegister` is called).

The *entry* argument points to the `DevRegEntry` structure which designates the device driver instance. The `DevRegEntry` structure is described below:

<code>dev_class</code>	Points to a string specifying the device class name.
<code>dev_ops</code>	Points to a structure of driver service routines.
<code>dev_id</code>	Points to a handle which is usually passed back to the driver (as first argument) each time a driver service routine is invoked.
<code>dev_node</code>	Specifies the device node in the tree which is serviced by the driver. Note that a driver client must switch to the DKI thread context when accessing the device tree structure (see section "Device Node Properties").

The *version* argument specifies the driver interface version implemented by the device driver instance.

The *dev_class* field and *version* argument specify APIs implemented by the driver. This type of API is provided to driver clients as a structure of indirect functions implementing the API service routines. For instance, the "uart" device API is specified by the `UartDevOps` structure, the "timer" device API is specified by the `TimerDevOps` structure. A pointer to this type of structure is obtained by a device driver client from the *dev_ops* field of the device registry entry.

The *shared* argument specifies whether the device driver instance may be shared between multiple clients. In other words, it indicates whether the registered instance may be looked up multiple times.

The *handler* argument specifies the driver handler which is invoked by the device registry module as an acknowledgement to a shut-down event (see `svDeviceEvent`). *handler* is called when the last reference to the device registry

entry goes away and a shut-down event has been previously signalled on the entry. *entry* is passed back to the driver as the argument of *handler*.

In case of success, a non zero *DevRegId* is returned, otherwise `NULL` is returned.

Device registry assigns a logical unit number to the physical device. The logical unit number is unique within *dev_class*. The device registry handles a logical unit counter per class. All counters are initialized to zero. Each time a new device entry's is allocated, the current counter value is assigned to the entry logical unit number and the counter is incremented. Thus, the logical device order within a class corresponds to the allocation order.

The logical unit number is mainly used by device driver clients to iterate through a given device class looking for a certain device. Using a valid logical unit number, a client is able to access the corresponding device registry entry, in particular, the *dev_node* field which points to the device node. The physical device corresponding to a given logical unit number may then be detected by the device node properties and the node position within the device tree.

Note - The device driver may allocate (and then register) multiple entries in order to declare different classes (APIs) for the same physical device. For instance, a driver may implement orthogonal sets of operations like power management which are additional to the main device functionality. Another example is a device driver servicing a multi-function device. This type of driver would register multiple entries: one entry per device function. From the client's point of view, this type of device would be visible as multiple (mono-function) independent devices.

`svDeviceRegister` adds a given device registry entry to the registry. The entry must be previously allocated by `svDeviceAlloc`. The entry becomes valid (and therefore visible for clients) only if there is no shut-down event signalled on the entry. Otherwise, the entry remains invalid and the device registry module invokes a handler previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The *handler* routine is called in the DK1 thread context.

Note - When *handler* is called the entry is no longer registered. Thus, it is useless to invoke `svDeviceUnregister` once *handler* has been called. On the other hand, in order to free memory resources, the driver should release the device registry entry invoking `svDeviceFree` once the entry is no longer used by the driver. Typically, the entry is no longer used by the driver when the connection to the parent bus/nexus driver is closed, and the parent driver is unable to invoke a call-back handler signalling a bus/nexus event.

`svDeviceUnregister` removes the device entry from the device registry (if the entry is valid). The device entry being removed is specified by the `dev_id` argument. `svDeviceUnregister` returns the following results when a valid entry is specified:

<code>K_OK</code>	The device entry has been successfully removed from the registry.
<code>K_EBUSY</code>	The device entry has not been removed from the registry because it is locked by a driver client. In other words, the <code>svDeviceLookup/Release</code> pairs are not balanced.

`svDeviceUnregister` returns `K_EBUSY` when an invalid entry is specified. The driver must remove all related device entries from the registry prior to being unloaded from the system. Typically, `svDeviceUnregister` is issued by the device driver (for each device driver instance) when the driver code is requested to be unloaded via `drv_unload`. The driver unload fails when the driver code is still being used by the system.

`svDeviceEvent` notifies the device registry module that a shutdown event has occurred (for example, that a hot-pluggable device has been removed). The shutdown event basically means that the driver instance should no longer be used by driver clients and the device entry must be removed from the registry (that is, the driver instance will disappear). The shutdown event is immediately propagated to all driver clients in order to stop the device usage and to release the device entry as soon as possible. The device is specified by the `dev_id` argument.

`svDeviceEvent` propagates the shutdown event invoking the `event_handler` routines specified by the driver clients in `svDeviceLookup`. The `event` and `arg` arguments are opaque for the device registry. They are passed (as arguments) to the client event handlers. `event` specifies the reason for the shut-down as follows:

`DEV_EVENT_SHUTDOWN` Normal device shut down

`DEV_EVENT_REMOVAL` Hot-plug (surprise) device removal

`arg` is `event` specific.

All driver clients are requested to release the device entry (as soon as possible) invoking the `svDeviceRelease` routine. Note that, prior to the shutdown event propagation, `svDeviceEvent` removes the device entry from the registry in order to prevent the entry from being found (and locked) by new driver clients. Once all driver client handlers are invoked, `svDeviceEvent` returns to the driver. Note that the real device shutdown has to be deferred until the `handler` routine invocation.

Once the device entry is released by the last driver client, the device registry module invokes *handler* previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The handler routine is called in the DKI thread context.

Typically, `svDeviceEvent` is used by a device driver servicing a hot-pluggable device. `svDeviceEvent` is invoked by the device driver when the driver is notified (by its parent driver) that the device has been removed from the bus.

Note - In the case of hot-plug removal, the device driver must still be operational until the device entry is released (ignoring all requests to the driver, for example). The driver is allowed to call `svDeviceEvent` on an invalid (unregistered) entry. In this type of case, *handler* will be invoked only when a registration attempt is made, that is, `svDeviceRegister` will be invoked. In addition, the registration will fail and the entry will remain invalid.

`svDeviceFree` releases a given device registry entry previously allocated by `svDeviceAlloc`. The device entry being released is specified by the *dev_id* argument. The entry being released must be invalid (unregistered). Thus, if the entry was registered by `svDeviceRegister`, the driver is allowed to release it if either the entry is successfully unregistered by `svDeviceUnregister` or *handler* (previously specified by `svDeviceAlloc`) is invoked.

`svDeviceLookup` searches the device entry in the registry matching the specified device class and logical unit.

<code>dev_class</code>	Specifies the device class.
<code>dev_version</code>	Specifies the minimum device driver interface version required.
<code>dev_unit</code>	Specifies the logical device unit in the class.
<code>cli_handler</code>	Specifies the event handler which is called when a device event is signalled.
<code>cli_cookie</code>	Specifies the first argument of <i>cli_handler</i> .
<code>cli_id</code>	Is an output argument identifying the client token on the matching device entry. This <i>cli_id</i> is passed back as an argument to <code>svDeviceEntry</code> to get the associated device entry, or to <code>svDeviceRelease</code> to release the device driver instance.

`svDeviceLookup` returns the following results:

K_OK	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was found in the registry.
K_EBUSY	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was not found in the registry. Or, the device is being used by another client and device sharing is not allowed.
K_EUNKNOWN	There is no device entry in the registry matching <i>dev_class</i> .
K_UNDEF	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Or, there is a device entry in the <i>dev_class</i> with a logical unit greater than <i>dev_unit</i> . In other words, <i>dev_unit</i> falls into a hole in the <i>dev_class</i> unit ranges. The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair found in the registry is implementing an older version of the interface than the one required.
K_ETOOMUCH	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Moreover <i>dev_unit</i> is greater than all existing logical units in the <i>dev_class</i> .

In case of success, the corresponding device entry is locked in the registry until a subsequent `svDeviceRelease` . Note that the device registry lock may or may not be exclusive, depending on the value of the *shared* argument specified in `svDeviceAlloc` . In other words, the device may be shared between multiple driver clients, if *shared* was `TRUE` at registration time. The device driver instance must not disappear while its device entry is locked in the registry.

In the case of a catastrophic device shutdown (for example, hot-plug removal) the device entry lock owners are notified (through the *cli_handler* routine invocation) that the device entry must be released as soon as possible.

The `svDeviceEntry` routine returns the device entry associated with a given client identifier. The *cli_id* argument specifies the client identifier previously returned by `svDeviceLookup` .

The `svDeviceRelease` routine releases the lock on the given device entry. The device entry is specified by the *cli_id* argument. Obviously, a device driver client should no longer access the device driver instance and the device node once the device entry has been released.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
svDeviceAlloc	+	+	-	+
svDeviceFree	+	+	-	+
svDeviceRegister	+	+	-	+
svDeviceUnregister	+	+	-	+
svDeviceEvent	+	+	+	-
svDeviceLookup	+	+	-	+
svDeviceRelease	+	+	-	+
svDeviceEntry	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`dtreeNodeRoot(9DKI)` , `svDkiThreadCall(9DKI)`

NAME	svDeviceRegister, svDeviceAlloc, svDeviceFree, svDeviceUnregister, svDeviceEvent, svDeviceLookup, svDeviceEntry, svDeviceRelease – device registry operations
SYNOPSIS	<pre>#include <dki/dki.h> void svDeviceRegister(DevRegId dev_id); DevRegId svDeviceAlloc(DevRegEntry * entry, unsigned int version, Bool shared, DevRelHandler handler); KnError svDeviceUnregister(DevRegId dev_id); void svDeviceEvent(DevRegId dev_id, DevEvent event, void* arg); void svDeviceFree(DevRegId dev_id); KnError svDeviceLookup(char * dev_class, unsigned int dev_version, unsigned int dev_unit, DevEventHandler cli_handler, void * cli_cookie, DevClientId * cli_id); DevRegEntry * svDeviceEntry(DevClientId cli_id); void svDeviceRelease(DevClientId cli_id); typedef struct { char* dev_class; void* dev_ops; void* dev_id; DevNode dev_node; } DevRegEntry;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to microkernel services implementing the device registry.
EXTENDED DESCRIPTION	<p>The device registry microkernel module implements a data base of driver instances servicing devices currently supported by the system. The device registry data base is populated by drivers which perform self-registration (using svDeviceRegister) at device initialization time. The device registry data base is accessed by driver clients in order to obtain a pointer to the driver instance servicing a given (logical) device. The device registry API is described below in detail.</p> <p>Note that only the svDeviceLookup , svDeviceRelease and svDeviceEntry microkernel calls should be used by driver clients. The rest of the API is dedicated to device drivers.</p>

`svDeviceAlloc` allocates a device registry entry for a given device driver instance. Note that the entry is allocated in an invalid state. This means that the entry is not included in the registry (it is not visible to clients via `svDeviceLookup`).

The entry becomes valid (visible to clients) when a subsequent `svDeviceRegister` is invoked. On the other hand, the driver is allowed to call `svDeviceEvent` on this type of entry. Shut down events signaled on an invalid entry are memorized by the device registry and they are processed when the entry becomes valid (that is, when `svDeviceRegister` is called).

The *entry* argument points to the `DevRegEntry` structure which designates the device driver instance. The `DevRegEntry` structure is described below:

<code>dev_class</code>	Points to a string specifying the device class name.
<code>dev_ops</code>	Points to a structure of driver service routines.
<code>dev_id</code>	Points to a handle which is usually passed back to the driver (as first argument) each time a driver service routine is invoked.
<code>dev_node</code>	Specifies the device node in the tree which is serviced by the driver. Note that a driver client must switch to the DKI thread context when accessing the device tree structure (see section "Device Node Properties").

The *version* argument specifies the driver interface version implemented by the device driver instance.

The *dev_class* field and *version* argument specify APIs implemented by the driver. This type of API is provided to driver clients as a structure of indirect functions implementing the API service routines. For instance, the "uart" device API is specified by the `UartDevOps` structure, the "timer" device API is specified by the `TimerDevOps` structure. A pointer to this type of structure is obtained by a device driver client from the *dev_ops* field of the device registry entry.

The *shared* argument specifies whether the device driver instance may be shared between multiple clients. In other words, it indicates whether the registered instance may be looked up multiple times.

The *handler* argument specifies the driver handler which is invoked by the device registry module as an acknowledgement to a shut-down event (see `svDeviceEvent`). *handler* is called when the last reference to the device registry

entry goes away and a shut-down event has been previously signalled on the entry. *entry* is passed back to the driver as the argument of *handler*.

In case of success, a non zero *DevRegId* is returned, otherwise `NULL` is returned.

Device registry assigns a logical unit number to the physical device. The logical unit number is unique within *dev_class*. The device registry handles a logical unit counter per class. All counters are initialized to zero. Each time a new device entry's is allocated, the current counter value is assigned to the entry logical unit number and the counter is incremented. Thus, the logical device order within a class corresponds to the allocation order.

The logical unit number is mainly used by device driver clients to iterate through a given device class looking for a certain device. Using a valid logical unit number, a client is able to access the corresponding device registry entry, in particular, the *dev_node* field which points to the device node. The physical device corresponding to a given logical unit number may then be detected by the device node properties and the node position within the device tree.

Note - The device driver may allocate (and then register) multiple entries in order to declare different classes (APIs) for the same physical device. For instance, a driver may implement orthogonal sets of operations like power management which are additional to the main device functionality. Another example is a device driver servicing a multi-function device. This type of driver would register multiple entries: one entry per device function. From the client's point of view, this type of device would be visible as multiple (mono-function) independent devices.

`svDeviceRegister` adds a given device registry entry to the registry. The entry must be previously allocated by `svDeviceAlloc`. The entry becomes valid (and therefore visible for clients) only if there is no shut-down event signalled on the entry. Otherwise, the entry remains invalid and the device registry module invokes a handler previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The *handler* routine is called in the DKI thread context.

Note - When *handler* is called the entry is no longer registered. Thus, it is useless to invoke `svDeviceUnregister` once *handler* has been called. On the other hand, in order to free memory resources, the driver should release the device registry entry invoking `svDeviceFree` once the entry is no longer used by the driver. Typically, the entry is no longer used by the driver when the connection to the parent bus/nexus driver is closed, and the parent driver is unable to invoke a call-back handler signalling a bus/nexus event.

`svDeviceUnregister` removes the device entry from the device registry (if the entry is valid). The device entry being removed is specified by the `dev_id` argument. `svDeviceUnregister` returns the following results when a valid entry is specified:

<code>K_OK</code>	The device entry has been successfully removed from the registry.
<code>K_EBUSY</code>	The device entry has not been removed from the registry because it is locked by a driver client. In other words, the <code>svDeviceLookup/Release</code> pairs are not balanced.

`svDeviceUnregister` returns `K_EBUSY` when an invalid entry is specified. The driver must remove all related device entries from the registry prior to being unloaded from the system. Typically, `svDeviceUnregister` is issued by the device driver (for each device driver instance) when the driver code is requested to be unloaded via `drv_unload`. The driver unload fails when the driver code is still being used by the system.

`svDeviceEvent` notifies the device registry module that a shutdown event has occurred (for example, that a hot-pluggable device has been removed). The shutdown event basically means that the driver instance should no longer be used by driver clients and the device entry must be removed from the registry (that is, the driver instance will disappear). The shutdown event is immediately propagated to all driver clients in order to stop the device usage and to release the device entry as soon as possible. The device is specified by the `dev_id` argument.

`svDeviceEvent` propagates the shutdown event invoking the `event_handler` routines specified by the driver clients in `svDeviceLookup`. The `event` and `arg` arguments are opaque for the device registry. They are passed (as arguments) to the client event handlers. `event` specifies the reason for the shut-down as follows:

`DEV_EVENT_SHUTDOWN` Normal device shut down

`DEV_EVENT_REMOVAL` Hot-plug (surprise) device removal

`arg` is `event` specific.

All driver clients are requested to release the device entry (as soon as possible) invoking the `svDeviceRelease` routine. Note that, prior to the shutdown event propagation, `svDeviceEvent` removes the device entry from the registry in order to prevent the entry from being found (and locked) by new driver clients. Once all driver client handlers are invoked, `svDeviceEvent` returns to the driver. Note that the real device shutdown has to be deferred until the `handler` routine invocation.

Once the device entry is released by the last driver client, the device registry module invokes *handler* previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The handler routine is called in the DKI thread context.

Typically, `svDeviceEvent` is used by a device driver servicing a hot-pluggable device. `svDeviceEvent` is invoked by the device driver when the driver is notified (by its parent driver) that the device has been removed from the bus.

Note - In the case of hot-plug removal, the device driver must still be operational until the device entry is released (ignoring all requests to the driver, for example). The driver is allowed to call `svDeviceEvent` on an invalid (unregistered) entry. In this type of case, *handler* will be invoked only when a registration attempt is made, that is, `svDeviceRegister` will be invoked. In addition, the registration will fail and the entry will remain invalid.

`svDeviceFree` releases a given device registry entry previously allocated by `svDeviceAlloc`. The device entry being released is specified by the *dev_id* argument. The entry being released must be invalid (unregistered). Thus, if the entry was registered by `svDeviceRegister`, the driver is allowed to release it if either the entry is successfully unregistered by `svDeviceUnregister` or *handler* (previously specified by `svDeviceAlloc`) is invoked.

`svDeviceLookup` searches the device entry in the registry matching the specified device class and logical unit.

<code>dev_class</code>	Specifies the device class.
<code>dev_version</code>	Specifies the minimum device driver interface version required.
<code>dev_unit</code>	Specifies the logical device unit in the class.
<code>cli_handler</code>	Specifies the event handler which is called when a device event is signalled.
<code>cli_cookie</code>	Specifies the first argument of <i>cli_handler</i> .
<code>cli_id</code>	Is an output argument identifying the client token on the matching device entry. This <i>cli_id</i> is passed back as an argument to <code>svDeviceEntry</code> to get the associated device entry, or to <code>svDeviceRelease</code> to release the device driver instance.

`svDeviceLookup` returns the following results:

K_OK	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was found in the registry.
K_EBUSY	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was not found in the registry. Or, the device is being used by another client and device sharing is not allowed.
K_EUNKNOWN	There is no device entry in the registry matching <i>dev_class</i> .
K_UNDEF	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Or, there is a device entry in the <i>dev_class</i> with a logical unit greater than <i>dev_unit</i> . In other words, <i>dev_unit</i> falls into a hole in the <i>dev_class</i> unit ranges. The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair found in the registry is implementing an older version of the interface than the one required.
K_ETOOMUCH	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Moreover <i>dev_unit</i> is greater than all existing logical units in the <i>dev_class</i> .

In case of success, the corresponding device entry is locked in the registry until a subsequent `svDeviceRelease` . Note that the device registry lock may or may not be exclusive, depending on the value of the *shared* argument specified in `svDeviceAlloc` . In other words, the device may be shared between multiple driver clients, if *shared* was `TRUE` at registration time. The device driver instance must not disappear while its device entry is locked in the registry.

In the case of a catastrophic device shutdown (for example, hot-plug removal) the device entry lock owners are notified (through the *cli_handler* routine invocation) that the device entry must be released as soon as possible.

The `svDeviceEntry` routine returns the device entry associated with a given client identifier. The *cli_id* argument specifies the client identifier previously returned by `svDeviceLookup` .

The `svDeviceRelease` routine releases the lock on the given device entry. The device entry is specified by the *cli_id* argument. Obviously, a device driver client should no longer access the device driver instance and the device node once the device entry has been released.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
svDeviceAlloc	+	+	-	+
svDeviceFree	+	+	-	+
svDeviceRegister	+	+	-	+
svDeviceUnregister	+	+	-	+
svDeviceEvent	+	+	+	-
svDeviceLookup	+	+	-	+
svDeviceRelease	+	+	-	+
svDeviceEntry	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`dtreeNodeRoot(9DKI)` , `svDkiThreadCall(9DKI)`

NAME	svDeviceRegister, svDeviceAlloc, svDeviceFree, svDeviceUnregister, svDeviceEvent, svDeviceLookup, svDeviceEntry, svDeviceRelease – device registry operations
SYNOPSIS	<pre>#include <dki/dki.h> void svDeviceRegister(DevRegId dev_id); DevRegId svDeviceAlloc(DevRegEntry * entry, unsigned int version, Bool shared, DevRelHandler handler); KnError svDeviceUnregister(DevRegId dev_id); void svDeviceEvent(DevRegId dev_id, DevEvent event, void* arg); void svDeviceFree(DevRegId dev_id); KnError svDeviceLookup(char * dev_class, unsigned int dev_version, unsigned int dev_unit, DevEventHandler cli_handler, void * cli_cookie, DevClientId * cli_id); DevRegEntry * svDeviceEntry(DevClientId cli_id); void svDeviceRelease(DevClientId cli_id); typedef struct { char* dev_class; void* dev_ops; void* dev_id; DevNode dev_node; } DevRegEntry;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to microkernel services implementing the device registry.
EXTENDED DESCRIPTION	<p>The device registry microkernel module implements a data base of driver instances servicing devices currently supported by the system. The device registry data base is populated by drivers which perform self-registration (using svDeviceRegister) at device initialization time. The device registry data base is accessed by driver clients in order to obtain a pointer to the driver instance servicing a given (logical) device. The device registry API is described below in detail.</p> <p>Note that only the svDeviceLookup, svDeviceRelease and svDeviceEntry microkernel calls should be used by driver clients. The rest of the API is dedicated to device drivers.</p>

`svDeviceAlloc` allocates a device registry entry for a given device driver instance. Note that the entry is allocated in an invalid state. This means that the entry is not included in the registry (it is not visible to clients via `svDeviceLookup`).

The entry becomes valid (visible to clients) when a subsequent `svDeviceRegister` is invoked. On the other hand, the driver is allowed to call `svDeviceEvent` on this type of entry. Shut down events signaled on an invalid entry are memorized by the device registry and they are processed when the entry becomes valid (that is, when `svDeviceRegister` is called).

The *entry* argument points to the `DevRegEntry` structure which designates the device driver instance. The `DevRegEntry` structure is described below:

<code>dev_class</code>	Points to a string specifying the device class name.
<code>dev_ops</code>	Points to a structure of driver service routines.
<code>dev_id</code>	Points to a handle which is usually passed back to the driver (as first argument) each time a driver service routine is invoked.
<code>dev_node</code>	Specifies the device node in the tree which is serviced by the driver. Note that a driver client must switch to the DKI thread context when accessing the device tree structure (see section "Device Node Properties").

The *version* argument specifies the driver interface version implemented by the device driver instance.

The *dev_class* field and *version* argument specify APIs implemented by the driver. This type of API is provided to driver clients as a structure of indirect functions implementing the API service routines. For instance, the "uart" device API is specified by the `UartDevOps` structure, the "timer" device API is specified by the `TimerDevOps` structure. A pointer to this type of structure is obtained by a device driver client from the *dev_ops* field of the device registry entry.

The *shared* argument specifies whether the device driver instance may be shared between multiple clients. In other words, it indicates whether the registered instance may be looked up multiple times.

The *handler* argument specifies the driver handler which is invoked by the device registry module as an acknowledgement to a shut-down event (see `svDeviceEvent`). *handler* is called when the last reference to the device registry

entry goes away and a shut-down event has been previously signalled on the entry. *entry* is passed back to the driver as the argument of *handler* .

In case of success, a non zero *DevRegId* is returned, otherwise `NULL` is returned.

Device registry assigns a logical unit number to the physical device. The logical unit number is unique within *dev_class* . The device registry handles a logical unit counter per class. All counters are initialized to zero. Each time a new device entry's is allocated, the current counter value is assigned to the entry logical unit number and the counter is incremented. Thus, the logical device order within a class corresponds to the allocation order.

The logical unit number is mainly used by device driver clients to iterate through a given device class looking for a certain device. Using a valid logical unit number, a client is able to access the corresponding device registry entry, in particular, the *dev_node* field which points to the device node. The physical device corresponding to a given logical unit number may then be detected by the device node properties and the node position within the device tree.

Note - The device driver may allocate (and then register) multiple entries in order to declare different classes (APIs) for the same physical device. For instance, a driver may implement orthogonal sets of operations like power management which are additional to the main device functionality. Another example is a device driver servicing a multi-function device. This type of driver would register multiple entries: one entry per device function. From the client's point of view, this type of device would be visible as multiple (mono-function) independent devices.

`svDeviceRegister` adds a given device registry entry to the registry. The entry must be previously allocated by `svDeviceAlloc` . The entry becomes valid (and therefore visible for clients) only if there is no shut-down event signalled on the entry. Otherwise, the entry remains invalid and the device registry module invokes a handler previously specified via `svDeviceAlloc` . A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler* . The *handler* routine is called in the DK1 thread context.

Note - When *handler* is called the entry is no longer registered. Thus, it is useless to invoke `svDeviceUnregister` once *handler* has been called. On the other hand, in order to free memory resources, the driver should release the device registry entry invoking `svDeviceFree` once the entry is no longer used by the driver. Typically, the entry is no longer used by the driver when the connection to the parent bus/nexus driver is closed, and the parent driver is unable to invoke a call-back handler signalling a bus/nexus event.

`svDeviceUnregister` removes the device entry from the device registry (if the entry is valid). The device entry being removed is specified by the `dev_id` argument. `svDeviceUnregister` returns the following results when a valid entry is specified:

<code>K_OK</code>	The device entry has been successfully removed from the registry.
<code>K_EBUSY</code>	The device entry has not been removed from the registry because it is locked by a driver client. In other words, the <code>svDeviceLookup/Release</code> pairs are not balanced.

`svDeviceUnregister` returns `K_EBUSY` when an invalid entry is specified. The driver must remove all related device entries from the registry prior to being unloaded from the system. Typically, `svDeviceUnregister` is issued by the device driver (for each device driver instance) when the driver code is requested to be unloaded via `drv_unload`. The driver unload fails when the driver code is still being used by the system.

`svDeviceEvent` notifies the device registry module that a shutdown event has occurred (for example, that a hot-pluggable device has been removed). The shutdown event basically means that the driver instance should no longer be used by driver clients and the device entry must be removed from the registry (that is, the driver instance will disappear). The shutdown event is immediately propagated to all driver clients in order to stop the device usage and to release the device entry as soon as possible. The device is specified by the `dev_id` argument.

`svDeviceEvent` propagates the shutdown event invoking the `event_handler` routines specified by the driver clients in `svDeviceLookup`. The `event` and `arg` arguments are opaque for the device registry. They are passed (as arguments) to the client event handlers. `event` specifies the reason for the shut-down as follows:

`DEV_EVENT_SHUTDOWN` Normal device shut down

`DEV_EVENT_REMOVAL` Hot-plug (surprise) device removal

`arg` is `event` specific.

All driver clients are requested to release the device entry (as soon as possible) invoking the `svDeviceRelease` routine. Note that, prior to the shutdown event propagation, `svDeviceEvent` removes the device entry from the registry in order to prevent the entry from being found (and locked) by new driver clients. Once all driver client handlers are invoked, `svDeviceEvent` returns to the driver. Note that the real device shutdown has to be deferred until the `handler` routine invocation.

Once the device entry is released by the last driver client, the device registry module invokes *handler* previously specified via `svDeviceAlloc`. A pointer to the `DevRegEntry` structure (previously specified via `svDeviceAlloc`) is passed back to the driver as the argument of *handler*. The handler routine is called in the DKI thread context.

Typically, `svDeviceEvent` is used by a device driver servicing a hot-pluggable device. `svDeviceEvent` is invoked by the device driver when the driver is notified (by its parent driver) that the device has been removed from the bus.

Note - In the case of hot-plug removal, the device driver must still be operational until the device entry is released (ignoring all requests to the driver, for example). The driver is allowed to call `svDeviceEvent` on an invalid (unregistered) entry. In this type of case, *handler* will be invoked only when a registration attempt is made, that is, `svDeviceRegister` will be invoked. In addition, the registration will fail and the entry will remain invalid.

`svDeviceFree` releases a given device registry entry previously allocated by `svDeviceAlloc`. The device entry being released is specified by the *dev_id* argument. The entry being released must be invalid (unregistered). Thus, if the entry was registered by `svDeviceRegister`, the driver is allowed to release it if either the entry is successfully unregistered by `svDeviceUnregister` or *handler* (previously specified by `svDeviceAlloc`) is invoked.

`svDeviceLookup` searches the device entry in the registry matching the specified device class and logical unit.

<code>dev_class</code>	Specifies the device class.
<code>dev_version</code>	Specifies the minimum device driver interface version required.
<code>dev_unit</code>	Specifies the logical device unit in the class.
<code>cli_handler</code>	Specifies the event handler which is called when a device event is signalled.
<code>cli_cookie</code>	Specifies the first argument of <i>cli_handler</i> .
<code>cli_id</code>	Is an output argument identifying the client token on the matching device entry. This <i>cli_id</i> is passed back as an argument to <code>svDeviceEntry</code> to get the associated device entry, or to <code>svDeviceRelease</code> to release the device driver instance.

`svDeviceLookup` returns the following results:

K_OK	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was found in the registry.
K_EBUSY	The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair was not found in the registry. Or, the device is being used by another client and device sharing is not allowed.
K_EUNKNOWN	There is no device entry in the registry matching <i>dev_class</i> .
K_UNDEF	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Or, there is a device entry in the <i>dev_class</i> with a logical unit greater than <i>dev_unit</i> . In other words, <i>dev_unit</i> falls into a hole in the <i>dev_class</i> unit ranges. The device entry matching the <i>dev_class</i> and <i>dev_unit</i> pair found in the registry is implementing an older version of the interface than the one required.
K_ETOOMUCH	There is no device entry in the registry matching <i>dev_unit</i> in the <i>dev_class</i> . Moreover <i>dev_unit</i> is greater than all existing logical units in the <i>dev_class</i> .

In case of success, the corresponding device entry is locked in the registry until a subsequent `svDeviceRelease` . Note that the device registry lock may or may not be exclusive, depending on the value of the *shared* argument specified in `svDeviceAlloc` . In other words, the device may be shared between multiple driver clients, if *shared* was `TRUE` at registration time. The device driver instance must not disappear while its device entry is locked in the registry.

In the case of a catastrophic device shutdown (for example, hot-plug removal) the device entry lock owners are notified (through the *cli_handler* routine invocation) that the device entry must be released as soon as possible.

The `svDeviceEntry` routine returns the device entry associated with a given client identifier. The *cli_id* argument specifies the client identifier previously returned by `svDeviceLookup` .

The `svDeviceRelease` routine releases the lock on the given device entry. The device entry is specified by the *cli_id* argument. Obviously, a device driver client should no longer access the device driver instance and the device node once the device entry has been released.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
svDeviceAlloc	+	+	-	+
svDeviceFree	+	+	-	+
svDeviceRegister	+	+	-	+
svDeviceUnregister	+	+	-	+
svDeviceEvent	+	+	+	-
svDeviceLookup	+	+	-	+
svDeviceRelease	+	+	-	+
svDeviceEntry	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`dtreeNodeRoot(9DKI)` , `svDkiThreadCall(9DKI)`

NAME	svDkiOpen, svDkiClose, svDkiEvent – system event management		
SYNOPSIS	<pre>#include <dki/dki.h> KnError svDkiOpen(DevNode dev_node, DkiEventHandler dev_evt_handler, DkiLoadHandler dev_load_handler, void * dev_cookie, DkiDevId * dev_id); void svDkiClose(DkiDevId dev_id); void svDkiEvent(DkiEvent event, void * arg);</pre>		
FEATURES	DKI		
DESCRIPTION	The system event management services are provided by the microkernel to the lowest layer drivers. They are used mainly to register event handlers for all running drivers, and to start propagating an event from the microkernel.		
EXTENDED DESCRIPTION	<p>Typically, a system reboot starts propagating a specific event from the microkernel to the lowest-layer level drivers. Those drivers then recursively propagate the event to the upper layer drivers by calling their <code>BusEventHandler</code> handler, registered at <code>open</code> time.</p> <p><code>svDkiOpen</code> must be issued by the lowest layer level drivers. It establishes a connection between the device driver and DKI.</p> <p>The <code>dev_node</code> argument specifies the device node (in the device tree) which is serviced by the device driver instance. In case of initialization, the device node is given as an argument of <code>drv_init</code> by the parent bus driver. In case of probing, the device node is either found (among existing child nodes attached to the parent node) or created (and attached to the parent node) by the device driver.</p> <p>The <code>dev_evt_handler</code> argument specifies the device driver handler which is invoked by the DKI when an event occurs. It takes three arguments. The first argument is the <code>dev_cookie</code>. The second one specifies the DKI event type. The third argument points to a structure which is event type specific.</p> <p>DKI events are defined below:</p> <table border="0"> <tr> <td style="vertical-align: top;">DKI_SYS_SHUTDOWN</td> <td>Notifies a device driver that the system is going to be shut down. The device driver should propagate the event to all client drivers if it is a bus driver. Otherwise, it should reset the device hardware and return from the event handler. Note that the driver must neither notify clients nor free allocated resources.</td> </tr> </table> <p>Typically, <code>dev_evt_handler</code> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level.</p>	DKI_SYS_SHUTDOWN	Notifies a device driver that the system is going to be shut down. The device driver should propagate the event to all client drivers if it is a bus driver. Otherwise, it should reset the device hardware and return from the event handler. Note that the driver must neither notify clients nor free allocated resources.
DKI_SYS_SHUTDOWN	Notifies a device driver that the system is going to be shut down. The device driver should propagate the event to all client drivers if it is a bus driver. Otherwise, it should reset the device hardware and return from the event handler. Note that the driver must neither notify clients nor free allocated resources.		

The *dev_load_handler* argument specifies the device driver handler which is invoked by the DKI when a new driver has been dynamically loaded. It is invoked passing *dev_cookie* as the only argument. Note that this *dev_load_handler* handler is optional.

Typically, it should be used only by bus drivers supporting dynamically loadable device drivers, and should be set to `NULL` by all other drivers. This type of bus driver handler should manage the newly loaded driver in a similar way to the driver's initialization at boot time. That is, associate the driver with a device node, and initialize it, in order to create a running instance of the newly loaded driver. Note that the *dev_load_handler* routine is invoked in the DKI thread context.

The *dev_cookie* argument specifies a device driver cookie. It is opaque for the DKI. *dev_cookie* is passed back to the driver when *dev_evt_handler* or *dev_load_handler* is invoked.

Upon successful completion, `svDkiOpen` returns `K_OK` and an identifier designating the DKI/device connection is returned in the *dev_id* argument. The *dev_id* is opaque for the driver, it must be passed back to the DKI as an argument of the `svDkiClose` service routine. In case of failure, an error code is returned as described below:

<code>K_EINVAL</code>	The <i>dev_node</i> argument provided is not a valid device tree node.
<code>K_EBUSY</code>	The <i>dev_node</i> device tree node provided is already in use (associated to another driver).
<code>K_ENOMEM</code>	The system is out of memory.

The `svDkiClose` routine releases the DKI/driver connection. It must be the last call issued by the driver.

`svDkiEvent` is used to start propagating a given event to the device driver hierarchy. It calls all `DkiEventHandler` handlers registered through `svDkiOpen` to signal all running drivers that the given event occurred. The *event* argument specifies the event to propagate. The *arg* argument specifies the specific parameter associated with the event type.

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svDkiOpen</code>	-	+	-	+

Allowed Calling Contexts

svDkiClose	-	+	-	+
svDkiEvent	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	svDkiOpen, svDkiClose, svDkiEvent – system event management
SYNOPSIS	<pre>#include <dki/dki.h> KnError svDkiOpen(DevNode dev_node, DkiEventHandler dev_evt_handler, DkiLoadHandler dev_load_handler, void * dev_cookie, DkiDevId * dev_id); void svDkiClose(DkiDevId dev_id); void svDkiEvent(DkiEvent event, void * arg);</pre>
FEATURES	DKI
DESCRIPTION	The system event management services are provided by the microkernel to the lowest layer drivers. They are used mainly to register event handlers for all running drivers, and to start propagating an event from the microkernel.
EXTENDED DESCRIPTION	<p>Typically, a system reboot starts propagating a specific event from the microkernel to the lowest-layer level drivers. Those drivers then recursively propagate the event to the upper layer drivers by calling their <code>BusEventHandler</code> handler, registered at open time.</p> <p><code>svDkiOpen</code> must be issued by the lowest layer level drivers. It establishes a connection between the device driver and DKI.</p> <p>The <code>dev_node</code> argument specifies the device node (in the device tree) which is serviced by the device driver instance. In case of initialization, the device node is given as an argument of <code>drv_init</code> by the parent bus driver. In case of probing, the device node is either found (among existing child nodes attached to the parent node) or created (and attached to the parent node) by the device driver.</p> <p>The <code>dev_evt_handler</code> argument specifies the device driver handler which is invoked by the DKI when an event occurs. It takes three arguments. The first argument is the <code>dev_cookie</code>. The second one specifies the DKI event type. The third argument points to a structure which is event type specific.</p> <p>DKI events are defined below:</p> <p><code>DKI_SYS_SHUTDOWN</code> Notifies a device driver that the system is going to be shut down. The device driver should propagate the event to all client drivers if it is a bus driver. Otherwise, it should reset the device hardware and return from the event handler. Note that the driver must neither notify clients nor free allocated resources.</p> <p>Typically, <code>dev_evt_handler</code> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level.</p>

The *dev_load_handler* argument specifies the device driver handler which is invoked by the DKI when a new driver has been dynamically loaded. It is invoked passing *dev_cookie* as the only argument. Note that this *dev_load_handler* handler is optional.

Typically, it should be used only by bus drivers supporting dynamically loadable device drivers, and should be set to `NULL` by all other drivers. This type of bus driver handler should manage the newly loaded driver in a similar way to the driver's initialization at boot time. That is, associate the driver with a device node, and initialize it, in order to create a running instance of the newly loaded driver. Note that the *dev_load_handler* routine is invoked in the DKI thread context.

The *dev_cookie* argument specifies a device driver cookie. It is opaque for the DKI. *dev_cookie* is passed back to the driver when *dev_evt_handler* or *dev_load_handler* is invoked.

Upon successful completion, `svDkiOpen` returns `K_OK` and an identifier designating the DKI/device connection is returned in the *dev_id* argument. The *dev_id* is opaque for the driver, it must be passed back to the DKI as an argument of the `svDkiClose` service routine. In case of failure, an error code is returned as described below:

- `K_EINVAL` The *dev_node* argument provided is not a valid device tree node.
- `K_EBUSY` The *dev_node* device tree node provided is already in use (associated to another driver).
- `K_ENOMEM` The system is out of memory.

The `svDkiClose` routine releases the DKI/driver connection. It must be the last call issued by the driver.

`svDkiEvent` is used to start propagating a given event to the device driver hierarchy. It calls all `DkiEventHandler` handlers registered through `svDkiOpen` to signal all running drivers that the given event occurred. The *event* argument specifies the event to propagate. The *arg* argument specifies the specific parameter associated with the event type.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svDkiOpen</code>	-	+	-	+

svDkiClose	-	+	-	+
svDkiEvent	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	svDkiOpen, svDkiClose, svDkiEvent – system event management		
SYNOPSIS	<pre>#include <dki/dki.h> KnError svDkiOpen(DevNode dev_node, DkiEventHandler dev_evt_handler, DkiLoadHandler dev_load_handler, void * dev_cookie, DkiDevId * dev_id); void svDkiClose(DkiDevId dev_id); void svDkiEvent(DkiEvent event, void * arg);</pre>		
FEATURES	DKI		
DESCRIPTION	The system event management services are provided by the microkernel to the lowest layer drivers. They are used mainly to register event handlers for all running drivers, and to start propagating an event from the microkernel.		
EXTENDED DESCRIPTION	<p>Typically, a system reboot starts propagating a specific event from the microkernel to the lowest-layer level drivers. Those drivers then recursively propagate the event to the upper layer drivers by calling their <code>BusEventHandler</code> handler, registered at <code>open</code> time.</p> <p><code>svDkiOpen</code> must be issued by the lowest layer level drivers. It establishes a connection between the device driver and DKI.</p> <p>The <code>dev_node</code> argument specifies the device node (in the device tree) which is serviced by the device driver instance. In case of initialization, the device node is given as an argument of <code>drv_init</code> by the parent bus driver. In case of probing, the device node is either found (among existing child nodes attached to the parent node) or created (and attached to the parent node) by the device driver.</p> <p>The <code>dev_evt_handler</code> argument specifies the device driver handler which is invoked by the DKI when an event occurs. It takes three arguments. The first argument is the <code>dev_cookie</code>. The second one specifies the DKI event type. The third argument points to a structure which is event type specific.</p> <p>DKI events are defined below:</p> <table border="0"> <tr> <td style="vertical-align: top;">DKI_SYS_SHUTDOWN</td> <td>Notifies a device driver that the system is going to be shut down. The device driver should propagate the event to all client drivers if it is a bus driver. Otherwise, it should reset the device hardware and return from the event handler. Note that the driver must neither notify clients nor free allocated resources.</td> </tr> </table> <p>Typically, <code>dev_evt_handler</code> is called as an interrupt handler and therefore the handler implementation must be restricted to the API allowed at interrupt level.</p>	DKI_SYS_SHUTDOWN	Notifies a device driver that the system is going to be shut down. The device driver should propagate the event to all client drivers if it is a bus driver. Otherwise, it should reset the device hardware and return from the event handler. Note that the driver must neither notify clients nor free allocated resources.
DKI_SYS_SHUTDOWN	Notifies a device driver that the system is going to be shut down. The device driver should propagate the event to all client drivers if it is a bus driver. Otherwise, it should reset the device hardware and return from the event handler. Note that the driver must neither notify clients nor free allocated resources.		

The *dev_load_handler* argument specifies the device driver handler which is invoked by the DKI when a new driver has been dynamically loaded. It is invoked passing *dev_cookie* as the only argument. Note that this *dev_load_handler* handler is optional.

Typically, it should be used only by bus drivers supporting dynamically loadable device drivers, and should be set to `NULL` by all other drivers. This type of bus driver handler should manage the newly loaded driver in a similar way to the driver's initialization at boot time. That is, associate the driver with a device node, and initialize it, in order to create a running instance of the newly loaded driver. Note that the *dev_load_handler* routine is invoked in the DKI thread context.

The *dev_cookie* argument specifies a device driver cookie. It is opaque for the DKI. *dev_cookie* is passed back to the driver when *dev_evt_handler* or *dev_load_handler* is invoked.

Upon successful completion, `svDkiOpen` returns `K_OK` and an identifier designating the DKI/device connection is returned in the *dev_id* argument. The *dev_id* is opaque for the driver, it must be passed back to the DKI as an argument of the `svDkiClose` service routine. In case of failure, an error code is returned as described below:

<code>K_EINVAL</code>	The <i>dev_node</i> argument provided is not a valid device tree node.
<code>K_EBUSY</code>	The <i>dev_node</i> device tree node provided is already in use (associated to another driver).
<code>K_ENOMEM</code>	The system is out of memory.

The `svDkiClose` routine releases the DKI/driver connection. It must be the last call issued by the driver.

`svDkiEvent` is used to start propagating a given event to the device driver hierarchy. It calls all `DkiEventHandler` handlers registered through `svDkiOpen` to signal all running drivers that the given event occurred. The *event* argument specifies the event to propagate. The *arg* argument specifies the specific parameter associated with the event type.

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svDkiOpen</code>	-	+	-	+

Allowed Calling Contexts

svDkiClose	-	+	-	+
svDkiEvent	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	svDkiThreadCall, svDkiThreadTrigger – call a routine in the DKI thread context
SYNOPSIS	<pre>#include <dki/dki.h> void svDkiThreadCall(DkiCallToken * token, DkiCallHandler handler, void * cookie); void svDkiThreadTrigger(DkiCallToken * token, DkiCallHandler handler, void * cookie);</pre>
FEATURES	DKI
DESCRIPTION	Provides synchronization services through the DKI thread.
EXTENDED DESCRIPTION	<p>The DKI thread is launched by the microkernel at initialization time.</p> <p>The purpose of this thread is to synchronize identical calls to DKI and DDI services issued from different threads at drivers' initialization time as well as at runtime. This is typically used for initialization and shutdown of drivers. By ensuring such synchronization, the DKI thread avoids using any other synchronization mechanism (locks) in the driver implementations.</p> <p>Refer to the tables in each section, that indicate the allowed calling level, to know which services should to be called in the context of the DKI thread to ensure coherency. The DKI thread is involved as a synchronization mechanism in the following two cases:</p> <ul style="list-style-type: none"> ■ <i>Normal case</i> <p>In the normal case, all calls related to initialization / shutdown of the drivers are done implicitly in the context of the DKI thread. That means that drivers should not worry about synchronization, and do nothing, because their routines are called directly from the DKI thread.</p> ■ <i>Specific cases</i> <p>There are two special cases in which a driver should use DKI thread services explicitly to ensure synchronization:</p> <ul style="list-style-type: none"> ■ <i>Hot-pluggable device drivers</i> <p>In the case of a hot-pluggable device driver, the initialization/shutdown process has to be executed at runtime and not as part of the kernel/drivers initialization process. In this type of case, drivers should use the DKI thread services below to synchronize explicitly with drivers already running.</p> ■ <i>Deferred driver initialization</i> <p>In some cases, a driver may defer its device initialization until it is opened. This is a way to resolve conflicts about usage of the same resource by multiple drivers. In that way, drivers sharing a resource can be loaded at same time, if they are not opened at the same time. In this type of deferred initialization scheme, the initialization/shutdown process must be executed at runtime (at time of open/close) and not as</p>

part of the kernel/drivers initialization process. Thus, these kinds of drivers should also use the DKI thread services below to synchronize explicitly with drivers already running.

`svDkiThreadCall` synchronously invokes a routine in the context of the DKI thread. Synchronously means that the caller is blocked until the invoked routine returns.

The *token* argument is the address of a `DkiCallToken` structure which must be allocated by the caller. This structure is opaque for the driver, and is used only by the DKI thread. Note that the same structure may be reused for subsequent calls. The *handler* argument specifies the routine to call. The *cookie* argument specifies the argument to pass to the *handler* routine when called.

`svDkiThreadTrigger` asynchronously invokes a routine in the context of the DKI thread. Asynchronously means that the function immediately returns to the caller without waiting for the invoked routine to return. If the driver needs to know when the handler returns, it should use any synchronization mechanism inside the handler itself. The *token* argument is the address of a `DkiCallToken` structure which must be allocated by the caller. This structure is opaque to the driver, and is used only by the DKI thread. Note that the same structure may be reused for subsequent calls, once the handler is invoked. The *handler* argument specifies the routine to call. The *cookie* argument specifies the argument to pass to the *handler* routine when called.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>svDkiThreadCall</code>	+	+	-	+
<code>svDkiThreadTrigger</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svDkiThreadCall, svDkiThreadTrigger – call a routine in the DKI thread context
SYNOPSIS	<pre>#include <dki/dki.h> void svDkiThreadCall(DkiCallToken * token, DkiCallHandler handler, void * cookie); void svDkiThreadTrigger(DkiCallToken * token, DkiCallHandler handler, void * cookie);</pre>
FEATURES	DKI
DESCRIPTION	Provides synchronization services through the DKI thread.
EXTENDED DESCRIPTION	<p>The DKI thread is launched by the microkernel at initialization time.</p> <p>The purpose of this thread is to synchronize identical calls to DKI and DDI services issued from different threads at drivers' initialization time as well as at runtime. This is typically used for initialization and shutdown of drivers. By ensuring such synchronization, the DKI thread avoids using any other synchronization mechanism (locks) in the driver implementations.</p> <p>Refer to the tables in each section, that indicate the allowed calling level, to know which services should to be called in the context of the DKI thread to ensure coherency. The DKI thread is involved as a synchronization mechanism in the following two cases:</p> <ul style="list-style-type: none"> ■ <i>Normal case</i> <p>In the normal case, all calls related to initialization / shutdown of the drivers are done implicitly in the context of the DKI thread. That means that drivers should not worry about synchronization, and do nothing, because their routines are called directly from the DKI thread.</p> ■ <i>Specific cases</i> <p>There are two special cases in which a driver should use DKI thread services explicitly to ensure synchronization:</p> <ul style="list-style-type: none"> ■ <i>Hot-pluggable device drivers</i> <p>In the case of a hot-pluggable device driver, the initialization/shutdown process has to be executed at runtime and not as part of the kernel/drivers initialization process. In this type of case, drivers should use the DKI thread services below to synchronize explicitly with drivers already running.</p> ■ <i>Deferred driver initialization</i> <p>In some cases, a driver may defer its device initialization until it is opened. This is a way to resolve conflicts about usage of the same resource by multiple drivers. In that way, drivers sharing a resource can be loaded at same time, if they are not opened at the same time. In this type of deferred initialization scheme, the initialization/shutdown process must be executed at runtime (at time of open/close) and not as</p>

part of the kernel/drivers initialization process. Thus, these kinds of drivers should also use the DKI thread services below to synchronize explicitly with drivers already running.

`svDkiThreadCall` synchronously invokes a routine in the context of the DKI thread. Synchronously means that the caller is blocked until the invoked routine returns.

The *token* argument is the address of a `DkiCallToken` structure which must be allocated by the caller. This structure is opaque for the driver, and is used only by the DKI thread. Note that the same structure may be reused for subsequent calls. The *handler* argument specifies the routine to call. The *cookie* argument specifies the argument to pass to the *handler* routine when called.

`svDkiThreadTrigger` asynchronously invokes a routine in the context of the DKI thread. Asynchronously means that the function immediately returns to the caller without waiting for the invoked routine to return. If the driver needs to know when the handler returns, it should use any synchronization mechanism inside the handler itself. The *token* argument is the address of a `DkiCallToken` structure which must be allocated by the caller. This structure is opaque to the driver, and is used only by the DKI thread. Note that the same structure may be reused for subsequent calls, once the handler is invoked. The *handler* argument specifies the routine to call. The *cookie* argument specifies the argument to pass to the *handler* routine when called.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>svDkiThreadCall</code>	+	+	-	+
<code>svDkiThreadTrigger</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svDriverRegister, svDriverLookupFirst, svDriverLookupNext, svDriverRelease, svDriverEntry, svDriverCap, svDriverUnregister – driver registry operations
SYNOPSIS	<pre> #include <dki/dki.h> KnError svDriverRegister(DrvRegEntry * drv_entry); DrvRegId svDriverLookupFirst(void); DrvRegId svDriverLookupNext(DrvRegId drv_id); void svDriverRelease(DrvRegId drv_id); DrvRegEntry * svDriverEntry(DrvRegId drv_id); KnCap * svDriverCap(DrvRegId drv_id); KnError svDriverUnregister(DrvRegId drv_id); typedef struct { char* drv_name; char* drv_info; char* bus_class; int bus_version; void (*drv_probe) (DevNode bus_node, void* bus_ops, void* bus_id); void (*drv_bind) (DevNode bus_node); void (*drv_init) (DevNode dev_node, void* bus_ops, void* bus_id); KnError (*drv_unload) (); } DrvRegEntry; </pre>

FEATURES

DKI

DESCRIPTION

The `driver_registry` module implements a data base of drivers which have registered within the system. The `driver_registry` data base is populated by drivers which perform self-registration (using `svDriverRegister`) at driver initialization time.

EXTENDED DESCRIPTION

`svDriverRegister` adds the driver entry to the driver registry. It returns `K_OK` in case of success, otherwise `K_ENOMEM` is returned to indicate that the system is out of memory. The `drv_entry` argument points to the `DrvRegEntry` structure specifying driver properties and static driver routines. The `DrvRegEntry` structure is described below.

<code>drv_name</code>	Points to a string specifying the driver name. For example, the driver file name.
<code>drv_info</code>	Points to a string specifying extra information about the driver component, such as version or author.
<code>bus_class</code>	Points to a string specifying the class of the parent driver API required for the driver, such as "pci".
<code>bus_version</code>	Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code> , the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.
<code>drv_probe</code>	Points to a static driver routine which performs the device enumeration/probing on the bus.
<code>drv_bind</code>	Points to a static driver routine which performs the driver-to-device binding.
<code>drv_init</code>	Points to a static driver routine which clones an instance of the driver for the given device.
<code>drv_unload</code>	Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.

The `bus_class` and `bus_version` fields specify a parent bus/nexus API required for the driver. This type of API is provided to the driver as a structure of indirect

functions implementing the API service routines. For instance, the *"pci"* bus API is specified by the `PciBusOps` structure, the *"isa"* bus API is specified by `IsaBusOps`. When a bus driver invokes the `drv_probe` or `drv_init` driver's routine, it provides a pointer to the structure of the bus service routines. The structure type corresponds to the `bus_class` field value.

Note that a bus driver may provide multiple APIs. A typical example is a PCI bus driver providing the common and PCI bus APIs. The common bus driver API is named by *"bus"* and specified by the `BusOps` structure. The PCI bus driver API is named by *"pci"* and specified by the `PciBusOps` structure. `BusOps` actually, provides a subset of services provided by `PciBusOps`. This type of bus driver is able to support drivers which use either common (*"bus"*) or PCI (*"pci"*) parent bus interfaces. When the bus driver invokes a child driver, it gives a pointer to either the `BusOps` or `PciBusOps` structure depending on the `bus_class` specified in the child driver registry entry.

Note also that a `PROP_DRIVER` may be adaptive to the parent bus API. In other words, such a driver is able to run on top of a number of different buses (for example, *"pci"* and *"isa"*). Typically, this type of driver is composed of two parts: bus class specific and bus class independent.

The bus class specific part of the driver code mainly deals with the device probing and initialization. In addition, it provides an abstraction layer in order to hide the bus class dependencies from the bus class independent part of the driver code. This kind of multi-bus driver should be registered multiple times in the driver registry. Each entry specifies a given bus class API (via the `bus_class` field) on top of which the driver may run. In order to determine to which bus class the driver is applied, the `drv_probe`, `drv_bind` and `drv_init` routines have to be entry specific (that is, bus class API specific). Under these conditions, when a given driver's routine is invoked by a bus driver, the driver detects the bus class to which it is applied and casts the `bus_ops` argument to the appropriate structure (for example, `PciBusOps`).

The `drv_probe` routine is invoked by a bus/nexus driver when the `bus_class` specified in the registry entry matches the bus/nexus driver class. `drv_probe` is called with three arguments.

The `dev_node` argument specifies the parent device node.

The `bus_ops / bus_id` pair specifies the parent device driver.

`bus_ops` points to a structure of service routines implementing a bus driver API. This structure is bus class specific and corresponds to the `bus_class` specified by the driver registry entry.

bus_id is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The `drv_probe` routine is optional. In cases where the `drv_probe` routine is not provided by the driver, the *drv_probe* field must be set to `NULL`.

The `drv_bind` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class.

`drv_bind` is called with one argument, *dev_node*, which specifies a given device node. The `drv_bind` routine enables the driver to perform a driver-to-device binding. Typically, the driver examines properties attached to the device node in order to determine the type of device at the node and to check whether the device may be serviced by the driver. If the check is positive, the driver binds the driver to the device node by attaching a `PROP_DRIVER` property to the device node. The property value specifies the driver name. The parent bus driver uses such a property to determine the name of the driver servicing the device. In other words, via the `PROP_DRIVER` property, the child driver gives its name to the parent bus driver, asking it to invoke the `drv_init` routine on that device. Note that, if a "driver" property is already present in the device node, then the `drv_bind` routine can not continue; `drv_bind` should not override existing driver-to-device binding.

The `drv_bind` routine is optional. If the `drv_bind` routine is not provided by the driver, the *drv_bind* field must be set to `NULL`.

The `drv_init` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class and a given driver is bound to a given device node.

`drv_init` is called with three arguments. The *dev_node* argument specifies the device node for which a device driver instance should be created. The *bus_ops* / *bus_id* pair specifies the parent device driver. *bus_ops* points to a structure of service routines implementing a bus driver API.

This structure is bus class specific and corresponds to the *bus_class* specified by the driver registry entry. *bus_id* is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The purpose of `drv_init` is to create an instance of the driver servicing the given device, to perform a device hardware initialization and to register the driver instance in the device registry.

Typically, `drv_init` would read the device and bus node properties in order to obtain the assigned bus resources and tunable parameters related to the bus/device. The `drv_init` routine is optional. In a case when the `drv_init` routine is not provided by the driver, the *drv_init* field must be set to `NULL`. A typical example of a probe-only driver is a self-identifying bus enumerator (for example, a PCI enumerator) which is implemented as driver. This type of

driver has the `drv_probe` routine which enumerates devices residing on the bus and creates device nodes. This type of driver obviously does not have the `drv_init` routine.

`drv_unload` is called by the driver registry module (more exactly by the `svDriverUnregister` routine) when somebody wishes to unload the driver code from the system. The purpose of `drv_unload` is to check that the device driver code is not currently being used. The driver must check, for each driver instance, whether it is locked in the device registry.

In case of success, all device instances are removed from the device registry and `K_OK` is returned. Otherwise, the device entries are unchanged in the device registry and `K_EBUSY` is returned.

The `drv_unload` routine is optional. In cases when `drv_unload` is not provided by the driver, the `drv_unload` field must be set to `NULL`. Note that, in this case, the driver code cannot be unloaded.

`svDriverLookupFirst` returns the first driver entry in the registry. When the registry is not empty a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

The driver entry lock is not exclusive. It is simply a counter which prevents the driver component from being unloaded when it is still being used by a bus/nexus driver for probing or initialization.

`svDriverLookupNext` returns the next driver entry in the registry. The current entry is specified by the `drv_id` argument. The current entry must be locked by a previously called `svDriverLookupFirst` or `svDriverLookupNext`. If the current entry is not the last one in the registry a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the next driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

`svDriverRelease` releases the lock of the driver entry specified by the `drv_id` argument.

`svDriverEntry` returns a pointer to the driver entry structure specified by the `drv_id` argument. The device entry being accessed must have been previously locked using `svDriverLookupFirst` or `svDriverLookupNext`. Note that the device entry structure is read-only.

svDriverCap returns a pointer to the driver actor capability. The driver entry is specified by the *drv_id* argument. The driver entry being accessed must have been previously locked using svDriverLookupFirst or svDriverLookupNext . Note that if a given driver entry is registered by a built-in driver, a NULL pointer is returned. In other words, the driver actor capability makes sense only for dynamically loaded drivers. The actor capability may be used by an application in order to delete the driver actor once the driver entry is unregistered. Note that the driver capability structure is read-only.

svDriverUnregister tries to remove the driver entry specified by the *dev_id* argument from the driver registry. The device entry being removed must have been previously locked using svDriverLookupFirst or svDriverLookupNext .

In case of success, K_OK is returned, otherwise K_EBUSY is returned. The K_EBUSY result means that either the driver entry is locked in the driver registry (that is, a static driver routine is currently being used by a bus/nexus driver) or an instance of the device driver is locked in the device registry (that is, there is a driver instance which is currently being used by a driver client). Note that when K_EBUSY is returned, the driver entry remains locked in the registry and should be unlocked explicitly by svDriverRelease .

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svDriverRegister	+	+	-	+
svDriverLookupFirst	+	+	-	+
svDriverLookupNext	+	+	-	+
svDriverRelease	+	+	-	+
svDriverEntry	+	+	-	-
svDriverCap	+	+	-	-
svDriverUnregister	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

dreeNodeRoot(9DKI) , svDkiThreadCall(9DKI)

NAME	svDriverRegister, svDriverLookupFirst, svDriverLookupNext, svDriverRelease, svDriverEntry, svDriverCap, svDriverUnregister – driver registry operations
SYNOPSIS	<pre> #include <dki/dki.h> KnError svDriverRegister(DrvRegEntry * drv_entry); DrvRegId svDriverLookupFirst(void); DrvRegId svDriverLookupNext(DrvRegId drv_id); void svDriverRelease(DrvRegId drv_id); DrvRegEntry * svDriverEntry(DrvRegId drv_id); KnCap * svDriverCap(DrvRegId drv_id); KnError svDriverUnregister(DrvRegId drv_id); typedef struct { char* drv_name; char* drv_info; char* bus_class; int bus_version; void (*drv_probe) (DevNode bus_node, void* bus_ops, void* bus_id); void (*drv_bind) (DevNode bus_node); void (*drv_init) (DevNode dev_node, void* bus_ops, void* bus_id); KnError (*drv_unload) (); } DrvRegEntry; </pre>

FEATURES
DESCRIPTION

DKI

The `driver_registry` module implements a data base of drivers which have registered within the system. The `driver_registry` data base is populated by drivers which perform self-registration (using `svDriverRegister`) at driver initialization time.

EXTENDED
DESCRIPTION

`svDriverRegister` adds the driver entry to the driver registry. It returns `K_OK` in case of success, otherwise `K_ENOMEM` is returned to indicate that the system is out of memory. The `drv_entry` argument points to the `DrvRegEntry` structure specifying driver properties and static driver routines. The `DrvRegEntry` structure is described below.

<code>drv_name</code>	Points to a string specifying the driver name. For example, the driver file name.
<code>drv_info</code>	Points to a string specifying extra information about the driver component, such as version or author.
<code>bus_class</code>	Points to a string specifying the class of the parent driver API required for the driver, such as "pci".
<code>bus_version</code>	Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code> , the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.
<code>drv_probe</code>	Points to a static driver routine which performs the device enumeration/probing on the bus.
<code>drv_bind</code>	Points to a static driver routine which performs the driver-to-device binding.
<code>drv_init</code>	Points to a static driver routine which clones an instance of the driver for the given device.
<code>drv_unload</code>	Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.

The `bus_class` and `bus_version` fields specify a parent bus/nexus API required for the driver. This type of API is provided to the driver as a structure of indirect

functions implementing the API service routines. For instance, the *"pci"* bus API is specified by the `PciBusOps` structure, the *"isa"* bus API is specified by `IsaBusOps`. When a bus driver invokes the `drv_probe` or `drv_init` driver's routine, it provides a pointer to the structure of the bus service routines. The structure type corresponds to the `bus_class` field value.

Note that a bus driver may provide multiple APIs. A typical example is a PCI bus driver providing the common and PCI bus APIs. The common bus driver API is named by *"bus"* and specified by the `BusOps` structure. The PCI bus driver API is named by *"pci"* and specified by the `PciBusOps` structure. `BusOps` actually provides a subset of services provided by `PciBusOps`. This type of bus driver is able to support drivers which use either common (*"bus"*) or PCI (*"pci"*) parent bus interfaces. When the bus driver invokes a child driver, it gives a pointer to either the `BusOps` or `PciBusOps` structure depending on the `bus_class` specified in the child driver registry entry.

Note also that a `PROP_DRIVER` may be adaptive to the parent bus API. In other words, such a driver is able to run on top of a number of different buses (for example, *"pci"* and *"isa"*). Typically, this type of driver is composed of two parts: bus class specific and bus class independent.

The bus class specific part of the driver code mainly deals with the device probing and initialization. In addition, it provides an abstraction layer in order to hide the bus class dependencies from the bus class independent part of the driver code. This kind of multi-bus driver should be registered multiple times in the driver registry. Each entry specifies a given bus class API (via the `bus_class` field) on top of which the driver may run. In order to determine to which bus class the driver is applied, the `drv_probe`, `drv_bind` and `drv_init` routines have to be entry specific (that is, bus class API specific). Under these conditions, when a given driver's routine is invoked by a bus driver, the driver detects the bus class to which it is applied and casts the `bus_ops` argument to the appropriate structure (for example, `PciBusOps`).

The `drv_probe` routine is invoked by a bus/nexus driver when the `bus_class` specified in the registry entry matches the bus/nexus driver class. `drv_probe` is called with three arguments.

The `dev_node` argument specifies the parent device node.

The `bus_ops / bus_id` pair specifies the parent device driver.

`bus_ops` points to a structure of service routines implementing a bus driver API. This structure is bus class specific and corresponds to the `bus_class` specified by the driver registry entry.

bus_id is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The `drv_probe` routine is optional. In cases where the `drv_probe` routine is not provided by the driver, the *drv_probe* field must be set to `NULL`.

The `drv_bind` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class.

`drv_bind` is called with one argument, *dev_node*, which specifies a given device node. The `drv_bind` routine enables the driver to perform a driver-to-device binding. Typically, the driver examines properties attached to the device node in order to determine the type of device at the node and to check whether the device may be serviced by the driver. If the check is positive, the driver binds the driver to the device node by attaching a `PROP_DRIVER` property to the device node. The property value specifies the driver name. The parent bus driver uses such a property to determine the name of the driver servicing the device. In other words, via the `PROP_DRIVER` property, the child driver gives its name to the parent bus driver, asking it to invoke the `drv_init` routine on that device. Note that, if a "driver" property is already present in the device node, then the `drv_bind` routine can not continue; `drv_bind` should not override existing driver-to-device binding.

The `drv_bind` routine is optional. If the `drv_bind` routine is not provided by the driver, the *drv_bind* field must be set to `NULL`.

The `drv_init` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class and a given driver is bound to a given device node.

`drv_init` is called with three arguments. The *dev_node* argument specifies the device node for which a device driver instance should be created. The *bus_ops* / *bus_id* pair specifies the parent device driver. *bus_ops* points to a structure of service routines implementing a bus driver API.

This structure is bus class specific and corresponds to the *bus_class* specified by the driver registry entry. *bus_id* is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The purpose of `drv_init` is to create an instance of the driver servicing the given device, to perform a device hardware initialization and to register the driver instance in the device registry.

Typically, `drv_init` would read the device and bus node properties in order to obtain the assigned bus resources and tunable parameters related to the bus/device. The `drv_init` routine is optional. In a case when the `drv_init` routine is not provided by the driver, the *drv_init* field must be set to `NULL`. A typical example of a probe-only driver is a self-identifying bus enumerator (for example, a PCI enumerator) which is implemented as driver. This type of

driver has the `drv_probe` routine which enumerates devices residing on the bus and creates device nodes. This type of driver obviously does not have the `drv_init` routine.

`drv_unload` is called by the driver registry module (more exactly by the `svDriverUnregister` routine) when somebody wishes to unload the driver code from the system. The purpose of `drv_unload` is to check that the device driver code is not currently being used. The driver must check, for each driver instance, whether it is locked in the device registry.

In case of success, all device instances are removed from the device registry and `K_OK` is returned. Otherwise, the device entries are unchanged in the device registry and `K_EBUSY` is returned.

The `drv_unload` routine is optional. In cases when `drv_unload` is not provided by the driver, the `drv_unload` field must be set to `NULL`. Note that, in this case, the driver code cannot be unloaded.

`svDriverLookupFirst` returns the first driver entry in the registry. When the registry is not empty a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

The driver entry lock is not exclusive. It is simply a counter which prevents the driver component from being unloaded when it is still being used by a bus/nexus driver for probing or initialization.

`svDriverLookupNext` returns the next driver entry in the registry. The current entry is specified by the `drv_id` argument. The current entry must be locked by a previously called `svDriverLookupFirst` or `svDriverLookupNext`. If the current entry is not the last one in the registry a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the next driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

`svDriverRelease` releases the lock of the driver entry specified by the `drv_id` argument.

`svDriverEntry` returns a pointer to the driver entry structure specified by the `drv_id` argument. The device entry being accessed must have been previously locked using `svDriverLookupFirst` or `svDriverLookupNext`. Note that the device entry structure is read-only.

svDriverCap returns a pointer to the driver actor capability. The driver entry is specified by the *drv_id* argument. The driver entry being accessed must have been previously locked using svDriverLookupFirst or svDriverLookupNext . Note that if a given driver entry is registered by a built-in driver, a NULL pointer is returned. In other words, the driver actor capability makes sense only for dynamically loaded drivers. The actor capability may be used by an application in order to delete the driver actor once the driver entry is unregistered. Note that the driver capability structure is read-only.

svDriverUnregister tries to remove the driver entry specified by the *dev_id* argument from the driver registry. The device entry being removed must have been previously locked using svDriverLookupFirst or svDriverLookupNext .

In case of success, K_OK is returned, otherwise K_EBUSY is returned. The K_EBUSY result means that either the driver entry is locked in the driver registry (that is, a static driver routine is currently being used by a bus/nexus driver) or an instance of the device driver is locked in the device registry (that is, there is a driver instance which is currently being used by a driver client). Note that when K_EBUSY is returned, the driver entry remains locked in the registry and should be unlocked explicitly by svDriverRelease .

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svDriverRegister	+	+	-	+
svDriverLookupFirst	+	+	-	+
svDriverLookupNext	+	+	-	+
svDriverRelease	+	+	-	+
svDriverEntry	+	+	-	-
svDriverCap	+	+	-	-
svDriverUnregister	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

dreeNodeRoot(9DKI) , svDkiThreadCall(9DKI)

NAME	svDriverRegister, svDriverLookupFirst, svDriverLookupNext, svDriverRelease, svDriverEntry, svDriverCap, svDriverUnregister – driver registry operations
SYNOPSIS	<pre> #include <dki/dki.h> KnError svDriverRegister(DrvRegEntry * drv_entry); DrvRegId svDriverLookupFirst(void); DrvRegId svDriverLookupNext(DrvRegId drv_id); void svDriverRelease(DrvRegId drv_id); DrvRegEntry * svDriverEntry(DrvRegId drv_id); KnCap * svDriverCap(DrvRegId drv_id); KnError svDriverUnregister(DrvRegId drv_id); typedef struct { char* drv_name; char* drv_info; char* bus_class; int bus_version; void (*drv_probe) (DevNode bus_node, void* bus_ops, void* bus_id); void (*drv_bind) (DevNode bus_node); void (*drv_init) (DevNode dev_node, void* bus_ops, void* bus_id); KnError (*drv_unload) (); } DrvRegEntry; </pre>

FEATURES

DKI

DESCRIPTION

The `driver_registry` module implements a data base of drivers which have registered within the system. The `driver_registry` data base is populated by drivers which perform self-registration (using `svDriverRegister`) at driver initialization time.

EXTENDED DESCRIPTION

`svDriverRegister` adds the driver entry to the driver registry. It returns `K_OK` in case of success, otherwise `K_ENOMEM` is returned to indicate that the system is out of memory. The `drv_entry` argument points to the `DrvRegEntry` structure specifying driver properties and static driver routines. The `DrvRegEntry` structure is described below.

<code>drv_name</code>	Points to a string specifying the driver name. For example, the driver file name.
<code>drv_info</code>	Points to a string specifying extra information about the driver component, such as version or author.
<code>bus_class</code>	Points to a string specifying the class of the parent driver API required for the driver, such as "pci".
<code>bus_version</code>	Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code> , the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.
<code>drv_probe</code>	Points to a static driver routine which performs the device enumeration/probing on the bus.
<code>drv_bind</code>	Points to a static driver routine which performs the driver-to-device binding.
<code>drv_init</code>	Points to a static driver routine which clones an instance of the driver for the given device.
<code>drv_unload</code>	Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.

The `bus_class` and `bus_version` fields specify a parent bus/nexus API required for the driver. This type of API is provided to the driver as a structure of indirect

functions implementing the API service routines. For instance, the *"pci"* bus API is specified by the `PciBusOps` structure, the *"isa"* bus API is specified by `IsaBusOps`. When a bus driver invokes the `drv_probe` or `drv_init` driver's routine, it provides a pointer to the structure of the bus service routines. The structure type corresponds to the `bus_class` field value.

Note that a bus driver may provide multiple APIs. A typical example is a PCI bus driver providing the common and PCI bus APIs. The common bus driver API is named by *"bus"* and specified by the `BusOps` structure. The PCI bus driver API is named by *"pci"* and specified by the `PciBusOps` structure. `BusOps` actually provides a subset of services provided by `PciBusOps`. This type of bus driver is able to support drivers which use either common (*"bus"*) or PCI (*"pci"*) parent bus interfaces. When the bus driver invokes a child driver, it gives a pointer to either the `BusOps` or `PciBusOps` structure depending on the `bus_class` specified in the child driver registry entry.

Note also that a `PROP_DRIVER` may be adaptive to the parent bus API. In other words, such a driver is able to run on top of a number of different buses (for example, *"pci"* and *"isa"*). Typically, this type of driver is composed of two parts: bus class specific and bus class independent.

The bus class specific part of the driver code mainly deals with the device probing and initialization. In addition, it provides an abstraction layer in order to hide the bus class dependencies from the bus class independent part of the driver code. This kind of multi-bus driver should be registered multiple times in the driver registry. Each entry specifies a given bus class API (via the `bus_class` field) on top of which the driver may run. In order to determine to which bus class the driver is applied, the `drv_probe`, `drv_bind` and `drv_init` routines have to be entry specific (that is, bus class API specific). Under these conditions, when a given driver's routine is invoked by a bus driver, the driver detects the bus class to which it is applied and casts the `bus_ops` argument to the appropriate structure (for example, `PciBusOps`).

The `drv_probe` routine is invoked by a bus/nexus driver when the `bus_class` specified in the registry entry matches the bus/nexus driver class. `drv_probe` is called with three arguments.

The `dev_node` argument specifies the parent device node.

The `bus_ops / bus_id` pair specifies the parent device driver.

`bus_ops` points to a structure of service routines implementing a bus driver API. This structure is bus class specific and corresponds to the `bus_class` specified by the driver registry entry.

bus_id is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The `drv_probe` routine is optional. In cases where the `drv_probe` routine is not provided by the driver, the *drv_probe* field must be set to `NULL`.

The `drv_bind` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class.

`drv_bind` is called with one argument, *dev_node*, which specifies a given device node. The `drv_bind` routine enables the driver to perform a driver-to-device binding. Typically, the driver examines properties attached to the device node in order to determine the type of device at the node and to check whether the device may be serviced by the driver. If the check is positive, the driver binds the driver to the device node by attaching a `PROP_DRIVER` property to the device node. The property value specifies the driver name. The parent bus driver uses such a property to determine the name of the driver servicing the device. In other words, via the `PROP_DRIVER` property, the child driver gives its name to the parent bus driver, asking it to invoke the `drv_init` routine on that device. Note that, if a "driver" property is already present in the device node, then the `drv_bind` routine can not continue; `drv_bind` should not override existing driver-to-device binding.

The `drv_bind` routine is optional. If the `drv_bind` routine is not provided by the driver, the *drv_bind* field must be set to `NULL`.

The `drv_init` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class and a given driver is bound to a given device node.

`drv_init` is called with three arguments. The *dev_node* argument specifies the device node for which a device driver instance should be created. The *bus_ops* / *bus_id* pair specifies the parent device driver. *bus_ops* points to a structure of service routines implementing a bus driver API.

This structure is bus class specific and corresponds to the *bus_class* specified by the driver registry entry. *bus_id* is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The purpose of `drv_init` is to create an instance of the driver servicing the given device, to perform a device hardware initialization and to register the driver instance in the device registry.

Typically, `drv_init` would read the device and bus node properties in order to obtain the assigned bus resources and tunable parameters related to the bus/device. The `drv_init` routine is optional. In a case when the `drv_init` routine is not provided by the driver, the *drv_init* field must be set to `NULL`. A typical example of a probe-only driver is a self-identifying bus enumerator (for example, a PCI enumerator) which is implemented as driver. This type of

driver has the `drv_probe` routine which enumerates devices residing on the bus and creates device nodes. This type of driver obviously does not have the `drv_init` routine.

`drv_unload` is called by the driver registry module (more exactly by the `svDriverUnregister` routine) when somebody wishes to unload the driver code from the system. The purpose of `drv_unload` is to check that the device driver code is not currently being used. The driver must check, for each driver instance, whether it is locked in the device registry.

In case of success, all device instances are removed from the device registry and `K_OK` is returned. Otherwise, the device entries are unchanged in the device registry and `K_EBUSY` is returned.

The `drv_unload` routine is optional. In cases when `drv_unload` is not provided by the driver, the `drv_unload` field must be set to `NULL`. Note that, in this case, the driver code cannot be unloaded.

`svDriverLookupFirst` returns the first driver entry in the registry. When the registry is not empty a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

The driver entry lock is not exclusive. It is simply a counter which prevents the driver component from being unloaded when it is still being used by a bus/nexus driver for probing or initialization.

`svDriverLookupNext` returns the next driver entry in the registry. The current entry is specified by the `drv_id` argument. The current entry must be locked by a previously called `svDriverLookupFirst` or `svDriverLookupNext`. If the current entry is not the last one in the registry a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the next driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

`svDriverRelease` releases the lock of the driver entry specified by the `drv_id` argument.

`svDriverEntry` returns a pointer to the driver entry structure specified by the `drv_id` argument. The device entry being accessed must have been previously locked using `svDriverLookupFirst` or `svDriverLookupNext`. Note that the device entry structure is read-only.

svDriverCap returns a pointer to the driver actor capability. The driver entry is specified by the *drv_id* argument. The driver entry being accessed must have been previously locked using svDriverLookupFirst or svDriverLookupNext . Note that if a given driver entry is registered by a built-in driver, a NULL pointer is returned. In other words, the driver actor capability makes sense only for dynamically loaded drivers. The actor capability may be used by an application in order to delete the driver actor once the driver entry is unregistered. Note that the driver capability structure is read-only.

svDriverUnregister tries to remove the driver entry specified by the *dev_id* argument from the driver registry. The device entry being removed must have been previously locked using svDriverLookupFirst or svDriverLookupNext .

In case of success, K_OK is returned, otherwise K_EBUSY is returned. The K_EBUSY result means that either the driver entry is locked in the driver registry (that is, a static driver routine is currently being used by a bus/nexus driver) or an instance of the device driver is locked in the device registry (that is, there is a driver instance which is currently being used by a driver client). Note that when K_EBUSY is returned, the driver entry remains locked in the registry and should be unlocked explicitly by svDriverRelease .

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svDriverRegister	+	+	-	+
svDriverLookupFirst	+	+	-	+
svDriverLookupNext	+	+	-	+
svDriverRelease	+	+	-	+
svDriverEntry	+	+	-	-
svDriverCap	+	+	-	-
svDriverUnregister	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

dreeNodeRoot(9DKI) , svDkiThreadCall(9DKI)

NAME	svDriverRegister, svDriverLookupFirst, svDriverLookupNext, svDriverRelease, svDriverEntry, svDriverCap, svDriverUnregister – driver registry operations
SYNOPSIS	<pre> #include <dki/dki.h> KnError svDriverRegister(DrvRegEntry * drv_entry); DrvRegId svDriverLookupFirst(void); DrvRegId svDriverLookupNext(DrvRegId drv_id); void svDriverRelease(DrvRegId drv_id); DrvRegEntry * svDriverEntry(DrvRegId drv_id); KnCap * svDriverCap(DrvRegId drv_id); KnError svDriverUnregister(DrvRegId drv_id); typedef struct { char* drv_name; char* drv_info; char* bus_class; int bus_version; void (*drv_probe) (DevNode bus_node, void* bus_ops, void* bus_id); void (*drv_bind) (DevNode bus_node); void (*drv_init) (DevNode dev_node, void* bus_ops, void* bus_id); KnError (*drv_unload) (); } DrvRegEntry; </pre>

FEATURES

DKI

DESCRIPTION

The `driver_registry` module implements a data base of drivers which have registered within the system. The `driver_registry` data base is populated by drivers which perform self-registration (using `svDriverRegister`) at driver initialization time.

EXTENDED DESCRIPTION

`svDriverRegister` adds the driver entry to the driver registry. It returns `K_OK` in case of success, otherwise `K_ENOMEM` is returned to indicate that the system is out of memory. The `drv_entry` argument points to the `DrvRegEntry` structure specifying driver properties and static driver routines. The `DrvRegEntry` structure is described below.

<code>drv_name</code>	Points to a string specifying the driver name. For example, the driver file name.
<code>drv_info</code>	Points to a string specifying extra information about the driver component, such as version or author.
<code>bus_class</code>	Points to a string specifying the class of the parent driver API required for the driver, such as "pci".
<code>bus_version</code>	Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code> , the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.
<code>drv_probe</code>	Points to a static driver routine which performs the device enumeration/probing on the bus.
<code>drv_bind</code>	Points to a static driver routine which performs the driver-to-device binding.
<code>drv_init</code>	Points to a static driver routine which clones an instance of the driver for the given device.
<code>drv_unload</code>	Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.

The `bus_class` and `bus_version` fields specify a parent bus/nexus API required for the driver. This type of API is provided to the driver as a structure of indirect

functions implementing the API service routines. For instance, the *"pci"* bus API is specified by the `PciBusOps` structure, the *"isa"* bus API is specified by `IsaBusOps`. When a bus driver invokes the `drv_probe` or `drv_init` driver's routine, it provides a pointer to the structure of the bus service routines. The structure type corresponds to the `bus_class` field value.

Note that a bus driver may provide multiple APIs. A typical example is a PCI bus driver providing the common and PCI bus APIs. The common bus driver API is named by *"bus"* and specified by the `BusOps` structure. The PCI bus driver API is named by *"pci"* and specified by the `PciBusOps` structure. `BusOps` actually provides a subset of services provided by `PciBusOps`. This type of bus driver is able to support drivers which use either common (*"bus"*) or PCI (*"pci"*) parent bus interfaces. When the bus driver invokes a child driver, it gives a pointer to either the `BusOps` or `PciBusOps` structure depending on the `bus_class` specified in the child driver registry entry.

Note also that a `PROP_DRIVER` may be adaptive to the parent bus API. In other words, such a driver is able to run on top of a number of different buses (for example, *"pci"* and *"isa"*). Typically, this type of driver is composed of two parts: bus class specific and bus class independent.

The bus class specific part of the driver code mainly deals with the device probing and initialization. In addition, it provides an abstraction layer in order to hide the bus class dependencies from the bus class independent part of the driver code. This kind of multi-bus driver should be registered multiple times in the driver registry. Each entry specifies a given bus class API (via the `bus_class` field) on top of which the driver may run. In order to determine to which bus class the driver is applied, the `drv_probe`, `drv_bind` and `drv_init` routines have to be entry specific (that is, bus class API specific). Under these conditions, when a given driver's routine is invoked by a bus driver, the driver detects the bus class to which it is applied and casts the `bus_ops` argument to the appropriate structure (for example, `PciBusOps`).

The `drv_probe` routine is invoked by a bus/nexus driver when the `bus_class` specified in the registry entry matches the bus/nexus driver class. `drv_probe` is called with three arguments.

The `dev_node` argument specifies the parent device node.

The `bus_ops / bus_id` pair specifies the parent device driver.

`bus_ops` points to a structure of service routines implementing a bus driver API. This structure is bus class specific and corresponds to the `bus_class` specified by the driver registry entry.

bus_id is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The `drv_probe` routine is optional. In cases where the `drv_probe` routine is not provided by the driver, the *drv_probe* field must be set to `NULL`.

The `drv_bind` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class.

`drv_bind` is called with one argument, *dev_node*, which specifies a given device node. The `drv_bind` routine enables the driver to perform a driver-to-device binding. Typically, the driver examines properties attached to the device node in order to determine the type of device at the node and to check whether the device may be serviced by the driver. If the check is positive, the driver binds the driver to the device node by attaching a `PROP_DRIVER` property to the device node. The property value specifies the driver name. The parent bus driver uses such a property to determine the name of the driver servicing the device. In other words, via the `PROP_DRIVER` property, the child driver gives its name to the parent bus driver, asking it to invoke the `drv_init` routine on that device. Note that, if a "driver" property is already present in the device node, then the `drv_bind` routine can not continue; `drv_bind` should not override existing driver-to-device binding.

The `drv_bind` routine is optional. If the `drv_bind` routine is not provided by the driver, the *drv_bind* field must be set to `NULL`.

The `drv_init` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class and a given driver is bound to a given device node.

`drv_init` is called with three arguments. The *dev_node* argument specifies the device node for which a device driver instance should be created. The *bus_ops* / *bus_id* pair specifies the parent device driver. *bus_ops* points to a structure of service routines implementing a bus driver API.

This structure is bus class specific and corresponds to the *bus_class* specified by the driver registry entry. *bus_id* is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The purpose of `drv_init` is to create an instance of the driver servicing the given device, to perform a device hardware initialization and to register the driver instance in the device registry.

Typically, `drv_init` would read the device and bus node properties in order to obtain the assigned bus resources and tunable parameters related to the bus/device. The `drv_init` routine is optional. In a case when the `drv_init` routine is not provided by the driver, the *drv_init* field must be set to `NULL`. A typical example of a probe-only driver is a self-identifying bus enumerator (for example, a PCI enumerator) which is implemented as driver. This type of

driver has the `drv_probe` routine which enumerates devices residing on the bus and creates device nodes. This type of driver obviously does not have the `drv_init` routine.

`drv_unload` is called by the driver registry module (more exactly by the `svDriverUnregister` routine) when somebody wishes to unload the driver code from the system. The purpose of `drv_unload` is to check that the device driver code is not currently being used. The driver must check, for each driver instance, whether it is locked in the device registry.

In case of success, all device instances are removed from the device registry and `K_OK` is returned. Otherwise, the device entries are unchanged in the device registry and `K_EBUSY` is returned.

The `drv_unload` routine is optional. In cases when `drv_unload` is not provided by the driver, the `drv_unload` field must be set to `NULL`. Note that, in this case, the driver code cannot be unloaded.

`svDriverLookupFirst` returns the first driver entry in the registry. When the registry is not empty a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

The driver entry lock is not exclusive. It is simply a counter which prevents the driver component from being unloaded when it is still being used by a bus/nexus driver for probing or initialization.

`svDriverLookupNext` returns the next driver entry in the registry. The current entry is specified by the `drv_id` argument. The current entry must be locked by a previously called `svDriverLookupFirst` or `svDriverLookupNext`. If the current entry is not the last one in the registry a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the next driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

`svDriverRelease` releases the lock of the driver entry specified by the `drv_id` argument.

`svDriverEntry` returns a pointer to the driver entry structure specified by the `drv_id` argument. The device entry being accessed must have been previously locked using `svDriverLookupFirst` or `svDriverLookupNext`. Note that the device entry structure is read-only.

svDriverCap returns a pointer to the driver actor capability. The driver entry is specified by the *drv_id* argument. The driver entry being accessed must have been previously locked using svDriverLookupFirst or svDriverLookupNext . Note that if a given driver entry is registered by a built-in driver, a NULL pointer is returned. In other words, the driver actor capability makes sense only for dynamically loaded drivers. The actor capability may be used by an application in order to delete the driver actor once the driver entry is unregistered. Note that the driver capability structure is read-only.

svDriverUnregister tries to remove the driver entry specified by the *dev_id* argument from the driver registry. The device entry being removed must have been previously locked using svDriverLookupFirst or svDriverLookupNext .

In case of success, K_OK is returned, otherwise K_EBUSY is returned. The K_EBUSY result means that either the driver entry is locked in the driver registry (that is, a static driver routine is currently being used by a bus/nexus driver) or an instance of the device driver is locked in the device registry (that is, there is a driver instance which is currently being used by a driver client). Note that when K_EBUSY is returned, the driver entry remains locked in the registry and should be unlocked explicitly by svDriverRelease .

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svDriverRegister	+	+	-	+
svDriverLookupFirst	+	+	-	+
svDriverLookupNext	+	+	-	+
svDriverRelease	+	+	-	+
svDriverEntry	+	+	-	-
svDriverCap	+	+	-	-
svDriverUnregister	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

dreeNodeRoot(9DKI) , svDkiThreadCall(9DKI)

NAME	svDriverRegister, svDriverLookupFirst, svDriverLookupNext, svDriverRelease, svDriverEntry, svDriverCap, svDriverUnregister – driver registry operations
SYNOPSIS	<pre> #include <dki/dki.h> KnError svDriverRegister(DrvRegEntry * drv_entry); DrvRegId svDriverLookupFirst(void); DrvRegId svDriverLookupNext(DrvRegId drv_id); void svDriverRelease(DrvRegId drv_id); DrvRegEntry * svDriverEntry(DrvRegId drv_id); KnCap * svDriverCap(DrvRegId drv_id); KnError svDriverUnregister(DrvRegId drv_id); typedef struct { char* drv_name; char* drv_info; char* bus_class; int bus_version; void (*drv_probe) (DevNode bus_node, void* bus_ops, void* bus_id); void (*drv_bind) (DevNode bus_node); void (*drv_init) (DevNode dev_node, void* bus_ops, void* bus_id); KnError (*drv_unload) (); } DrvRegEntry; </pre>

FEATURES	DKI																
DESCRIPTION	The <code>driver_registry</code> module implements a data base of drivers which have registered within the system. The <code>driver_registry</code> data base is populated by drivers which perform self-registration (using <code>svDriverRegister</code>) at driver initialization time.																
EXTENDED DESCRIPTION	<p><code>svDriverRegister</code> adds the driver entry to the driver registry. It returns <code>K_OK</code> in case of success, otherwise <code>K_ENOMEM</code> is returned to indicate that the system is out of memory. The <code>drv_entry</code> argument points to the <code>DrvRegEntry</code> structure specifying driver properties and static driver routines. The <code>DrvRegEntry</code> structure is described below.</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>drv_name</code></td> <td>Points to a string specifying the driver name. For example, the driver file name.</td> </tr> <tr> <td><code>drv_info</code></td> <td>Points to a string specifying extra information about the driver component, such as version or author.</td> </tr> <tr> <td><code>bus_class</code></td> <td>Points to a string specifying the class of the parent driver API required for the driver, such as "pci".</td> </tr> <tr> <td><code>bus_version</code></td> <td>Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code>, the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.</td> </tr> <tr> <td><code>drv_probe</code></td> <td>Points to a static driver routine which performs the device enumeration/probing on the bus.</td> </tr> <tr> <td><code>drv_bind</code></td> <td>Points to a static driver routine which performs the driver-to-device binding.</td> </tr> <tr> <td><code>drv_init</code></td> <td>Points to a static driver routine which clones an instance of the driver for the given device.</td> </tr> <tr> <td><code>drv_unload</code></td> <td>Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.</td> </tr> </table> <p>The <code>bus_class</code> and <code>bus_version</code> fields specify a parent bus/nexus API required for the driver. This type of API is provided to the driver as a structure of indirect</p>	<code>drv_name</code>	Points to a string specifying the driver name. For example, the driver file name.	<code>drv_info</code>	Points to a string specifying extra information about the driver component, such as version or author.	<code>bus_class</code>	Points to a string specifying the class of the parent driver API required for the driver, such as "pci".	<code>bus_version</code>	Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code> , the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.	<code>drv_probe</code>	Points to a static driver routine which performs the device enumeration/probing on the bus.	<code>drv_bind</code>	Points to a static driver routine which performs the driver-to-device binding.	<code>drv_init</code>	Points to a static driver routine which clones an instance of the driver for the given device.	<code>drv_unload</code>	Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.
<code>drv_name</code>	Points to a string specifying the driver name. For example, the driver file name.																
<code>drv_info</code>	Points to a string specifying extra information about the driver component, such as version or author.																
<code>bus_class</code>	Points to a string specifying the class of the parent driver API required for the driver, such as "pci".																
<code>bus_version</code>	Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code> , the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.																
<code>drv_probe</code>	Points to a static driver routine which performs the device enumeration/probing on the bus.																
<code>drv_bind</code>	Points to a static driver routine which performs the driver-to-device binding.																
<code>drv_init</code>	Points to a static driver routine which clones an instance of the driver for the given device.																
<code>drv_unload</code>	Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.																

functions implementing the API service routines. For instance, the *"pci"* bus API is specified by the `PciBusOps` structure, the *"isa"* bus API is specified by `IsaBusOps`. When a bus driver invokes the `drv_probe` or `drv_init` driver's routine, it provides a pointer to the structure of the bus service routines. The structure type corresponds to the `bus_class` field value.

Note that a bus driver may provide multiple APIs. A typical example is a PCI bus driver providing the common and PCI bus APIs. The common bus driver API is named by *"bus"* and specified by the `BusOps` structure. The PCI bus driver API is named by *"pci"* and specified by the `PciBusOps` structure. `BusOps` actually, provides a subset of services provided by `PciBusOps`. This type of bus driver is able to support drivers which use either common (*"bus"*) or PCI (*"pci"*) parent bus interfaces. When the bus driver invokes a child driver, it gives a pointer to either the `BusOps` or `PciBusOps` structure depending on the `bus_class` specified in the child driver registry entry.

Note also that a `PROP_DRIVER` may be adaptive to the parent bus API. In other words, such a driver is able to run on top of a number of different buses (for example, *"pci"* and *"isa"*). Typically, this type of driver is composed of two parts: bus class specific and bus class independent.

The bus class specific part of the driver code mainly deals with the device probing and initialization. In addition, it provides an abstraction layer in order to hide the bus class dependencies from the bus class independent part of the driver code. This kind of multi-bus driver should be registered multiple times in the driver registry. Each entry specifies a given bus class API (via the `bus_class` field) on top of which the driver may run. In order to determine to which bus class the driver is applied, the `drv_probe`, `drv_bind` and `drv_init` routines have to be entry specific (that is, bus class API specific). Under these conditions, when a given driver's routine is invoked by a bus driver, the driver detects the bus class to which it is applied and casts the `bus_ops` argument to the appropriate structure (for example, `PciBusOps`).

The `drv_probe` routine is invoked by a bus/nexus driver when the `bus_class` specified in the registry entry matches the bus/nexus driver class. `drv_probe` is called with three arguments.

The `dev_node` argument specifies the parent device node.

The `bus_ops / bus_id` pair specifies the parent device driver.

`bus_ops` points to a structure of service routines implementing a bus driver API. This structure is bus class specific and corresponds to the `bus_class` specified by the driver registry entry.

bus_id is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The `drv_probe` routine is optional. In cases where the `drv_probe` routine is not provided by the driver, the *drv_probe* field must be set to `NULL`.

The `drv_bind` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class.

`drv_bind` is called with one argument, *dev_node*, which specifies a given device node. The `drv_bind` routine enables the driver to perform a driver-to-device binding. Typically, the driver examines properties attached to the device node in order to determine the type of device at the node and to check whether the device may be serviced by the driver. If the check is positive, the driver binds the driver to the device node by attaching a `PROP_DRIVER` property to the device node. The property value specifies the driver name. The parent bus driver uses such a property to determine the name of the driver servicing the device. In other words, via the `PROP_DRIVER` property, the child driver gives its name to the parent bus driver, asking it to invoke the `drv_init` routine on that device. Note that, if a "driver" property is already present in the device node, then the `drv_bind` routine can not continue; `drv_bind` should not override existing driver-to-device binding.

The `drv_bind` routine is optional. If the `drv_bind` routine is not provided by the driver, the *drv_bind* field must be set to `NULL`.

The `drv_init` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class and a given driver is bound to a given device node.

`drv_init` is called with three arguments. The *dev_node* argument specifies the device node for which a device driver instance should be created. The *bus_ops* / *bus_id* pair specifies the parent device driver. *bus_ops* points to a structure of service routines implementing a bus driver API.

This structure is bus class specific and corresponds to the *bus_class* specified by the driver registry entry. *bus_id* is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The purpose of `drv_init` is to create an instance of the driver servicing the given device, to perform a device hardware initialization and to register the driver instance in the device registry.

Typically, `drv_init` would read the device and bus node properties in order to obtain the assigned bus resources and tunable parameters related to the bus/device. The `drv_init` routine is optional. In a case when the `drv_init` routine is not provided by the driver, the *drv_init* field must be set to `NULL`. A typical example of a probe-only driver is a self-identifying bus enumerator (for example, a PCI enumerator) which is implemented as driver. This type of

driver has the `drv_probe` routine which enumerates devices residing on the bus and creates device nodes. This type of driver obviously does not have the `drv_init` routine.

`drv_unload` is called by the driver registry module (more exactly by the `svDriverUnregister` routine) when somebody wishes to unload the driver code from the system. The purpose of `drv_unload` is to check that the device driver code is not currently being used. The driver must check, for each driver instance, whether it is locked in the device registry.

In case of success, all device instances are removed from the device registry and `K_OK` is returned. Otherwise, the device entries are unchanged in the device registry and `K_EBUSY` is returned.

The `drv_unload` routine is optional. In cases when `drv_unload` is not provided by the driver, the `drv_unload` field must be set to `NULL`. Note that, in this case, the driver code cannot be unloaded.

`svDriverLookupFirst` returns the first driver entry in the registry. When the registry is not empty a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

The driver entry lock is not exclusive. It is simply a counter which prevents the driver component from being unloaded when it is still being used by a bus/nexus driver for probing or initialization.

`svDriverLookupNext` returns the next driver entry in the registry. The current entry is specified by the `drv_id` argument. The current entry must be locked by a previously called `svDriverLookupFirst` or `svDriverLookupNext`. If the current entry is not the last one in the registry a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the next driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

`svDriverRelease` releases the lock of the driver entry specified by the `drv_id` argument.

`svDriverEntry` returns a pointer to the driver entry structure specified by the `drv_id` argument. The device entry being accessed must have been previously locked using `svDriverLookupFirst` or `svDriverLookupNext`. Note that the device entry structure is read-only.

svDriverCap returns a pointer to the driver actor capability. The driver entry is specified by the *drv_id* argument. The driver entry being accessed must have been previously locked using svDriverLookupFirst or svDriverLookupNext . Note that if a given driver entry is registered by a built-in driver, a NULL pointer is returned. In other words, the driver actor capability makes sense only for dynamically loaded drivers. The actor capability may be used by an application in order to delete the driver actor once the driver entry is unregistered. Note that the driver capability structure is read-only.

svDriverUnregister tries to remove the driver entry specified by the *dev_id* argument from the driver registry. The device entry being removed must have been previously locked using svDriverLookupFirst or svDriverLookupNext .

In case of success, K_OK is returned, otherwise K_EBUSY is returned. The K_EBUSY result means that either the driver entry is locked in the driver registry (that is, a static driver routine is currently being used by a bus/nexus driver) or an instance of the device driver is locked in the device registry (that is, there is a driver instance which is currently being used by a driver client). Note that when K_EBUSY is returned, the driver entry remains locked in the registry and should be unlocked explicitly by svDriverRelease .

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svDriverRegister	+	+	-	+
svDriverLookupFirst	+	+	-	+
svDriverLookupNext	+	+	-	+
svDriverRelease	+	+	-	+
svDriverEntry	+	+	-	-
svDriverCap	+	+	-	-
svDriverUnregister	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

dreeNodeRoot(9DKI) , svDkiThreadCall(9DKI)

NAME	svDriverRegister, svDriverLookupFirst, svDriverLookupNext, svDriverRelease, svDriverEntry, svDriverCap, svDriverUnregister – driver registry operations
SYNOPSIS	<pre> #include <dki/dki.h> KnError svDriverRegister(DrvRegEntry * drv_entry); DrvRegId svDriverLookupFirst(void); DrvRegId svDriverLookupNext(DrvRegId drv_id); void svDriverRelease(DrvRegId drv_id); DrvRegEntry * svDriverEntry(DrvRegId drv_id); KnCap * svDriverCap(DrvRegId drv_id); KnError svDriverUnregister(DrvRegId drv_id); typedef struct { char* drv_name; char* drv_info; char* bus_class; int bus_version; void (*drv_probe) (DevNode bus_node, void* bus_ops, void* bus_id); void (*drv_bind) (DevNode bus_node); void (*drv_init) (DevNode dev_node, void* bus_ops, void* bus_id); KnError (*drv_unload) (); } DrvRegEntry; </pre>

FEATURES

DKI

DESCRIPTION

The `driver_registry` module implements a data base of drivers which have registered within the system. The `driver_registry` data base is populated by drivers which perform self-registration (using `svDriverRegister`) at driver initialization time.

EXTENDED DESCRIPTION

`svDriverRegister` adds the driver entry to the driver registry. It returns `K_OK` in case of success, otherwise `K_ENOMEM` is returned to indicate that the system is out of memory. The `drv_entry` argument points to the `DrvRegEntry` structure specifying driver properties and static driver routines. The `DrvRegEntry` structure is described below.

<code>drv_name</code>	Points to a string specifying the driver name. For example, the driver file name.
<code>drv_info</code>	Points to a string specifying extra information about the driver component, such as version or author.
<code>bus_class</code>	Points to a string specifying the class of the parent driver API required for the driver, such as "pci".
<code>bus_version</code>	Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code> , the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.
<code>drv_probe</code>	Points to a static driver routine which performs the device enumeration/probing on the bus.
<code>drv_bind</code>	Points to a static driver routine which performs the driver-to-device binding.
<code>drv_init</code>	Points to a static driver routine which clones an instance of the driver for the given device.
<code>drv_unload</code>	Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.

The `bus_class` and `bus_version` fields specify a parent bus/nexus API required for the driver. This type of API is provided to the driver as a structure of indirect

functions implementing the API service routines. For instance, the *"pci"* bus API is specified by the `PciBusOps` structure, the *"isa"* bus API is specified by `IsaBusOps`. When a bus driver invokes the `drv_probe` or `drv_init` driver's routine, it provides a pointer to the structure of the bus service routines. The structure type corresponds to the `bus_class` field value.

Note that a bus driver may provide multiple APIs. A typical example is a PCI bus driver providing the common and PCI bus APIs. The common bus driver API is named by *"bus"* and specified by the `BusOps` structure. The PCI bus driver API is named by *"pci"* and specified by the `PciBusOps` structure. `BusOps` actually, provides a subset of services provided by `PciBusOps`. This type of bus driver is able to support drivers which use either common (*"bus"*) or PCI (*"pci"*) parent bus interfaces. When the bus driver invokes a child driver, it gives a pointer to either the `BusOps` or `PciBusOps` structure depending on the `bus_class` specified in the child driver registry entry.

Note also that a `PROP_DRIVER` may be adaptive to the parent bus API. In other words, such a driver is able to run on top of a number of different buses (for example, *"pci"* and *"isa"*). Typically, this type of driver is composed of two parts: bus class specific and bus class independent.

The bus class specific part of the driver code mainly deals with the device probing and initialization. In addition, it provides an abstraction layer in order to hide the bus class dependencies from the bus class independent part of the driver code. This kind of multi-bus driver should be registered multiple times in the driver registry. Each entry specifies a given bus class API (via the `bus_class` field) on top of which the driver may run. In order to determine to which bus class the driver is applied, the `drv_probe`, `drv_bind` and `drv_init` routines have to be entry specific (that is, bus class API specific). Under these conditions, when a given driver's routine is invoked by a bus driver, the driver detects the bus class to which it is applied and casts the `bus_ops` argument to the appropriate structure (for example, `PciBusOps`).

The `drv_probe` routine is invoked by a bus/nexus driver when the `bus_class` specified in the registry entry matches the bus/nexus driver class. `drv_probe` is called with three arguments.

The `dev_node` argument specifies the parent device node.

The `bus_ops / bus_id` pair specifies the parent device driver.

`bus_ops` points to a structure of service routines implementing a bus driver API. This structure is bus class specific and corresponds to the `bus_class` specified by the driver registry entry.

bus_id is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The `drv_probe` routine is optional. In cases where the `drv_probe` routine is not provided by the driver, the *drv_probe* field must be set to `NULL`.

The `drv_bind` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class.

`drv_bind` is called with one argument, *dev_node*, which specifies a given device node. The `drv_bind` routine enables the driver to perform a driver-to-device binding. Typically, the driver examines properties attached to the device node in order to determine the type of device at the node and to check whether the device may be serviced by the driver. If the check is positive, the driver binds the driver to the device node by attaching a `PROP_DRIVER` property to the device node. The property value specifies the driver name. The parent bus driver uses such a property to determine the name of the driver servicing the device. In other words, via the `PROP_DRIVER` property, the child driver gives its name to the parent bus driver, asking it to invoke the `drv_init` routine on that device. Note that, if a "driver" property is already present in the device node, then the `drv_bind` routine can not continue; `drv_bind` should not override existing driver-to-device binding.

The `drv_bind` routine is optional. If the `drv_bind` routine is not provided by the driver, the *drv_bind* field must be set to `NULL`.

The `drv_init` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class and a given driver is bound to a given device node.

`drv_init` is called with three arguments. The *dev_node* argument specifies the device node for which a device driver instance should be created. The *bus_ops* / *bus_id* pair specifies the parent device driver. *bus_ops* points to a structure of service routines implementing a bus driver API.

This structure is bus class specific and corresponds to the *bus_class* specified by the driver registry entry. *bus_id* is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The purpose of `drv_init` is to create an instance of the driver servicing the given device, to perform a device hardware initialization and to register the driver instance in the device registry.

Typically, `drv_init` would read the device and bus node properties in order to obtain the assigned bus resources and tunable parameters related to the bus/device. The `drv_init` routine is optional. In a case when the `drv_init` routine is not provided by the driver, the *drv_init* field must be set to `NULL`. A typical example of a probe-only driver is a self-identifying bus enumerator (for example, a PCI enumerator) which is implemented as driver. This type of

driver has the `drv_probe` routine which enumerates devices residing on the bus and creates device nodes. This type of driver obviously does not have the `drv_init` routine.

`drv_unload` is called by the driver registry module (more exactly by the `svDriverUnregister` routine) when somebody wishes to unload the driver code from the system. The purpose of `drv_unload` is to check that the device driver code is not currently being used. The driver must check, for each driver instance, whether it is locked in the device registry.

In case of success, all device instances are removed from the device registry and `K_OK` is returned. Otherwise, the device entries are unchanged in the device registry and `K_EBUSY` is returned.

The `drv_unload` routine is optional. In cases when `drv_unload` is not provided by the driver, the `drv_unload` field must be set to `NULL`. Note that, in this case, the driver code cannot be unloaded.

`svDriverLookupFirst` returns the first driver entry in the registry. When the registry is not empty a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

The driver entry lock is not exclusive. It is simply a counter which prevents the driver component from being unloaded when it is still being used by a bus/nexus driver for probing or initialization.

`svDriverLookupNext` returns the next driver entry in the registry. The current entry is specified by the `drv_id` argument. The current entry must be locked by a previously called `svDriverLookupFirst` or `svDriverLookupNext`. If the current entry is not the last one in the registry a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the next driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

`svDriverRelease` releases the lock of the driver entry specified by the `drv_id` argument.

`svDriverEntry` returns a pointer to the driver entry structure specified by the `drv_id` argument. The device entry being accessed must have been previously locked using `svDriverLookupFirst` or `svDriverLookupNext`. Note that the device entry structure is read-only.

svDriverCap returns a pointer to the driver actor capability. The driver entry is specified by the *drv_id* argument. The driver entry being accessed must have been previously locked using svDriverLookupFirst or svDriverLookupNext . Note that if a given driver entry is registered by a built-in driver, a NULL pointer is returned. In other words, the driver actor capability makes sense only for dynamically loaded drivers. The actor capability may be used by an application in order to delete the driver actor once the driver entry is unregistered. Note that the driver capability structure is read-only.

svDriverUnregister tries to remove the driver entry specified by the *dev_id* argument from the driver registry. The device entry being removed must have been previously locked using svDriverLookupFirst or svDriverLookupNext .

In case of success, K_OK is returned, otherwise K_EBUSY is returned. The K_EBUSY result means that either the driver entry is locked in the driver registry (that is, a static driver routine is currently being used by a bus/nexus driver) or an instance of the device driver is locked in the device registry (that is, there is a driver instance which is currently being used by a driver client). Note that when K_EBUSY is returned, the driver entry remains locked in the registry and should be unlocked explicitly by svDriverRelease .

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svDriverRegister	+	+	-	+
svDriverLookupFirst	+	+	-	+
svDriverLookupNext	+	+	-	+
svDriverRelease	+	+	-	+
svDriverEntry	+	+	-	-
svDriverCap	+	+	-	-
svDriverUnregister	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

dreeNodeRoot(9DKI) , svDkiThreadCall(9DKI)

NAME	svDriverRegister, svDriverLookupFirst, svDriverLookupNext, svDriverRelease, svDriverEntry, svDriverCap, svDriverUnregister – driver registry operations
SYNOPSIS	<pre> #include <dki/dki.h> KnError svDriverRegister(DrvRegEntry * drv_entry); DrvRegId svDriverLookupFirst(void); DrvRegId svDriverLookupNext(DrvRegId drv_id); void svDriverRelease(DrvRegId drv_id); DrvRegEntry * svDriverEntry(DrvRegId drv_id); KnCap * svDriverCap(DrvRegId drv_id); KnError svDriverUnregister(DrvRegId drv_id); typedef struct { char* drv_name; char* drv_info; char* bus_class; int bus_version; void (*drv_probe) (DevNode bus_node, void* bus_ops, void* bus_id); void (*drv_bind) (DevNode bus_node); void (*drv_init) (DevNode dev_node, void* bus_ops, void* bus_id); KnError (*drv_unload) (); } DrvRegEntry; </pre>

FEATURES	DKI																
DESCRIPTION	The <code>driver_registry</code> module implements a data base of drivers which have registered within the system. The <code>driver_registry</code> data base is populated by drivers which perform self-registration (using <code>svDriverRegister</code>) at driver initialization time.																
EXTENDED DESCRIPTION	<p><code>svDriverRegister</code> adds the driver entry to the driver registry. It returns <code>K_OK</code> in case of success, otherwise <code>K_ENOMEM</code> is returned to indicate that the system is out of memory. The <code>drv_entry</code> argument points to the <code>DrvRegEntry</code> structure specifying driver properties and static driver routines. The <code>DrvRegEntry</code> structure is described below.</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>drv_name</code></td> <td>Points to a string specifying the driver name. For example, the driver file name.</td> </tr> <tr> <td><code>drv_info</code></td> <td>Points to a string specifying extra information about the driver component, such as version or author.</td> </tr> <tr> <td><code>bus_class</code></td> <td>Points to a string specifying the class of the parent driver API required for the driver, such as "pci".</td> </tr> <tr> <td><code>bus_version</code></td> <td>Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code>, the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.</td> </tr> <tr> <td><code>drv_probe</code></td> <td>Points to a static driver routine which performs the device enumeration/probing on the bus.</td> </tr> <tr> <td><code>drv_bind</code></td> <td>Points to a static driver routine which performs the driver-to-device binding.</td> </tr> <tr> <td><code>drv_init</code></td> <td>Points to a static driver routine which clones an instance of the driver for the given device.</td> </tr> <tr> <td><code>drv_unload</code></td> <td>Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.</td> </tr> </table> <p>The <code>bus_class</code> and <code>bus_version</code> fields specify a parent bus/nexus API required for the driver. This type of API is provided to the driver as a structure of indirect</p>	<code>drv_name</code>	Points to a string specifying the driver name. For example, the driver file name.	<code>drv_info</code>	Points to a string specifying extra information about the driver component, such as version or author.	<code>bus_class</code>	Points to a string specifying the class of the parent driver API required for the driver, such as "pci".	<code>bus_version</code>	Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code> , the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.	<code>drv_probe</code>	Points to a static driver routine which performs the device enumeration/probing on the bus.	<code>drv_bind</code>	Points to a static driver routine which performs the driver-to-device binding.	<code>drv_init</code>	Points to a static driver routine which clones an instance of the driver for the given device.	<code>drv_unload</code>	Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.
<code>drv_name</code>	Points to a string specifying the driver name. For example, the driver file name.																
<code>drv_info</code>	Points to a string specifying extra information about the driver component, such as version or author.																
<code>bus_class</code>	Points to a string specifying the class of the parent driver API required for the driver, such as "pci".																
<code>bus_version</code>	Specifies the minimum version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than <code>bus_version</code> , the corresponding driver component will never be called by this type of bus/nexus driver. In other words, the <code>drv_probe</code> and <code>drv_init</code> routines will never be invoked by this type of bus/nexus driver.																
<code>drv_probe</code>	Points to a static driver routine which performs the device enumeration/probing on the bus.																
<code>drv_bind</code>	Points to a static driver routine which performs the driver-to-device binding.																
<code>drv_init</code>	Points to a static driver routine which clones an instance of the driver for the given device.																
<code>drv_unload</code>	Points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.																

functions implementing the API service routines. For instance, the *"pci"* bus API is specified by the `PciBusOps` structure, the *"isa"* bus API is specified by `IsaBusOps`. When a bus driver invokes the `drv_probe` or `drv_init` driver's routine, it provides a pointer to the structure of the bus service routines. The structure type corresponds to the `bus_class` field value.

Note that a bus driver may provide multiple APIs. A typical example is a PCI bus driver providing the common and PCI bus APIs. The common bus driver API is named by *"bus"* and specified by the `BusOps` structure. The PCI bus driver API is named by *"pci"* and specified by the `PciBusOps` structure. `BusOps` actually provides a subset of services provided by `PciBusOps`. This type of bus driver is able to support drivers which use either common (*"bus"*) or PCI (*"pci"*) parent bus interfaces. When the bus driver invokes a child driver, it gives a pointer to either the `BusOps` or `PciBusOps` structure depending on the `bus_class` specified in the child driver registry entry.

Note also that a `PROP_DRIVER` may be adaptive to the parent bus API. In other words, such a driver is able to run on top of a number of different buses (for example, *"pci"* and *"isa"*). Typically, this type of driver is composed of two parts: bus class specific and bus class independent.

The bus class specific part of the driver code mainly deals with the device probing and initialization. In addition, it provides an abstraction layer in order to hide the bus class dependencies from the bus class independent part of the driver code. This kind of multi-bus driver should be registered multiple times in the driver registry. Each entry specifies a given bus class API (via the `bus_class` field) on top of which the driver may run. In order to determine to which bus class the driver is applied, the `drv_probe`, `drv_bind` and `drv_init` routines have to be entry specific (that is, bus class API specific). Under these conditions, when a given driver's routine is invoked by a bus driver, the driver detects the bus class to which it is applied and casts the `bus_ops` argument to the appropriate structure (for example, `PciBusOps`).

The `drv_probe` routine is invoked by a bus/nexus driver when the `bus_class` specified in the registry entry matches the bus/nexus driver class. `drv_probe` is called with three arguments.

The `dev_node` argument specifies the parent device node.

The `bus_ops / bus_id` pair specifies the parent device driver.

`bus_ops` points to a structure of service routines implementing a bus driver API. This structure is bus class specific and corresponds to the `bus_class` specified by the driver registry entry.

bus_id is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The `drv_probe` routine is optional. In cases where the `drv_probe` routine is not provided by the driver, the *drv_probe* field must be set to `NULL`.

The `drv_bind` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class.

`drv_bind` is called with one argument, *dev_node*, which specifies a given device node. The `drv_bind` routine enables the driver to perform a driver-to-device binding. Typically, the driver examines properties attached to the device node in order to determine the type of device at the node and to check whether the device may be serviced by the driver. If the check is positive, the driver binds the driver to the device node by attaching a `PROP_DRIVER` property to the device node. The property value specifies the driver name. The parent bus driver uses such a property to determine the name of the driver servicing the device. In other words, via the `PROP_DRIVER` property, the child driver gives its name to the parent bus driver, asking it to invoke the `drv_init` routine on that device. Note that, if a "driver" property is already present in the device node, then the `drv_bind` routine can not continue; `drv_bind` should not override existing driver-to-device binding.

The `drv_bind` routine is optional. If the `drv_bind` routine is not provided by the driver, the *drv_bind* field must be set to `NULL`.

The `drv_init` routine is invoked by a bus/nexus driver when the *bus_class* specified in the registry entry matches the bus/nexus driver class and a given driver is bound to a given device node.

`drv_init` is called with three arguments. The *dev_node* argument specifies the device node for which a device driver instance should be created. The *bus_ops* / *bus_id* pair specifies the parent device driver. *bus_ops* points to a structure of service routines implementing a bus driver API.

This structure is bus class specific and corresponds to the *bus_class* specified by the driver registry entry. *bus_id* is opaque to the driver. It must be passed back to the bus driver when the `open` bus service routine is invoked. The purpose of `drv_init` is to create an instance of the driver servicing the given device, to perform a device hardware initialization and to register the driver instance in the device registry.

Typically, `drv_init` would read the device and bus node properties in order to obtain the assigned bus resources and tunable parameters related to the bus/device. The `drv_init` routine is optional. In a case when the `drv_init` routine is not provided by the driver, the *drv_init* field must be set to `NULL`. A typical example of a probe-only driver is a self-identifying bus enumerator (for example, a PCI enumerator) which is implemented as driver. This type of

driver has the `drv_probe` routine which enumerates devices residing on the bus and creates device nodes. This type of driver obviously does not have the `drv_init` routine.

`drv_unload` is called by the driver registry module (more exactly by the `svDriverUnregister` routine) when somebody wishes to unload the driver code from the system. The purpose of `drv_unload` is to check that the device driver code is not currently being used. The driver must check, for each driver instance, whether it is locked in the device registry.

In case of success, all device instances are removed from the device registry and `K_OK` is returned. Otherwise, the device entries are unchanged in the device registry and `K_EBUSY` is returned.

The `drv_unload` routine is optional. In cases when `drv_unload` is not provided by the driver, the `drv_unload` field must be set to `NULL`. Note that, in this case, the driver code cannot be unloaded.

`svDriverLookupFirst` returns the first driver entry in the registry. When the registry is not empty a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

The driver entry lock is not exclusive. It is simply a counter which prevents the driver component from being unloaded when it is still being used by a bus/nexus driver for probing or initialization.

`svDriverLookupNext` returns the next driver entry in the registry. The current entry is specified by the `drv_id` argument. The current entry must be locked by a previously called `svDriverLookupFirst` or `svDriverLookupNext`. If the current entry is not the last one in the registry a non zero `DrvRegId` (designating the entry) is returned, otherwise `NULL` is returned. In case of success, the next driver entry is locked in the registry. It should be unlocked by a subsequent invocation of `svDriverRelease` or `svDriverUnregister`.

`svDriverRelease` releases the lock of the driver entry specified by the `drv_id` argument.

`svDriverEntry` returns a pointer to the driver entry structure specified by the `drv_id` argument. The device entry being accessed must have been previously locked using `svDriverLookupFirst` or `svDriverLookupNext`. Note that the device entry structure is read-only.

svDriverCap returns a pointer to the driver actor capability. The driver entry is specified by the *drv_id* argument. The driver entry being accessed must have been previously locked using svDriverLookupFirst or svDriverLookupNext . Note that if a given driver entry is registered by a built-in driver, a NULL pointer is returned. In other words, the driver actor capability makes sense only for dynamically loaded drivers. The actor capability may be used by an application in order to delete the driver actor once the driver entry is unregistered. Note that the driver capability structure is read-only.

svDriverUnregister tries to remove the driver entry specified by the *dev_id* argument from the driver registry. The device entry being removed must have been previously locked using svDriverLookupFirst or svDriverLookupNext .

In case of success, K_OK is returned, otherwise K_EBUSY is returned. The K_EBUSY result means that either the driver entry is locked in the driver registry (that is, a static driver routine is currently being used by a bus/nexus driver) or an instance of the device driver is locked in the device registry (that is, there is a driver instance which is currently being used by a driver client). Note that when K_EBUSY is returned, the driver entry remains locked in the registry and should be unlocked explicitly by svDriverRelease .

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svDriverRegister	+	+	-	+
svDriverLookupFirst	+	+	-	+
svDriverLookupNext	+	+	-	+
svDriverRelease	+	+	-	+
svDriverEntry	+	+	-	-
svDriverCap	+	+	-	-
svDriverUnregister	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

dreeNodeRoot(9DKI) , svDkiThreadCall(9DKI)

NAME	svIntrAttach, svIntrDetach, svSoftIntrAttach, svSoftIntrDetach, svTimerIntrAttach, svTimerIntrDetach, svIntrCtxGet – interrupts management
FEATURES	DKI
DESCRIPTION	<p>Provides interrupts management services.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ svIntrAttach_usparc(9DKI) ■ svIntrDetach_usparc(9DKI) ■ svSoftIntrAttach_usparc(9DKI) ■ svSoftIntrDetach_usparc(9DKI) ■ svTimerIntrAttach_usparc(9DKI) ■ svTimerIntrDetach_usparc(9DKI) ■ svIntrCtxGet_usparc(9DKI) ■ svIntrAttach_powerpc(9DKI) ■ svIntrDetach_powerpc(9DKI) ■ svIntrCtxGet_powerpc(9DKI) ■ svIntrAttach_x86(9DKI) ■ svIntrDetach_x86(9DKI) ■ svIntrCtxGet_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svIntrAttach_powerpc, svIntrDetach_powerpc, svIntrCtxGet_powerpc – PowerPC interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(CpuIntr intr, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx); typedef struct { void (*unmask) (CpuIntrId intrId); void (*mask) (CpuIntrId intrId); } CpuIntrOps;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to PowerPC DKI interrupts services.
EXTENDED DESCRIPTION	<p>For PowerPC family processors, interrupts are defined as all asynchronous exceptions. These interrupts are identified using unsigned integer numbers which correspond to the exception number. The microkernel provides services to allow device drivers to manage PowerPC interrupts, mainly to attach/detach handlers to these interrupts.</p> <p>svIntrAttach attaches a handler to a given CPU interrupt. The <i>intr</i> argument specifies the interrupt to attach to. The <i>intrHandler</i> argument specifies the handler to call back when the specified interrupt occurs. The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On success, <code>K_OK</code> is returned and the services defined on the attached interrupt are returned in the <i>intrOps</i> parameter. An identifier for the attached interrupt is also returned in <i>intrId</i>. This identifier must be used as the first parameter to further calls to <i>intrOps</i> services.</p> <p>Services available on an attached interrupt are defined by the <code>CpuIntrOps</code> structure as follows:</p> <p>The <code>CpuIntrOps.unmask</code> routine enables the interrupt at CPU level identified by <i>intrId</i>.</p> <p>The <code>CpuIntrOps.mask</code> routine disables the interrupt at CPU level identified by <i>intrId</i>.</p>

Note that `SYSTEM_RESET_INTR` and `MACHINE_CHECK_INTR` are non-maskable interrupts. Thus `unmask/mask` have no effect on these interrupts.

Note also, that all other interrupts are masked using the same bit in the PowerPC processor MSR register. Thus there is no way to `unmask/mask` only one of these interrupts separately from the others. When an interrupt occurs, the attached `CpuIntrHandler` is invoked with interrupts masked at processor level. This behaves in the same way as if `imsIntrMask_f()` was just called prior to invoking the handler (`imsIntrMaskCount_f` is positive). It is up to the interrupt handler called to use `unmask` to allow the interrupt to be nested or not. Typically, a host bus driver handler should handle PowerPC external interrupts as follows:

- Identify the interrupt source (through a PIC or special cycle).
- Unmask interrupts at processor level (`unmask`).
- Call handlers attached to the identified source.
- Optionally `mask` interrupts to do critical tasks, such as notifying the end of the interrupt to a PIC.
- Return to the DKI.

On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The interrupt specified in <code>intr</code> is invalid (not in the list above).
<code>K_EFAIL</code>	The interrupt specified in <code>intr</code> is valid, but a handler is already attached to it.
<code>K_ENOMEM</code>	The system is out of physical memory.

`svIntrDetach` detaches an interrupt handler previously connected by `svIntrAttach`. The `intrId` argument identifies the attached interrupt handler, and was previously returned by `svIntrAttach`.

`svIntrCtxGet` retrieves the current level interrupt context. It is mainly to be used for profiling purposes. On success, `K_OK` is returned and a pointer to the current level interrupt saved context is returned in the `intrCtx` argument. The CPU context saved on interrupt contains the volatile general purpose and floating point registers, condition register, Machine Status Register, instruction pointer, link register and the interrupt number itself. On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The interrupt level is zero (not called from an interrupt handler).
-----------------------	---

K_ENOTAVAILABLE There is no context available for the currently handled interrupt.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svIntrAttach	+	+	-	+
CpuIntrOps.unmask	+	+	+	-
CpuIntrOps.mask	+	+	+	-
svIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap(9DKI)`, `svDkiThreadCall(9DKI)`, `imsIntrMask_f(9DKI)`

NAME	svIntrAttach_usparc, svIntrDetach_usparc, svSoftIntrAttach_usparc, svSoftIntrDetach_usparc, svTimerIntrAttach_uparc, svTimerIntrDetach_uparc, svIntrCtxGet_usparc – UltraSPARC interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(unsigned int intrNumb, unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrIdp); void svIntrDetach(CpuIntrId intrId); KnError svSoftIntrAttach(unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svSoftIntrDetach(CpuIntrId intrId); KnError svTimerIntrAttach(CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svTimerIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx);</pre>
FEATURES	DKI
DESCRIPTION	<p>Provides UltraSPARC interrupts management services.</p> <p>An UltraSPARC processor is an implementation of the SPARC-V9 CPU architecture. As specified by the SPARC-V9 architecture, there are fifteen CPU interrupt sources assigned to the fifteen CPU interrupt levels from 1 up to 15. However, the UltraSPARC processor uses these fifteen interrupt levels only for software generated interrupts. The hardware interrupts are delivered to the processor using the "Mondo" interrupt transfer mechanism. The hardware interrupt source is designated by the interrupt number. Typically, the interrupt number is 11-bit width and composed of the interrupt group number (5 MSB) and the interrupt offset number (6 LSB).</p> <p>The "Mondo" interrupt dispatch handler is built into the microkernel. The microkernel handles a mapping between the hardware interrupt numbers and software interrupt levels. When a mondo interrupt packet is received by the microkernel, the interrupt request descriptor is queued and an associated software interrupt is triggered. The microkernel software interrupt handler then dequeues the interrupt request descriptor and invokes a handler associated to the given interrupt number. In this way, a driver interrupt handler is always invoked in the UltraSPARC software interrupt context. This provides the interrupt handler with an environment analogous to the SPARC-V8 one (interrupt levels).</p> <p>The microkernel provides services which allow device drivers to manage UltraSPARC interrupts, mainly to attach/detach handlers to the CPU interrupts.</p> <pre>typedef void (*CpuIntrHandler)(void*);</pre>

svIntrAttach

svIntrAttach attaches a given handler to a given interrupt number at a given processor interrupt level.

The *intrNumb* argument specifies the interrupt number to which to attach.

The *intrLevel* argument specifies the interrupt level to which to attach.

The *intrHandler* argument specifies the handler to call back when the given interrupt occurs.

The *intrCookie* argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.

On success, *K_OK* is returned and the services defined on the attached interrupt are returned in the *intrOps* parameter. An identifier for the attached interrupt is also returned in *intrId*. This identifier must be used as first argument to subsequent calls to *intrOps* services.

On failure, an error code is returned as follows:

<i>K_EINVAL</i>	The specified interrupt number or interrupt level are invalid.
<i>K_BUSY</i>	Another handler is already attached to the given <i>intrNumb</i> .
<i>K_ENOMEM</i>	The system is out of memory.

Services available on an attached interrupt are defined by the *CpuIntrOps* structure as follows:

```
typedef struct CpuIntrOps {
    void
    (*mask) (CpuIntrId intrId);

    void
    (*unmask) (CpuIntrId intrId);

    void
    (*enable) (CpuIntrId intrId);

    void
    (*disable) (CpuIntrId intrId);

    void
    (*trigger) (CpuIntrId intrId);
} CpuIntrOps;
```

The `CpuIntrOps.mask` routine disables the interrupt at CPU level identified by `intrId`. In other words, the PIL register is set to the level corresponding to `intrId`. Note that the original value of the PIL register is saved by DKI in order to be restored later by the `CpuIntrOps.unmask` routine.

The `CpuIntrOps.unmask` routine enables the interrupt at CPU level identified by `intrId`.

In other words, the PIL register is restored to the original value saved by the previously called `CpuIntrOps.mask` routine.

The `mask / unmask` pair may be called from base level only and must not be nested. The `mask/unmask` pair is typically used to implement a critical section of code which needs to be protected against the interrupt. Note that with respect to the SPARC-V9 architecture, when an interrupt level `N` is masked, all interrupts with a level less than `N` are also masked. Thus, there is no way to `mask` only one CPU interrupt level except the lowest one.

The `enable / disable` pair may only be called from the attached interrupt handler. When an interrupt occurs, the attached `CpuIntrHandler` is invoked with the interrupt masked at processor level. This behaves in exactly the same way as if `CpuIntrOps.disable` was called just prior to the handler invocation. In other words, the PIL register is set to the interrupt level and the original interrupt processor level (which was when the interrupt occurred) is saved by DKI. Note that the interrupt handler must return to DKI in the same context as it was called, that is with the interrupt disabled at processor level.

On the other hand, the called interrupt handler may use the `enable/disable` pair to allow the interrupt to be nested. This feature is typically used by a host bus driver when the bus interrupts are multiplexed, that is, multiple bus interrupts are reported at the same CPU interrupt level. Typically, an interrupt handler of this type of host bus driver would take the following actions:

- Identify the bus interrupt source (through a PIC or special cycle).
- Disable the bus interrupt source at bus level (through PIC).
- Enable interrupt at processor level (`enable`).
- Call handlers attached to the identified bus interrupt source.
- Disable interrupt at processor level (`disable`).
- Acknowledge (if needed) and enable the bus interrupt source at bus level (through PIC).
- Return to the DKI.

The `CpuIntrOps.trigger` routine allows the interrupt to be triggered by software. Basically, this routine acts like the mondo interrupt dispatcher except the interrupt number is obtained from `intrId` rather than from the mondo interrupt packet. The `CpuIntrOps.trigger` routine is mainly dedicated to the

	software interrupts attached by <code>svSoftIntrAttach</code> . However, it may be also used for hardware interrupts, for instance, for debugging or diagnostic purposes.				
svIntrDetach	<p><code>svIntrDetach</code> detaches an interrupt handler previously connected by <code>svIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svIntrAttach</code> .</p>				
svSoftIntrDetach	<p><code>svSoftIntrAttach</code> attaches a given software interrupt handler to a given processor interrupt level.</p> <p>The <i>intrLevel</i> argument specifies the interrupt level to which to attach.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On success, <code>K_OK</code> is returned and the services defined on the attached interrupt are returned in the <i>intrOps</i> parameter.</p> <p>An identifier for the attached interrupt is also returned in <i>intrId</i> . This identifier must be used as the first argument to subsequent calls to <i>intrOps</i> services.</p> <p>On failure, an error code is returned as follows:</p> <table border="0"> <tr> <td><code>K_EINVAL</code></td> <td>The specified interrupt level is invalid.</td> </tr> <tr> <td><code>K_ENOMEM</code></td> <td>The system is out of memory.</td> </tr> </table> <p>Services available on an attached interrupt are defined by the <code>CpuIntrOps</code> structure described above.</p>	<code>K_EINVAL</code>	The specified interrupt level is invalid.	<code>K_ENOMEM</code>	The system is out of memory.
<code>K_EINVAL</code>	The specified interrupt level is invalid.				
<code>K_ENOMEM</code>	The system is out of memory.				
svSoftIntrDetach	<p><code>svSoftIntrDetach</code> detaches a software interrupt handler previously connected by <code>svSoftIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svSoftIntrAttach</code> .</p>				
svTimerIntrAttach	<p><code>svTimerIntrAttach</code> attaches a given interrupt handler to the UltraSPARC tick-counter interrupt.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque the microkernel.</p>				

On success, `K_OK` is returned and the services defined on the attached interrupt are returned in the `intrOps` parameter. An identifier for the attached interrupt is also returned in `intrId`. This identifier must be used as the first argument to subsequent calls to `intrOps` services.

On failure, an error code is returned as follows:

- `K_BUSY` Another handler is already attached to the tick-counter interrupt.
- `K_ENOMEM` The system is out of memory.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure described above.

svTimerIntrDetach

`svTimerIntrDetach` detaches the interrupt handler previously connected by `svTimerIntrAttach`.

The `intrId` argument identifies the attached interrupt handler, previously returned by `svTimerIntrAttach`.

svIntrCtxGet

`svIntrCtxGet` retrieves the current level interrupt context. It is typically used for profiling purposes.

On success, `K_OK` is returned and a pointer to the recently saved interrupt context is returned in the `intrCtx` argument. The CPU context saved on interrupt has the same structure as the thread context saved on exception or trap. It contains the global registers `%g1-%g7`, the output registers of the interrupted window `%o0-%o7` and the following processor registers: `%tstate`, `%pc`, `%npc`, `%tt`, `%y`. In addition, the thread context contains the number of outstanding windows and the pointer to the outstanding windows buffer if the number of windows is greater than zero.

On failure, an error code is returned as follows:

- `K_EINVAL` Interrupt level is zero (not called from an interrupt handler).
- `K_ENOTAVAILABLE` There is no context available for currently handled interrupt.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each interrupt management service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svIntrAttach</code>	+	+	-	+

svIntrDetach	+	+	-	+
svIntrSoftAttach	+	+	-	+
svIntrSoftDetach	+	+	-	+
svTimerIntrAttach	+	+	-	+
svTimerIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-
CpuIntrOps.mask	+	+	-	-
CpuIntrOps.unmask	+	+	-	-
CpuIntrOps.enable	-	-	+	-
CpuIntrOps.disable	-	-	+	-
CpuIntrOps.trigger	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME	svIntrAttach_x86, svIntrDetach_x86, svIntrCtxGet_x86 – Intel x86 interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(CpuIntr intr, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx); typedef struct CpuIntrOps { void (*enable) (CpuIntrId intrId); void (*disable) (CpuIntrId intrId); } CpuIntrOps;</pre>
FEATURES	DKI
DESCRIPTION	Provides Intel x86 interrupts management services.
EXTENDED DESCRIPTION	<p>For Intel family processors, there are two kinds of interrupts:</p> <p>Exceptions Exceptions are generated and detected by the processor when executing instructions, such as division by 0. These exceptions are classified as faults, traps and aborts. Another type of exception exists: Programmed exceptions generated by INTx instructions.</p> <p>Interrupts Interrupts are generated by events external to the processor, such as requests to service peripheral devices.</p> <p>There are two further types of interrupts:</p> <p>Maskable Maskable interrupts are received on the INTR input pin of the Intel x86 processor. They do not occur if the IF flag of the EFLAGS register is not set.</p> <p>Non Maskable Non maskable interrupts are received on the NMI input pin of the Intel x86 processor. There is no mechanism to prevent non maskable interrupts.</p>

The DKI does not provide support for exceptions, only for interrupts. The processor associates an identifying number with each different interrupt (non maskable and maskable). This number is called a vector in the range from 0-255 with the range of 0-31 reserved for exceptions. The DKI associates an identifying numeric constant with each interrupt.

The following numeric constants are available:

- NMI_INT (non maskable).
- INTR_BASE_NUMBER (maskable)

The DKI provides support for 64 interrupts, ranging from INTR_BASE_NUMBER to INTR_BASE_NUMBER + 64.

The DKI provides services to allow host bus drivers to manage Intel ix86 interrupts, mainly to attach/detach handlers to/from these interrupts.

`svIntrAttach`

`svIntrAttach` attaches a handler to a given CPU interrupt.

The *intr* argument specifies the interrupt to which to attach.

The *intrHandler* argument specifies the handler to call back when the specified interrupt occurs.

The *intrCookie* argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.

On success, `K_OK` is returned and services defined on the attached interrupt are returned in the *intrOps* parameter. An identifier for the attached interrupt is also returned in *intrId*. This identifier must be used as the first parameter to further calls to *intrOps* services.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure as follows:

The `CpuIntrOps.enable` routine enables the interrupt at CPU level identified by *intrId*.

The `CpuIntrOps.disable` routine disables the interrupt at CPU level identified by *intrId*.

Note that `NMI_INT` is a non maskable interrupt. Thus, `enable/disable` have no effect on these interrupts..

Note also that all other interrupts are masked using the same bit on an Intel x86 processor `EFLAGS` register. Thus there is no way to `enable/disable` a single interrupt separately from the others.

When an interrupt occurs, the attached *CpuIntrHandler* is invoked with interrupts masked at processor level. This produces the same effect as if `imsIntrMask_f()` was called just prior to invoking the handler (`imsIntrMaskCount_f` is positive). Two parameters are passed to this handler:

- Provided *intrCookie* parameter
- Detected *intr*

It is up to the called interrupt handler to use `enable` in order to allow or disallow interrupt nesting.

Typically, a host bus driver handler should handle x86 external interrupts as follows:

- Enable interrupts at processor level (`enable`).
- Call handler(s) attached to the identified cpu interrupt.
- Optionally `disable` interrupts in order to perform critical tasks, such as notifying the end of an interrupt to a PIC.
- Return to the DKI.

On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The interrupt specified in <i>intr</i> is invalid (not in the list above). Or, the interrupt specified in <i>intr</i> is valid, but a handler is already attached to it.
-----------------------	--

`svIntrDetach`

`svIntrDetach` detaches an interrupt handler previously connected by `svIntrAttach`.

The *intrId* argument identifies the attached interrupt handler, and was previously returned by `svIntrAttach`.

`svIntrCtxGet`

`svIntrCtxGet` retrieves the current level interrupt context. It is mainly to be used for profiling purposes.

On success, `K_OK` is returned and a pointer to the current level of interrupt saved context is returned in the *intrCtx* argument. The CPU context saved on interrupt contains the interrupted PC composed of the `eip` and `cs` registers, the `EFLAGS` register, and the stack pointer.

On failure, an error code is returned as follows:

K_EINVAL Interrupt level is zero (not called from an interrupt handler).
 K_ENOTAVAILABLE There is no context available for the currently handled interrupt.

Intel x86 Interrupts Management Allowed Calling Contexts

The table below specifies the contexts in which a caller is allowed to invoke each services.

Services	Base level	DKI thread	Interrupt	Blocking
svIntrAttach	+	+	-	+
CpuIntrOps.enable	+	+	+	-
CpuIntrOps.disable	+	+	+	-
svIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	svIntrAttach, svIntrDetach, svSoftIntrAttach, svSoftIntrDetach, svTimerIntrAttach, svTimerIntrDetach, svIntrCtxGet – interrupts management
FEATURES	DKI
DESCRIPTION	<p>Provides interrupts management services.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ svIntrAttach_usparc(9DKI) ■ svIntrDetach_usparc(9DKI) ■ svSoftIntrAttach_usparc(9DKI) ■ svSoftIntrDetach_usparc(9DKI) ■ svTimerIntrAttach_usparc(9DKI) ■ svTimerIntrDetach_usparc(9DKI) ■ svIntrCtxGet_usparc(9DKI) ■ svIntrAttach_powerpc(9DKI) ■ svIntrDetach_powerpc(9DKI) ■ svIntrCtxGet_powerpc(9DKI) ■ svIntrAttach_x86(9DKI) ■ svIntrDetach_x86(9DKI) ■ svIntrCtxGet_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svIntrAttach_powerpc, svIntrDetach_powerpc, svIntrCtxGet_powerpc – PowerPC interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(CpuIntr intr, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx); typedef struct { void (*unmask) (CpuIntrId intrId); void (*mask) (CpuIntrId intrId); } CpuIntrOps;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to PowerPC DKI interrupts services.
EXTENDED DESCRIPTION	<p>For PowerPC family processors, interrupts are defined as all asynchronous exceptions. These interrupts are identified using unsigned integer numbers which correspond to the exception number. The microkernel provides services to allow device drivers to manage PowerPC interrupts, mainly to attach/detach handlers to these interrupts.</p> <p>svIntrAttach attaches a handler to a given CPU interrupt. The <i>intr</i> argument specifies the interrupt to attach to. The <i>intrHandler</i> argument specifies the handler to call back when the specified interrupt occurs. The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On success, K_OK is returned and the services defined on the attached interrupt are returned in the <i>intrOps</i> parameter. An identifier for the attached interrupt is also returned in <i>intrId</i>. This identifier must be used as the first parameter to further calls to <i>intrOps</i> services.</p> <p>Services available on an attached interrupt are defined by the CpuIntrOps structure as follows:</p> <p>The CpuIntrOps.unmask routine enables the interrupt at CPU level identified by <i>intrId</i>.</p> <p>The CpuIntrOps.mask routine disables the interrupt at CPU level identified by <i>intrId</i>.</p>

Note that `SYSTEM_RESET_INTR` and `MACHINE_CHECK_INTR` are non-maskable interrupts. Thus `unmask/mask` have no effect on these interrupts.

Note also, that all other interrupts are masked using the same bit in the PowerPC processor MSR register. Thus there is no way to `unmask/mask` only one of these interrupts separately from the others. When an interrupt occurs, the attached `CpuIntrHandler` is invoked with interrupts masked at processor level. This behaves in the same way as if `imsIntrMask_f()` was just called prior to invoking the handler (`imsIntrMaskCount_f` is positive). It is up to the interrupt handler called to use `unmask` to allow the interrupt to be nested or not. Typically, a host bus driver handler should handle PowerPC external interrupts as follows:

- Identify the interrupt source (through a PIC or special cycle).
- Unmask interrupts at processor level (`unmask`).
- Call handlers attached to the identified source.
- Optionally `mask` interrupts to do critical tasks, such as notifying the end of the interrupt to a PIC.
- Return to the DKI.

On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The interrupt specified in <code>intr</code> is invalid (not in the list above).
<code>K_EFAIL</code>	The interrupt specified in <code>intr</code> is valid, but a handler is already attached to it.
<code>K_ENOMEM</code>	The system is out of physical memory.

`svIntrDetach` detaches an interrupt handler previously connected by `svIntrAttach`. The `intrId` argument identifies the attached interrupt handler, and was previously returned by `svIntrAttach`.

`svIntrCtxGet` retrieves the current level interrupt context. It is mainly to be used for profiling purposes. On success, `K_OK` is returned and a pointer to the current level interrupt saved context is returned in the `intrCtx` argument. The CPU context saved on interrupt contains the volatile general purpose and floating point registers, condition register, Machine Status Register, instruction pointer, link register and the interrupt number itself. On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The interrupt level is zero (not called from an interrupt handler).
-----------------------	---

K_ENOTAVAILABLE There is no context available for the currently handled interrupt.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svIntrAttach	+	+	-	+
CpuIntrOps.unmask	+	+	+	-
CpuIntrOps.mask	+	+	+	-
svIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap(9DKI)` , `svDkiThreadCall(9DKI)` , `imsIntrMask_f(9DKI)`

NAME	svIntrAttach_usparc, svIntrDetach_usparc, svSoftIntrAttach_usparc, svSoftIntrDetach_usparc, svTimerIntrAttach_uparc, svTimerIntrDetach_uparc, svIntrCtxGet_usparc – UltraSPARC interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(unsigned int intrNumb, unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrIdp); void svIntrDetach(CpuIntrId intrId); KnError svSoftIntrAttach(unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svSoftIntrDetach(CpuIntrId intrId); KnError svTimerIntrAttach(CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svTimerIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx);</pre>
FEATURES	DKI
DESCRIPTION	<p>Provides UltraSPARC interrupts management services.</p> <p>An UltraSPARC processor is an implementation of the SPARC-V9 CPU architecture. As specified by the SPARC-V9 architecture, there are fifteen CPU interrupt sources assigned to the fifteen CPU interrupt levels from 1 up to 15. However, the UltraSPARC processor uses these fifteen interrupt levels only for software generated interrupts. The hardware interrupts are delivered to the processor using the "Mondo" interrupt transfer mechanism. The hardware interrupt source is designated by the interrupt number. Typically, the interrupt number is 11-bit width and composed of the interrupt group number (5 MSB) and the interrupt offset number (6 LSB).</p> <p>The "Mondo" interrupt dispatch handler is built into the microkernel. The microkernel handles a mapping between the hardware interrupt numbers and software interrupt levels. When a mondo interrupt packet is received by the microkernel, the interrupt request descriptor is queued and an associated software interrupt is triggered. The microkernel software interrupt handler then dequeues the interrupt request descriptor and invokes a handler associated to the given interrupt number. In this way, a driver interrupt handler is always invoked in the UltraSPARC software interrupt context. This provides the interrupt handler with an environment analogous to the SPARC-V8 one (interrupt levels).</p> <p>The microkernel provides services which allow device drivers to manage UltraSPARC interrupts, mainly to attach/detach handlers to the CPU interrupts.</p> <pre>typedef void (*CpuIntrHandler)(void*);</pre>

svIntrAttach

svIntrAttach attaches a given handler to a given interrupt number at a given processor interrupt level.

The *intrNumb* argument specifies the interrupt number to which to attach.

The *intrLevel* argument specifies the interrupt level to which to attach.

The *intrHandler* argument specifies the handler to call back when the given interrupt occurs.

The *intrCookie* argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.

On success, K_OK is returned and the services defined on the attached interrupt are returned in the *intrOps* parameter. An identifier for the attached interrupt is also returned in *intrId*. This identifier must be used as first argument to subsequent calls to *intrOps* services.

On failure, an error code is returned as follows:

K_EINVAL	The specified interrupt number or interrupt level are invalid.
K_BUSY	Another handler is already attached to the given <i>intrNumb</i> .
K_ENOMEM	The system is out of memory.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure as follows:

```
typedef struct CpuIntrOps {
    void
    (*mask) (CpuIntrId intrId);

    void
    (*unmask) (CpuIntrId intrId);

    void
    (*enable) (CpuIntrId intrId);

    void
    (*disable) (CpuIntrId intrId);

    void
    (*trigger) (CpuIntrId intrId);
} CpuIntrOps;
```

The `CpuIntrOps.mask` routine disables the interrupt at CPU level identified by `intrId`. In other words, the PIL register is set to the level corresponding to `intrId`. Note that the original value of the PIL register is saved by DKI in order to be restored later by the `CpuIntrOps.unmask` routine.

The `CpuIntrOps.unmask` routine enables the interrupt at CPU level identified by `intrId`.

In other words, the PIL register is restored to the original value saved by the previously called `CpuIntrOps.mask` routine.

The `mask / unmask` pair may be called from base level only and must not be nested. The `mask/unmask` pair is typically used to implement a critical section of code which needs to be protected against the interrupt. Note that with respect to the SPARC-V9 architecture, when an interrupt level `N` is masked, all interrupts with a level less than `N` are also masked. Thus, there is no way to `mask` only one CPU interrupt level except the lowest one.

The `enable / disable` pair may only be called from the attached interrupt handler. When an interrupt occurs, the attached `CpuIntrHandler` is invoked with the interrupt masked at processor level. This behaves in exactly the same way as if `CpuIntrOps.disable` was called just prior to the handler invocation. In other words, the PIL register is set to the interrupt level and the original interrupt processor level (which was when the interrupt occurred) is saved by DKI. Note that the interrupt handler must return to DKI in the same context as it was called, that is with the interrupt disabled at processor level.

On the other hand, the called interrupt handler may use the `enable/disable` pair to allow the interrupt to be nested. This feature is typically used by a host bus driver when the bus interrupts are multiplexed, that is, multiple bus interrupts are reported at the same CPU interrupt level. Typically, an interrupt handler of this type of host bus driver would take the following actions:

- Identify the bus interrupt source (through a PIC or special cycle).
- Disable the bus interrupt source at bus level (through PIC).
- Enable interrupt at processor level (`enable`).
- Call handlers attached to the identified bus interrupt source.
- Disable interrupt at processor level (`disable`).
- Acknowledge (if needed) and enable the bus interrupt source at bus level (through PIC).
- Return to the DKI.

The `CpuIntrOps.trigger` routine allows the interrupt to be triggered by software. Basically, this routine acts like the mondo interrupt dispatcher except the interrupt number is obtained from `intrId` rather than from the mondo interrupt packet. The `CpuIntrOps.trigger` routine is mainly dedicated to the

	software interrupts attached by <code>svSoftIntrAttach</code> . However, it may be also used for hardware interrupts, for instance, for debugging or diagnostic purposes.				
svIntrDetach	<p><code>svIntrDetach</code> detaches an interrupt handler previously connected by <code>svIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svIntrAttach</code> .</p>				
svSoftIntrDetach	<p><code>svSoftIntrAttach</code> attaches a given software interrupt handler to a given processor interrupt level.</p> <p>The <i>intrLevel</i> argument specifies the interrupt level to which to attach.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On success, <code>K_OK</code> is returned and the services defined on the attached interrupt are returned in the <i>intrOps</i> parameter.</p> <p>An identifier for the attached interrupt is also returned in <i>intrId</i> . This identifier must be used as the first argument to subsequent calls to <i>intrOps</i> services.</p> <p>On failure, an error code is returned as follows:</p> <table border="0"> <tr> <td><code>K_EINVAL</code></td> <td>The specified interrupt level is invalid.</td> </tr> <tr> <td><code>K_ENOMEM</code></td> <td>The system is out of memory.</td> </tr> </table> <p>Services available on an attached interrupt are defined by the <code>CpuIntrOps</code> structure described above.</p>	<code>K_EINVAL</code>	The specified interrupt level is invalid.	<code>K_ENOMEM</code>	The system is out of memory.
<code>K_EINVAL</code>	The specified interrupt level is invalid.				
<code>K_ENOMEM</code>	The system is out of memory.				
svSoftIntrDetach	<p><code>svSoftIntrDetach</code> detaches a software interrupt handler previously connected by <code>svSoftIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svSoftIntrAttach</code> .</p>				
svTimerIntrAttach	<p><code>svTimerIntrAttach</code> attaches a given interrupt handler to the UltraSPARC tick-counter interrupt.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque the microkernel.</p>				

On success, `K_OK` is returned and the services defined on the attached interrupt are returned in the `intrOps` parameter. An identifier for the attached interrupt is also returned in `intrId`. This identifier must be used as the first argument to subsequent calls to `intrOps` services.

On failure, an error code is returned as follows:

- `K_BUSY` Another handler is already attached to the tick-counter interrupt.
- `K_ENOMEM` The system is out of memory.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure described above.

svTimerIntrDetach

`svTimerIntrDetach` detaches the interrupt handler previously connected by `svTimerIntrAttach`.

The `intrId` argument identifies the attached interrupt handler, previously returned by `svTimerIntrAttach`.

svIntrCtxGet

`svIntrCtxGet` retrieves the current level interrupt context. It is typically used for profiling purposes.

On success, `K_OK` is returned and a pointer to the recently saved interrupt context is returned in the `intrCtx` argument. The CPU context saved on interrupt has the same structure as the thread context saved on exception or trap. It contains the global registers `%g1-%g7`, the output registers of the interrupted window `%o0-%o7` and the following processor registers: `%tstate`, `%pc`, `%npc`, `%tt`, `%y`. In addition, the thread context contains the number of outstanding windows and the pointer to the outstanding windows buffer if the number of windows is greater than zero.

On failure, an error code is returned as follows:

- `K_EINVAL` Interrupt level is zero (not called from an interrupt handler).
- `K_ENOTAVAILABLE` There is no context available for currently handled interrupt.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each interrupt management service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svIntrAttach</code>	+	+	-	+

svIntrDetach	+	+	-	+
svIntrSoftAttach	+	+	-	+
svIntrSoftDetach	+	+	-	+
svTimerIntrAttach	+	+	-	+
svTimerIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-
CpuIntrOps.mask	+	+	-	-
CpuIntrOps.unmask	+	+	-	-
CpuIntrOps.enable	-	-	+	-
CpuIntrOps.disable	-	-	+	-
CpuIntrOps.trigger	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME	svIntrAttach_x86, svIntrDetach_x86, svIntrCtxGet_x86 – Intel x86 interrupts management	
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(CpuIntr intr, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx); typedef struct CpuIntrOps { void (*enable) (CpuIntrId intrId); void (*disable) (CpuIntrId intrId); } CpuIntrOps;</pre>	
FEATURES	DKI	
DESCRIPTION	Provides Intel x86 interrupts management services.	
EXTENDED DESCRIPTION	For Intel family processors, there are two kinds of interrupts:	
	Exceptions	Exceptions are generated and detected by the processor when executing instructions, such as division by 0. These exceptions are classified as faults, traps and aborts. Another type of exception exists: Programmed exceptions generated by INTx instructions.
	Interrupts	Interrupts are generated by events external to the processor, such as requests to service peripheral devices.
	There are two further types of interrupts:	
	Maskable	Maskable interrupts are received on the INTR input pin of the Intel x86 processor. They do not occur if the IF flag of the EFLAGS register is not set.
	Non Maskable	Non maskable interrupts are received on the NMI input pin of the Intel x86 processor. There is no mechanism to prevent non maskable interrupts.

The DKI does not provide support for exceptions, only for interrupts. The processor associates an identifying number with each different interrupt (non maskable and maskable). This number is called a vector in the range from 0-255 with the range of 0-31 reserved for exceptions. The DKI associates an identifying numeric constant with each interrupt.

The following numeric constants are available:

- NMI_INT (non maskable).
- INTR_BASE_NUMBER (maskable)

The DKI provides support for 64 interrupts, ranging from INTR_BASE_NUMBER to INTR_BASE_NUMBER + 64.

The DKI provides services to allow host bus drivers to manage Intel ix86 interrupts, mainly to attach/detach handlers to/from these interrupts.

`svIntrAttach`

`svIntrAttach` attaches a handler to a given CPU interrupt.

The *intr* argument specifies the interrupt to which to attach.

The *intrHandler* argument specifies the handler to call back when the specified interrupt occurs.

The *intrCookie* argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.

On success, `K_OK` is returned and services defined on the attached interrupt are returned in the *intrOps* parameter. An identifier for the attached interrupt is also returned in *intrId*. This identifier must be used as the first parameter to further calls to *intrOps* services.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure as follows:

The `CpuIntrOps.enable` routine enables the interrupt at CPU level identified by *intrId*.

The `CpuIntrOps.disable` routine disables the interrupt at CPU level identified by *intrId*.

Note that `NMI_INT` is a non maskable interrupt. Thus, `enable/disable` have no effect on these interrupts..

Note also that all other interrupts are masked using the same bit on an Intel x86 processor `EFLAGS` register. Thus there is no way to `enable/disable` a single interrupt separately from the others.

When an interrupt occurs, the attached *CpuIntrHandler* is invoked with interrupts masked at processor level. This produces the same effect as if `imsIntrMask_f()` was called just prior to invoking the handler (`imsIntrMaskCount_f` is positive). Two parameters are passed to this handler:

- Provided *intrCookie* parameter
- Detected *intr*

It is up to the called interrupt handler to use `enable` in order to allow or disallow interrupt nesting.

Typically, a host bus driver handler should handle x86 external interrupts as follows:

- Enable interrupts at processor level (`enable`).
- Call handler(s) attached to the identified cpu interrupt.
- Optionally `disable` interrupts in order to perform critical tasks, such as notifying the end of an interrupt to a PIC.
- Return to the DKI.

On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The interrupt specified in <i>intr</i> is invalid (not in the list above). Or, the interrupt specified in <i>intr</i> is valid, but a handler is already attached to it.
-----------------------	--

`svIntrDetach`

`svIntrDetach` detaches an interrupt handler previously connected by `svIntrAttach`.

The *intrId* argument identifies the attached interrupt handler, and was previously returned by `svIntrAttach`.

`svIntrCtxGet`

`svIntrCtxGet` retrieves the current level interrupt context. It is mainly to be used for profiling purposes.

On success, `K_OK` is returned and a pointer to the current level of interrupt saved context is returned in the *intrCtx* argument. The CPU context saved on interrupt contains the interrupted PC composed of the `eip` and `cs` registers, the `EFLAGS` register, and the stack pointer.

On failure, an error code is returned as follows:

K_EINVAL Interrupt level is zero (not called from an interrupt handler).
 K_ENOTAVAILABLE There is no context available for the currently handled interrupt.

Intel x86 Interrupts Management Allowed Calling Contexts

The table below specifies the contexts in which a caller is allowed to invoke each services.

Services	Base level	DKI thread	Interrupt	Blocking
svIntrAttach	+	+	-	+
CpuIntrOps.enable	+	+	+	-
CpuIntrOps.disable	+	+	+	-
svIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME	svIntrAttach, svIntrDetach, svSoftIntrAttach, svSoftIntrDetach, svTimerIntrAttach, svTimerIntrDetach, svIntrCtxGet – interrupts management
FEATURES	DKI
DESCRIPTION	<p>Provides interrupts management services.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ svIntrAttach_usparc(9DKI) ■ svIntrDetach_usparc(9DKI) ■ svSoftIntrAttach_usparc(9DKI) ■ svSoftIntrDetach_usparc(9DKI) ■ svTimerIntrAttach_usparc(9DKI) ■ svTimerIntrDetach_usparc(9DKI) ■ svIntrCtxGet_usparc(9DKI) ■ svIntrAttach_powerpc(9DKI) ■ svIntrDetach_powerpc(9DKI) ■ svIntrCtxGet_powerpc(9DKI) ■ svIntrAttach_x86(9DKI) ■ svIntrDetach_x86(9DKI) ■ svIntrCtxGet_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svIntrAttach_powerpc, svIntrDetach_powerpc, svIntrCtxGet_powerpc – PowerPC interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(CpuIntr intr, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx); typedef struct { void (*unmask) (CpuIntrId intrId); void (*mask) (CpuIntrId intrId); } CpuIntrOps;</pre>
FEATURES	DKI
DESCRIPTION	Provides access to PowerPC DKI interrupts services.
EXTENDED DESCRIPTION	<p>For PowerPC family processors, interrupts are defined as all asynchronous exceptions. These interrupts are identified using unsigned integer numbers which correspond to the exception number. The microkernel provides services to allow device drivers to manage PowerPC interrupts, mainly to attach/detach handlers to these interrupts.</p> <p>svIntrAttach attaches a handler to a given CPU interrupt. The <i>intr</i> argument specifies the interrupt to attach to. The <i>intrHandler</i> argument specifies the handler to call back when the specified interrupt occurs. The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On success, <code>K_OK</code> is returned and the services defined on the attached interrupt are returned in the <i>intrOps</i> parameter. An identifier for the attached interrupt is also returned in <i>intrId</i>. This identifier must be used as the first parameter to further calls to <i>intrOps</i> services.</p> <p>Services available on an attached interrupt are defined by the <code>CpuIntrOps</code> structure as follows:</p> <p>The <code>CpuIntrOps.unmask</code> routine enables the interrupt at CPU level identified by <i>intrId</i>.</p> <p>The <code>CpuIntrOps.mask</code> routine disables the interrupt at CPU level identified by <i>intrId</i>.</p>

Note that `SYSTEM_RESET_INTR` and `MACHINE_CHECK_INTR` are non-maskable interrupts. Thus `unmask/mask` have no effect on these interrupts.

Note also, that all other interrupts are masked using the same bit in the PowerPC processor MSR register. Thus there is no way to `unmask/mask` only one of these interrupts separately from the others. When an interrupt occurs, the attached `CpuIntrHandler` is invoked with interrupts masked at processor level. This behaves in the same way as if `imsIntrMask_f()` was just called prior to invoking the handler (`imsIntrMaskCount_f` is positive). It is up to the interrupt handler called to use `unmask` to allow the interrupt to be nested or not. Typically, a host bus driver handler should handle PowerPC external interrupts as follows:

- Identify the interrupt source (through a PIC or special cycle).
- Unmask interrupts at processor level (`unmask`).
- Call handlers attached to the identified source.
- Optionally `mask` interrupts to do critical tasks, such as notifying the end of the interrupt to a PIC.
- Return to the DKI.

On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The interrupt specified in <code>intr</code> is invalid (not in the list above).
<code>K_EFAIL</code>	The interrupt specified in <code>intr</code> is valid, but a handler is already attached to it.
<code>K_ENOMEM</code>	The system is out of physical memory.

`svIntrDetach` detaches an interrupt handler previously connected by `svIntrAttach`. The `intrId` argument identifies the attached interrupt handler, and was previously returned by `svIntrAttach`.

`svIntrCtxGet` retrieves the current level interrupt context. It is mainly to be used for profiling purposes. On success, `K_OK` is returned and a pointer to the current level interrupt saved context is returned in the `intrCtx` argument. The CPU context saved on interrupt contains the volatile general purpose and floating point registers, condition register, Machine Status Register, instruction pointer, link register and the interrupt number itself. On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The interrupt level is zero (not called from an interrupt handler).
-----------------------	---

K_ENOTAVAILABLE There is no context available for the currently handled interrupt.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svIntrAttach	+	+	-	+
CpuIntrOps.unmask	+	+	+	-
CpuIntrOps.mask	+	+	+	-
svIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap(9DKI)` , `svDkiThreadCall(9DKI)` , `imsIntrMask_f(9DKI)`

NAME	svIntrAttach_usparc, svIntrDetach_usparc, svSoftIntrAttach_usparc, svSoftIntrDetach_usparc, svTimerIntrAttach_uparc, svTimerIntrDetach_uparc, svIntrCtxGet_usparc – UltraSPARC interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(unsigned int intrNumb, unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrIdp); void svIntrDetach(CpuIntrId intrId); KnError svSoftIntrAttach(unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svSoftIntrDetach(CpuIntrId intrId); KnError svTimerIntrAttach(CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svTimerIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx);</pre>
FEATURES	DKI
DESCRIPTION	<p>Provides UltraSPARC interrupts management services.</p> <p>An UltraSPARC processor is an implementation of the SPARC-V9 CPU architecture. As specified by the SPARC-V9 architecture, there are fifteen CPU interrupt sources assigned to the fifteen CPU interrupt levels from 1 up to 15. However, the UltraSPARC processor uses these fifteen interrupt levels only for software generated interrupts. The hardware interrupts are delivered to the processor using the "Mondo" interrupt transfer mechanism. The hardware interrupt source is designated by the interrupt number. Typically, the interrupt number is 11-bit width and composed of the interrupt group number (5 MSB) and the interrupt offset number (6 LSB).</p> <p>The "Mondo" interrupt dispatch handler is built into the microkernel. The microkernel handles a mapping between the hardware interrupt numbers and software interrupt levels. When a mondo interrupt packet is received by the microkernel, the interrupt request descriptor is queued and an associated software interrupt is triggered. The microkernel software interrupt handler then dequeues the interrupt request descriptor and invokes a handler associated to the given interrupt number. In this way, a driver interrupt handler is always invoked in the UltraSPARC software interrupt context. This provides the interrupt handler with an environment analogous to the SPARC-V8 one (interrupt levels).</p> <p>The microkernel provides services which allow device drivers to manage UltraSPARC interrupts, mainly to attach/detach handlers to the CPU interrupts.</p> <pre>typedef void (*CpuIntrHandler)(void*);</pre>

svIntrAttach

svIntrAttach attaches a given handler to a given interrupt number at a given processor interrupt level.

The *intrNumb* argument specifies the interrupt number to which to attach.

The *intrLevel* argument specifies the interrupt level to which to attach.

The *intrHandler* argument specifies the handler to call back when the given interrupt occurs.

The *intrCookie* argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.

On success, K_OK is returned and the services defined on the attached interrupt are returned in the *intrOps* parameter. An identifier for the attached interrupt is also returned in *intrId*. This identifier must be used as first argument to subsequent calls to *intrOps* services.

On failure, an error code is returned as follows:

K_EINVAL	The specified interrupt number or interrupt level are invalid.
K_BUSY	Another handler is already attached to the given <i>intrNumb</i> .
K_ENOMEM	The system is out of memory.

Services available on an attached interrupt are defined by the CpuIntrOps structure as follows:

```
typedef struct CpuIntrOps {
    void
    (*mask) (CpuIntrId intrId);

    void
    (*unmask) (CpuIntrId intrId);

    void
    (*enable) (CpuIntrId intrId);

    void
    (*disable) (CpuIntrId intrId);

    void
    (*trigger) (CpuIntrId intrId);
} CpuIntrOps;
```

The `CpuIntrOps.mask` routine disables the interrupt at CPU level identified by `intrId`. In other words, the PIL register is set to the level corresponding to `intrId`. Note that the original value of the PIL register is saved by DKI in order to be restored later by the `CpuIntrOps.unmask` routine.

The `CpuIntrOps.unmask` routine enables the interrupt at CPU level identified by `intrId`.

In other words, the PIL register is restored to the original value saved by the previously called `CpuIntrOps.mask` routine.

The `mask / unmask` pair may be called from base level only and must not be nested. The `mask/unmask` pair is typically used to implement a critical section of code which needs to be protected against the interrupt. Note that with respect to the SPARC-V9 architecture, when an interrupt level `N` is masked, all interrupts with a level less than `N` are also masked. Thus, there is no way to `mask` only one CPU interrupt level except the lowest one.

The `enable / disable` pair may only be called from the attached interrupt handler. When an interrupt occurs, the attached `CpuIntrHandler` is invoked with the interrupt masked at processor level. This behaves in exactly the same way as if `CpuIntrOps.disable` was called just prior to the handler invocation. In other words, the PIL register is set to the interrupt level and the original interrupt processor level (which was when the interrupt occurred) is saved by DKI. Note that the interrupt handler must return to DKI in the same context as it was called, that is with the interrupt disabled at processor level.

On the other hand, the called interrupt handler may use the `enable/disable` pair to allow the interrupt to be nested. This feature is typically used by a host bus driver when the bus interrupts are multiplexed, that is, multiple bus interrupts are reported at the same CPU interrupt level. Typically, an interrupt handler of this type of host bus driver would take the following actions:

- Identify the bus interrupt source (through a PIC or special cycle).
- Disable the bus interrupt source at bus level (through PIC).
- Enable interrupt at processor level (`enable`).
- Call handlers attached to the identified bus interrupt source.
- Disable interrupt at processor level (`disable`).
- Acknowledge (if needed) and enable the bus interrupt source at bus level (through PIC).
- Return to the DKI.

The `CpuIntrOps.trigger` routine allows the interrupt to be triggered by software. Basically, this routine acts like the mondo interrupt dispatcher except the interrupt number is obtained from `intrId` rather than from the mondo interrupt packet. The `CpuIntrOps.trigger` routine is mainly dedicated to the

	software interrupts attached by <code>svSoftIntrAttach</code> . However, it may be also used for hardware interrupts, for instance, for debugging or diagnostic purposes.				
svIntrDetach	<p><code>svIntrDetach</code> detaches an interrupt handler previously connected by <code>svIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svIntrAttach</code> .</p>				
svSoftIntrDetach	<p><code>svSoftIntrAttach</code> attaches a given software interrupt handler to a given processor interrupt level.</p> <p>The <i>intrLevel</i> argument specifies the interrupt level to which to attach.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On success, <code>K_OK</code> is returned and the services defined on the attached interrupt are returned in the <i>intrOps</i> parameter.</p> <p>An identifier for the attached interrupt is also returned in <i>intrId</i> . This identifier must be used as the first argument to subsequent calls to <i>intrOps</i> services.</p> <p>On failure, an error code is returned as follows:</p> <table border="0"> <tr> <td><code>K_EINVAL</code></td> <td>The specified interrupt level is invalid.</td> </tr> <tr> <td><code>K_ENOMEM</code></td> <td>The system is out of memory.</td> </tr> </table> <p>Services available on an attached interrupt are defined by the <code>CpuIntrOps</code> structure described above.</p>	<code>K_EINVAL</code>	The specified interrupt level is invalid.	<code>K_ENOMEM</code>	The system is out of memory.
<code>K_EINVAL</code>	The specified interrupt level is invalid.				
<code>K_ENOMEM</code>	The system is out of memory.				
svSoftIntrDetach	<p><code>svSoftIntrDetach</code> detaches a software interrupt handler previously connected by <code>svSoftIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svSoftIntrAttach</code> .</p>				
svTimerIntrAttach	<p><code>svTimerIntrAttach</code> attaches a given interrupt handler to the UltraSPARC tick-counter interrupt.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque the microkernel.</p>				

On success, `K_OK` is returned and the services defined on the attached interrupt are returned in the `intrOps` parameter. An identifier for the attached interrupt is also returned in `intrId`. This identifier must be used as the first argument to subsequent calls to `intrOps` services.

On failure, an error code is returned as follows:

- `K_BUSY` Another handler is already attached to the tick-counter interrupt.
- `K_ENOMEM` The system is out of memory.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure described above.

svTimerIntrDetach

`svTimerIntrDetach` detaches the interrupt handler previously connected by `svTimerIntrAttach`.

The `intrId` argument identifies the attached interrupt handler, previously returned by `svTimerIntrAttach`.

svIntrCtxGet

`svIntrCtxGet` retrieves the current level interrupt context. It is typically used for profiling purposes.

On success, `K_OK` is returned and a pointer to the recently saved interrupt context is returned in the `intrCtx` argument. The CPU context saved on interrupt has the same structure as the thread context saved on exception or trap. It contains the global registers `%g1-%g7`, the output registers of the interrupted window `%o0-%o7` and the following processor registers: `%tstate`, `%pc`, `%npc`, `%tt`, `%y`. In addition, the thread context contains the number of outstanding windows and the pointer to the outstanding windows buffer if the number of windows is greater than zero.

On failure, an error code is returned as follows:

- `K_EINVAL` Interrupt level is zero (not called from an interrupt handler).
- `K_ENOTAVAILABLE` There is no context available for currently handled interrupt.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each interrupt management service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svIntrAttach</code>	+	+	-	+

svIntrDetach	+	+	-	+
svIntrSoftAttach	+	+	-	+
svIntrSoftDetach	+	+	-	+
svTimerIntrAttach	+	+	-	+
svTimerIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-
CpuIntrOps.mask	+	+	-	-
CpuIntrOps.unmask	+	+	-	-
CpuIntrOps.enable	-	-	+	-
CpuIntrOps.disable	-	-	+	-
CpuIntrOps.trigger	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME	svIntrAttach_x86, svIntrDetach_x86, svIntrCtxGet_x86 – Intel x86 interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(CpuIntr intr, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx); typedef struct CpuIntrOps { void (*enable) (CpuIntrId intrId); void (*disable) (CpuIntrId intrId); } CpuIntrOps;</pre>
FEATURES	DKI
DESCRIPTION	Provides Intel x86 interrupts management services.
EXTENDED DESCRIPTION	<p>For Intel family processors, there are two kinds of interrupts:</p> <p>Exceptions Exceptions are generated and detected by the processor when executing instructions, such as division by 0. These exceptions are classified as faults, traps and aborts. Another type of exception exists: Programmed exceptions generated by INTx instructions.</p> <p>Interrupts Interrupts are generated by events external to the processor, such as requests to service peripheral devices.</p> <p>There are two further types of interrupts:</p> <p>Maskable Maskable interrupts are received on the INTR input pin of the Intel x86 processor. They do not occur if the IF flag of the EFLAGS register is not set.</p> <p>Non Maskable Non maskable interrupts are received on the NMI input pin of the Intel x86 processor. There is no mechanism to prevent non maskable interrupts.</p>

The DKI does not provide support for exceptions, only for interrupts. The processor associates an identifying number with each different interrupt (non maskable and maskable). This number is called a vector in the range from 0-255 with the range of 0-31 reserved for exceptions. The DKI associates an identifying numeric constant with each interrupt.

The following numeric constants are available:

- NMI_INT (non maskable).
- INTR_BASE_NUMBER (maskable)

The DKI provides support for 64 interrupts, ranging from INTR_BASE_NUMBER to INTR_BASE_NUMBER + 64.

The DKI provides services to allow host bus drivers to manage Intel ix86 interrupts, mainly to attach/detach handlers to/from these interrupts.

`svIntrAttach`

`svIntrAttach` attaches a handler to a given CPU interrupt.

The *intr* argument specifies the interrupt to which to attach.

The *intrHandler* argument specifies the handler to call back when the specified interrupt occurs.

The *intrCookie* argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.

On success, `K_OK` is returned and services defined on the attached interrupt are returned in the *intrOps* parameter. An identifier for the attached interrupt is also returned in *intrId*. This identifier must be used as the first parameter to further calls to *intrOps* services.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure as follows:

The `CpuIntrOps.enable` routine enables the interrupt at CPU level identified by *intrId*.

The `CpuIntrOps.disable` routine disables the interrupt at CPU level identified by *intrId*.

Note that `NMI_INT` is a non maskable interrupt. Thus, `enable/disable` have no effect on these interrupts..

Note also that all other interrupts are masked using the same bit on an Intel x86 processor `EFLAGS` register. Thus there is no way to `enable/disable` a single interrupt separately from the others.

When an interrupt occurs, the attached *CpuIntrHandler* is invoked with interrupts masked at processor level. This produces the same effect as if `imsIntrMask_f()` was called just prior to invoking the handler (`imsIntrMaskCount_f` is positive). Two parameters are passed to this handler:

- Provided *intrCookie* parameter
- Detected *intr*

It is up to the called interrupt handler to use `enable` in order to allow or disallow interrupt nesting.

Typically, a host bus driver handler should handle x86 external interrupts as follows:

- Enable interrupts at processor level (`enable`).
- Call handler(s) attached to the identified cpu interrupt.
- Optionally `disable` interrupts in order to perform critical tasks, such as notifying the end an of interrupt to a PIC.
- Return to the DKI.

On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The interrupt specified in <i>intr</i> is invalid (not in the list above). Or, the interrupt specified in <i>intr</i> is valid, but a handler is already attached to it.
-----------------------	--

`svIntrDetach`

`svIntrDetach` detaches an interrupt handler previously connected by `svIntrAttach`.

The *intrId* argument identifies the attached interrupt handler, and was previously returned by `svIntrAttach`.

`svIntrCtxGet`

`svIntrCtxGet` retrieves the current level interrupt context. It is mainly to be used for profiling purposes.

On success, `K_OK` is returned and a pointer to the current level of interrupt saved context is returned in the *intrCtx* argument. The CPU context saved on interrupt contains the interrupted PC composed of the `eip` and `cs` registers, the `EFLAGS` register, and the stack pointer.

On failure, an error code is returned as follows:

K_EINVAL Interrupt level is zero (not called from an interrupt handler).
 K_ENOTAVAILABLE There is no context available for the currently handled interrupt.

Intel x86 Interrupts Management Allowed Calling Contexts

The table below specifies the contexts in which a caller is allowed to invoke each services.

Services	Base level	DKI thread	Interrupt	Blocking
svIntrAttach	+	+	-	+
CpuIntrOps.enable	+	+	+	-
CpuIntrOps.disable	+	+	+	-
svIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME	svMemAlloc, svMemFree – A general purpose memory allocator															
SYNOPSIS	<pre>#include <dki/dki.h> void * svMemAlloc(unsigned int size); void * svMemFree(void * mem, unsigned int size);</pre>															
FEATURES	DKI															
DESCRIPTION	Provides general purpose memory allocation.															
EXTENDED DESCRIPTION	<p>The microkernel provides general purpose memory management services for device drivers that need to allocate and free sized pieces of memory dynamically in a supervisor address space. As the initialization scheme is mainly dynamic, device drivers need to dynamically allocate/free small pieces of supervisor data. Typically, a device driver needs to dynamically allocate data associated with each instance to be registered in the Device Registry at initialization time.</p> <p>Moreover, most of the DDI services called from base level by the driver clients leads to dynamic allocation/freeing of a number of linked list elements for internal management purposes. Note that the memory allocated using these services is anonymous; it is not associated with any actor context. For this reason, all the memory allocated must be explicitly freed by the drivers before they terminate, as the kernel cannot do it at actor deletion time.</p> <p>svMemAlloc allocates memory of a given size which is accessible within the supervisor address space. The size argument specifies the required memory size. In case of success, svMemAlloc returns a pointer to the supervisor memory allocated, otherwise a NULL pointer is returned. The allocated memory returned by svMemAlloc is suitably aligned for any use.</p> <p>The svMemFree routine frees memory previously allocated using svMemAlloc. The mem argument specifies the pointer to the memory previously returned by svMemAlloc. The size argument must be exactly the same as the size required when previously allocating the memory using svMemAlloc.</p>															
Allowed Calling Contexts	<p>The following table specifies the contexts in which a caller is allowed to invoke each service.</p> <table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Services</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Base level</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">DKI thread</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Interrupt</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Blocking</th> </tr> </thead> <tbody> <tr> <td style="border-bottom: 1px solid black;">svMemAlloc</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">-</td> <td style="border-bottom: 1px solid black;">+</td> </tr> <tr> <td style="border-bottom: 1px solid black;">svMemFree</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">-</td> <td style="border-bottom: 1px solid black;">+</td> </tr> </tbody> </table>	Services	Base level	DKI thread	Interrupt	Blocking	svMemAlloc	+	+	-	+	svMemFree	+	+	-	+
Services	Base level	DKI thread	Interrupt	Blocking												
svMemAlloc	+	+	-	+												
svMemFree	+	+	-	+												
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:															

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svMemAlloc, svMemFree – A general purpose memory allocator															
SYNOPSIS	<pre>#include <dki/dki.h> void * svMemAlloc(unsigned int size); void * svMemFree(void * mem, unsigned int size);</pre>															
FEATURES	DKI															
DESCRIPTION	Provides general purpose memory allocation.															
EXTENDED DESCRIPTION	<p>The microkernel provides general purpose memory management services for device drivers that need to allocate and free sized pieces of memory dynamically in a supervisor address space. As the initialization scheme is mainly dynamic, device drivers need to dynamically allocate/free small pieces of supervisor data. Typically, a device driver needs to dynamically allocate data associated with each instance to be registered in the Device Registry at initialization time.</p> <p>Moreover, most of the DDI services called from base level by the driver clients leads to dynamic allocation/freeing of a number of linked list elements for internal management purposes. Note that the memory allocated using these services is anonymous; it is not associated with any actor context. For this reason, all the memory allocated must be explicitly freed by the drivers before they terminate, as the kernel cannot do it at actor deletion time.</p> <p>svMemAlloc allocates memory of a given size which is accessible within the supervisor address space. The size argument specifies the required memory size. In case of success, svMemAlloc returns a pointer to the supervisor memory allocated, otherwise a NULL pointer is returned. The allocated memory returned by svMemAlloc is suitably aligned for any use.</p> <p>The svMemFree routine frees memory previously allocated using svMemAlloc. The mem argument specifies the pointer to the memory previously returned by svMemAlloc. The size argument must be exactly the same as the size required when previously allocating the memory using svMemAlloc.</p>															
Allowed Calling Contexts	<p>The following table specifies the contexts in which a caller is allowed to invoke each service.</p> <table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Services</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Base level</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">DKI thread</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Interrupt</th> <th style="border-top: 1px solid black; border-bottom: 1px solid black;">Blocking</th> </tr> </thead> <tbody> <tr> <td style="border-bottom: 1px solid black;">svMemAlloc</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">-</td> <td style="border-bottom: 1px solid black;">+</td> </tr> <tr> <td style="border-bottom: 1px solid black;">svMemFree</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">+</td> <td style="border-bottom: 1px solid black;">-</td> <td style="border-bottom: 1px solid black;">+</td> </tr> </tbody> </table>	Services	Base level	DKI thread	Interrupt	Blocking	svMemAlloc	+	+	-	+	svMemFree	+	+	-	+
Services	Base level	DKI thread	Interrupt	Blocking												
svMemAlloc	+	+	-	+												
svMemFree	+	+	-	+												
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:															

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svPhysAlloc, svPhysFree – A special purpose physical memory allocator
SYNOPSIS	<pre>#include <dki/dki.h> KnError svPhysAlloc(KnPhMemChunk * chunk, KnPhMemAlign * align); void * svPhysFree(KnPhMemChunk * chunk); typedef struct { PhAddr paddr; PhSize psize; VmAddr vaddr; } KnPhMemChunk; typedef struct { PhAddr alignment; PhAddr floating; } KnPhMemAlign;</pre>
FEATURES	DKI
DESCRIPTION	Provides special purpose physical memory allocation.
EXTENDED DESCRIPTION	<p>In order to satisfy all physical memory constraints imposed by the different I/O buses, mainly for DMA purposes, the DKI provides an interface to allocate and free special purpose physical memory that satisfies the given constraints.</p> <p>Typically, different I/O buses may impose different constraints on the memory used by their devices for Direct Memory Access, such as alignment, specific boundary crossing, maximum size, or specific locations within the physical memory space.</p> <p>svPhysAlloc allocates an amount of contiguous physical memory that satisfies the required constraints. The <i>chunk</i> argument points to the KnPhMemChunk structure. The <i>psize</i> field is an input argument specifying the size in bytes of the memory to be allocated. The <i>paddr</i> field is an input/output argument specifying the physical start address of the allocated memory chunk.</p> <p>The <i>vaddr</i> field is not used by svPhysAlloc.</p> <p>The <i>align</i> argument points to the KnPhMemAlign structure.</p> <p>The KnPhMemAlign structure specifies the constraints on the memory chunk being allocated.</p> <p>The <i>alignment</i> field is a mask which specifies constraints on the start address of the memory chunk being allocated. A bit set within <i>alignment</i> specifies that the corresponding bit within the start address may take any value, i.e. there are</p>

no specific constraints on that bit. A bit cleared within *alignment* specifies that the corresponding bit within the start address must take the same value as the corresponding bit of the *paddr* field. This mask allows the caller to specify an alignment for the start of the allocated memory, by zeroing the required number of least significant bits. By resetting the required number of most significant bits, the caller may also indicate which part of the physical space the memory should be allocated to.

The *floating* field is a mask which indicates which bits of the returned address can vary while walking through the allocated memory for the required size. In other words, bits cleared in the mask must be constant for all addresses in the range of the allocated memory. This mask may be used to specify that the amount of memory allocated must not span across a given address boundary.

Note that the *align* argument may be set to `NULL` specifying that there are no constraints on the memory chunks allocated. In case of success, the starting address of the allocated memory is returned in the *paddr* field and the function returns `K_OK`. The returned address is aligned on a page boundary and is therefore suitably aligned to be used subsequently in mapping services. The size of effectively allocated memory is rounded up to the next page boundary.

On failure, an error code is returned as follows:

`K_ENOMEM` There is no contiguous physical memory which satisfies the requirements specified by both significant and floating arguments.

`K_ESIZE` The size argument is equal to zero.

`svPhysFree` releases physical memory previously allocated with `svPhysAlloc`.

The *chunk* argument points to the `KnPhMemChunk` structure which specifies the physical chunk start addresses and the chunk size.

Note - The *psize* field must have the value previously specified in `svPhysAlloc`. The *paddr* field must have the value previously returned by `svPhysAlloc`.

Typically, a driver uses the same `KnPhMemChunk` structure for the `svPhysAlloc` and `svPhysFree` calls and the structure fields are not modified by the driver once `svPhysAlloc` has been done.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysAlloc	+	+	-	+
svPhysFree	+	+	-	+

EXAMPLES

As an example, to allocate a 16KB contiguous memory region in the first Megabyte of physical space, that is aligned on a 4KB address, which does not span a 64KB boundary, the call should look like:

```
KnPhMemChunk chunk;
KnPhMemAlign align;

chunk.psize      = 0x00004000;          /* psize      = 16KB */
align.alignment  = 0x000ff000;          /* alignment  = (1024KB-1) & ~(4KB-1) */
align.floating   = 0x0000ffff;          /* floating   = 64KB-1 */

res = svPhysAlloc(&chunk, &align);
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svPhysAlloc, svPhysFree – A special purpose physical memory allocator
SYNOPSIS	<pre>#include <dki/dki.h> KnError svPhysAlloc(KnPhMemChunk * chunk, KnPhMemAlign * align); void * svPhysFree(KnPhMemChunk * chunk); typedef struct { PhAddr paddr; PhSize psize; VmAddr vaddr; } KnPhMemChunk; typedef struct { PhAddr alignment; PhAddr floating; } KnPhMemAlign;</pre>
FEATURES	DKI
DESCRIPTION	Provides special purpose physical memory allocation.
EXTENDED DESCRIPTION	<p>In order to satisfy all physical memory constraints imposed by the different I/O buses, mainly for DMA purposes, the DKI provides an interface to allocate and free special purpose physical memory that satisfies the given constraints.</p> <p>Typically, different I/O buses may impose different constraints on the memory used by their devices for Direct Memory Access, such as alignment, specific boundary crossing, maximum size, or specific locations within the physical memory space.</p> <p>svPhysAlloc allocates an amount of contiguous physical memory that satisfies the required constraints. The <i>chunk</i> argument points to the KnPhMemChunk structure. The <i>psize</i> field is an input argument specifying the size in bytes of the memory to be allocated. The <i>paddr</i> field is an input/output argument specifying the physical start address of the allocated memory chunk.</p> <p>The <i>vaddr</i> field is not used by svPhysAlloc.</p> <p>The <i>align</i> argument points to the KnPhMemAlign structure.</p> <p>The KnPhMemAlign structure specifies the constraints on the memory chunk being allocated.</p> <p>The <i>alignment</i> field is a mask which specifies constraints on the start address of the memory chunk being allocated. A bit set within <i>alignment</i> specifies that the corresponding bit within the start address may take any value, i.e. there are</p>

no specific constraints on that bit. A bit cleared within *alignment* specifies that the corresponding bit within the start address must take the same value as the corresponding bit of the *paddr* field. This mask allows the caller to specify an alignment for the start of the allocated memory, by zeroing the required number of least significant bits. By resetting the required number of most significant bits, the caller may also indicate which part of the physical space the memory should be allocated to.

The *floating* field is a mask which indicates which bits of the returned address can vary while walking through the allocated memory for the required size. In other words, bits cleared in the mask must be constant for all addresses in the range of the allocated memory. This mask may be used to specify that the amount of memory allocated must not span across a given address boundary.

Note that the *align* argument may be set to `NULL` specifying that there are no constraints on the memory chunks allocated. In case of success, the starting address of the allocated memory is returned in the *paddr* field and the function returns `K_OK`. The returned address is aligned on a page boundary and is therefore suitably aligned to be used subsequently in mapping services. The size of effectively allocated memory is rounded up to the next page boundary.

On failure, an error code is returned as follows:

`K_ENOMEM` There is no contiguous physical memory which satisfies the requirements specified by both significant and floating arguments.

`K_ESIZE` The size argument is equal to zero.

`svPhysFree` releases physical memory previously allocated with `svPhysAlloc`.

The *chunk* argument points to the `KnPhMemChunk` structure which specifies the physical chunk start addresses and the chunk size.

Note - The *psize* field must have the value previously specified in `svPhysAlloc`. The *paddr* field must have the value previously returned by `svPhysAlloc`.

Typically, a driver uses the same `KnPhMemChunk` structure for the `svPhysAlloc` and `svPhysFree` calls and the structure fields are not modified by the driver once `svPhysAlloc` has been done.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysAlloc	+	+	-	+
svPhysFree	+	+	-	+

EXAMPLES

As an example, to allocate a 16KB contiguous memory region in the first Megabyte of physical space, that is aligned on a 4KB address, which does not span a 64KB boundary, the call should look like:

```
KnPhMemChunk chunk;
KnPhMemAlign align;

chunk.psize      = 0x00004000;          /* psize      = 16KB   */
align.alignment  = 0x000ff000;          /* alignment  = (1024KB-1) & ~(4KB-1) */
align.floating   = 0x0000ffff;          /* floating   = 64KB-1 */

res = svPhysAlloc(&chunk, &align);
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svPhysMap, svPhysUnmap, vmMapToPhys – physical to virtual memory mapping
FEATURES	DKI
DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. This services should be used mainly by primary bus drivers to map bus I/O space, or DMA memory.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ svPhysMap_usparc(9DKI) ■ svPhysUnmap_usparc(9DKI) ■ vmMapToPhys_usparc(9DKI) ■ svPhysMap_powerpc(9DKI) ■ svPhysUnmap_powerpc(9DKI) ■ vmMapToPhys_powerpc(9DKI) ■ svPhysMap_x86(9DKI) ■ svPhysUnmap_x86(9DKI) ■ vmMapToPhys_x86(9DKI)
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p>

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svPhysMap_powerpc, svPhysUnmap_powerpc, vmMapToPhys_powerpc – PowerPC physical to virtual memory mapping						
Synopsis	<pre>#include <dki/dki.h> KnError svPhysMap(KnPhMemChunk * chunk, PteCntlBits cntlBits); void * svPhysUnmap(KnPhMemChunk * chunk); KnError vmMapToPhys(KnCap * actor, KnPhMemChunk * chunk, PteCntlBits cntlBits); typedef struct { PhAddr paddr; /* physical start addree */ PhSize psize; /* size */ VmAddr vaddr; /* virtual start address */ } KnPhMemChunk;</pre>						
FEATURES	DKI						
DESCRIPTION	Provides physical to virtual memory mapping services.						
EXTENDED DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. These services should be used mainly by primary bus drivers to map bus I/O space or DMA memory.</p> <p>svPhysMap allocates a range of virtual addresses within the supervisor address space and maps it to a given range of physical addresses.</p> <p>The <i>chunk</i> argument points to the KnPhMemChunk structure.</p> <p>The <i>paddr</i> and <i>psize</i> fields specify the chunk of physical memory being mapped to.</p> <p>The <i>vaddr</i> field is used by svPhysMap to return the virtual address to which the physical one is mapped.</p> <p>For PowerPC family processors, the <i>cntlBits</i> argument should be constructed by 'oring' the following values:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top;">PTE_READ_WRITE</td> <td>The mapped memory allows you to perform write accesses.</td> </tr> <tr> <td style="vertical-align: top;">PTE_CACHE_WRITE_THROUGH</td> <td>The memory is mapped with Write-Through attribute (bit <i>w</i>). Thus a store operation updates the cache if necessary, and in addition the update is written to memory.</td> </tr> <tr> <td style="vertical-align: top;">PTE_CACHE_DISABLE</td> <td>The memory is mapped with Cache-Inhibited attribute (bit <i>I</i>). Thus</td> </tr> </table>	PTE_READ_WRITE	The mapped memory allows you to perform write accesses.	PTE_CACHE_WRITE_THROUGH	The memory is mapped with Write-Through attribute (bit <i>w</i>). Thus a store operation updates the cache if necessary, and in addition the update is written to memory.	PTE_CACHE_DISABLE	The memory is mapped with Cache-Inhibited attribute (bit <i>I</i>). Thus
PTE_READ_WRITE	The mapped memory allows you to perform write accesses.						
PTE_CACHE_WRITE_THROUGH	The memory is mapped with Write-Through attribute (bit <i>w</i>). Thus a store operation updates the cache if necessary, and in addition the update is written to memory.						
PTE_CACHE_DISABLE	The memory is mapped with Cache-Inhibited attribute (bit <i>I</i>). Thus						

	accesses are made to main memory, bypassing the caches.
PTE_MEMORY_COHERENCY	The memory is mapped with Memory-Coherency attribute (bit M). This enforces coherency of memory shared between processors in a system. When performing an access to memory, there is a hardware indication to the rest of the system that the access is global. Other processors affected by the access must then respond to this global access. Typically, this is used for a snooping bus design.
PTE_MEMORY_GUARDED	The memory is mapped with Memory-Guarded attribute (bit G). This prevents the processor from making out-of-order access to that memory (that is, access not directly dictated by the program). This may be useful if there are holes in physical memory, or to prevent these accesses to certain peripheral devices.

If *cntlBits* is equal to zero, the memory is mapped as read-only, not guarded, with cache enabled in write-back mode, and no coherency is enforced.

Note - Any combination where PTE_CACHE_WRITE_THROUGH and PTE_CACHE_DISABLE are both set is not supported.

On success K_OK is returned, otherwise a negative error code is returned:

K_ENOMEM	The system is out of memory.
K_ESIZE	The <i>psize</i> argument is equal to zero.

`svPhysUnmap` unmaps the physical memory chunk previously mapped by `svPhysMap`. The *chunk* argument points to the `KnPhMemChunk` structure which specifies the virtual and physical chunk start addresses and the chunk size. Note that the *paddr* and *psize* fields must have the values previously specified in `svPhysMap`. The *vaddr* field must have the value previously returned by `svPhysMap`.

Note - Typically, a driver uses the same `KnPhMemChunk` structure for the `svPhysMap` and `svPhysUnmap` calls and the structure fields are not modified by the driver once `svPhysMap` has been done.

`vmMapToPhys` maps a given physical memory chunk to the target actor address space. The *actor* argument specifies the target actor capability.

If *actor* is `K_MYACTOR`, the address space of the current actor is used.

If *actor* is `K_SVACTOR`, the supervisor address space is used.

The *chunk* argument points to the `KnPhMemChunk` structure which specifies the physical and virtual chunk start addresses and the chunk size.

The specified virtual address range must be allocated using the `K_RESERVED` option (see `rgnAllocate(2)`) prior to the invocation of `vmMapToPhys`.

The specified actor can be a supervisor actor as well as a user actor.

The mapping produced by `vmMapToPhys` can only be invalidated by `rgnFree`. *vaddr*, *paddr* and *psize* must be page-aligned.

For PowerPC family processors, the *cntlBits* is defined in the same way as for the `svPhysMap` routine above.

On success `K_OK` is returned, otherwise a negative error code is returned:

<code>K_EFAULT</code>	The <i>actor</i> argument points outside the caller's address space.
<code>K_EINVAL</code>	An inconsistent actor capability was provided.
<code>K_EUNKNOWN</code>	<i>actorcap</i> does not specify a reachable actor.
<code>K_ENOMEM</code>	The system is out of memory.
<code>K_ESIZE</code>	The <i>psize</i> argument is equal to zero.
<code>K_EROUND</code>	<i>vaddr</i> or <i>paddr</i> or <i>psize</i> is not page-aligned.
<code>K_EADDR</code>	Some or all addresses from the target virtual address range are out of a region allocated with the <code>K_RESERVED</code> option.

Note - For performance reasons, the current implementation does not guarantee that any attempt to map a physical address to a virtual address out of a region allocated with `K_RESERVED` option would effectively produce a `K_EADDR` error.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysMap	+	+	-	+
svPhysUnmap	+	+	-	+
vmMapToPhys	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svPhysAlloc(9DKI)` , `svDkiThreadCall(9DKI)` , `rgnAllocate(2K)`

NAME	svPhysMap_usparc, svPhysUnmap_usparc, vmMapToPhys_usparc – UltraSPARC physical to virtual memory mapping
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svPhysMap(KnPhMemChunk * chunk, KnPteAttr attr); void svPhysUnmap(KnPhMemChunk * chunk); KnError vmMapToPhys(KnCap * actor, KnPhMemChunk * chunk, KnPteAttr attr);</pre>
FEATURES	DKI
DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. These services should be used mainly by a host bus driver to map bus I/O space or DMA memory.</p> <p>svPhysMap allocates a range of virtual addresses within the supervisor address space and maps it to a given range of physical addresses.</p> <p>The <i>chunk</i> argument points to the KnPhMemChunk structure shown below:</p> <pre>typedef struct { PhAddr paddr; /* physical start address */ PhSize psize; /* size */ VmAddr vaddr; /* virtual start address */ } KnPhMemChunk;</pre> <p>The <i>paddr</i> and <i>psize</i> fields specify the chunk of physical memory being mapped to. The <i>vaddr</i> field is used by svPhysMap to return the virtual address to which the physical one is mapped.</p> <p>The <i>attr</i> argument specifies the mapping attributes. <i>attr</i> is a bit-mask composed of two independent parts:</p> <ul style="list-style-type: none"> Data access attributes Instruction access attributes <p>Basically, each part of the attributes is a sub-set of bits defined by the translation table entry (TTE) of the UltraSPARC MMU.</p> <p>A combination of the following attributes may be specified for data access:</p> <p>PTE_DATTR_G — global The PTE_DATTR_G bit set allows the mapping to be shared among all (user and supervisor) contexts.</p> <p>PTE_DATTR_W — writable The PTE_DATTR_W bit set grants write permission for the mapping.</p> <p>PTE_DATTR_P — privileged</p>

The `PTE_DATTR_P` bit set restricts access to the mapping for the supervisor only.

`PTE_DATTR_E` — side-effect

The `PTE_DATTR_E` bit set makes noncacheable memory accesses to be strongly ordered against other E-bit accesses, and noncacheable stores are not merged. This bit should be set for I/O devices having side-effects. Note that the E-bit does not force an uncacheable access. It is expected that the `PTE_DATTR_CV` and `PTE_DATTR_CP` bits will be set to zero when the E-bit is set.

`PTE_DATTR_CV` — L1-cacheable

The `PTE_DATTR_CV` bit set allows data to be cached in the (L1) CPU data cache. Note that if the `PTE_DATTR_CV` bit is set, the `PTE_DATTR_CP` bit must also be set.

`PTE_DATTR_CP` — L2-cacheable

The `PTE_DATTR_CP` bit set allows data to be cached in the (L2) external cache.

`PTE_DATTR_IE` — invert endianness

The `PTE_DATTR_IE` bit set causes data accesses to the mapping to be processed with inverse endianness from that specified by the instruction.

`PTE_DATTR_V` — valid

The `PTE_DATTR_V` bit set enables data accesses to the mapping. If this bit is not set, all other bits (described above) are ignored and a data access to the mapping will result in a data access exception.

Combinations of the following attributes may be specified for instruction access:

`PTE_IATTR_G` — global

The `PTE_IATTR_G` bit set allows the mapping to be shared among all (user and supervisor) contexts.

`PTE_IATTR_P` — privileged

The `PTE_IATTR_P` bit set restricts access to the mapping to the supervisor only.

`PTE_IATTR_CV` — L1-cacheable

The `PTE_IATTR_CV` bit set allows instructions to be cached in the (L1) CPU instruction cache. Note that if the `PTE_IATTR_CV` bit is set, the `PTE_IATTR_CP` bit must also be set.

`PTE_IATTR_CP` — L2-cacheable

The `PTE_IATTR_CP` bit set allows instructions to be cached in the (L2) external cache.

PTE_IATTR_V — valid

The PTE_IATTR_V bit set enables instructions to be obtained from the mapping. If this bit is not set, all other bits (described above) are ignored and an instruction to fetch from the mapping will result in an instruction access exception.

On success svPhysMap returns K_OK , otherwise a negative error code is returned as follows:

K_ENOMEM The system is out of memory.

K_ESIZE The *psize* argument is equal to zero.

svPhysUnmap

svPhysUnmap unmaps the physical memory chunk previously mapped by svPhysMap . The *chunk* argument points to the KnPhMemChunk structure which specifies the virtual and physical chunk start addresses and the chunk size. Note that the *paddr* and *psize* fields must have the values previously specified in svPhysMap . The *vaddr* fields must have the value previously returned by svPhysMap . Typically, a driver uses the same KnPhMemChunk structure for the svPhysMap and svPhysUnmap calls and the structure fields are not modified by the driver once svPhysMap has been performed.

vmMapToPhys

vmMapToPhys maps a given physical memory chunk to the target actor address space.

The *actor* argument specifies the target actor capability. If *actor* is K_MYACTOR , the address space of the current actor is used. If *actor* is K_SVACTOR , the supervisor address space is used.

The *chunk* argument points to the KnPhMemChunk structure which specifies the physical and virtual chunk start addresses and the chunk size. The specified virtual address range must be allocated using the K_RESERVED option (see *rgnAllocate(2K)* prior to the invocation of vmMapToPhys . The specified actor can be a supervisor actor as well as an user actor. The mapping produced by vmMapToPhys can only be invalidated by *rgnFree* .

vaddr , *paddr* and *psize* must be page-aligned.

For UltraSPARC family processors, the *attr* argument is defined in the same way as for the svPhysMap routine above.

On success K_OK is returned, otherwise a negative error code is returned as follows:

K_EFAULT The *actor* argument points outside of the caller's address space.

K_EINVAL	An inconsistent actor capability was specified.
K_EUNKNOWN	<i>actorcap</i> does not specify a reachable actor.
K_ENOMEM	The system is out of memory.
K_ESIZE	The <i>psize</i> argument is equal to zero.
K_EROUND	<i>vaddr</i> or <i>paddr</i> or <i>psize</i> is not page-aligned.
K_EADDR	Some or all addresses from the target virtual address range are out of a region allocated with the <code>K_RESERVED</code> option. For performance reasons, the current implementation does not guarantee that any attempt to map a physical address to a virtual address outside a region allocated with <code>K_RESERVED</code> option would effectively produce a <code>K_EADDR</code> error.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysMap	+	+	-	+
svPhysUnmap	+	+	-	+
vmMapToPhys	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svPhysAlloc(9DKI)` , `svDkiThreadCall(9DKI)` , `rgnAllocate(2K)`

NAME	svPhysMap_x86, svPhysUnmap_x86, vmMapToPhys_x86 – Intel x86 physical to virtual memory mapping								
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svPhysMap(KnPhMemChunk * chunk, PteCntlBits cntlBits); void svPhysUnmap(KnPhMemChunk * chunk); KnError vmMapToPhys(KnCap * actor, KnPhMemChunk * chunk, PteCntlBits cntlBits); typedef struct { PhAddr paddr; /* physical start address */ PhSize psize; /* size */ VmAddr vaddr; /* virtual start address */ } KnPhMemChunk;</pre>								
FEATURES	DKI								
EXTENDED DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. These services should be used mainly by primary bus drivers to map bus I/O space, or DMA memory.</p> <p>svPhysMap</p> <p>svPhysMap allocates a range of virtual addresses within the supervisor address space and maps it to a given range of physical addresses.</p> <p>The <i>chunk</i> argument points to the KnPhMemChunk structure.</p> <p>The <i>paddr</i> and <i>psize</i> fields specify the chunk of physical memory being mapped to.</p> <p>The <i>vaddr</i> field is used by svPhysMap to return the virtual address to which the physical one is mapped.</p> <p>For Intel ix86 family processor, the cntlBits argument is a direct mapping of the Page Control bits of Page Table Entries, as described in the Intel 386/486 Programming Reference Manual. Therefore, the cntlBits argument should be constructed by “oring” the following values:</p> <table border="0"> <tr> <td style="padding-right: 20px;">PTE_PRESENT</td> <td>The mapped memory must be present in physical memory.</td> </tr> <tr> <td>PTE_READ_WRITE</td> <td>The mapped memory allows write accesses to be performed.</td> </tr> <tr> <td>PTE_USER_SUPERVISOR</td> <td>The mapped memory can be used in user space.</td> </tr> <tr> <td>PTE_WRITE_TRANSPARENT</td> <td>The memory is mapped with the Write-Through attribute (bit PWT). Thus, a store operation updates the cache if necessary, and in addition the update is written to memory.</td> </tr> </table>	PTE_PRESENT	The mapped memory must be present in physical memory.	PTE_READ_WRITE	The mapped memory allows write accesses to be performed.	PTE_USER_SUPERVISOR	The mapped memory can be used in user space.	PTE_WRITE_TRANSPARENT	The memory is mapped with the Write-Through attribute (bit PWT). Thus, a store operation updates the cache if necessary, and in addition the update is written to memory.
PTE_PRESENT	The mapped memory must be present in physical memory.								
PTE_READ_WRITE	The mapped memory allows write accesses to be performed.								
PTE_USER_SUPERVISOR	The mapped memory can be used in user space.								
PTE_WRITE_TRANSPARENT	The memory is mapped with the Write-Through attribute (bit PWT). Thus, a store operation updates the cache if necessary, and in addition the update is written to memory.								

PTE_CACHE_DISABLE The memory is mapped with caching disabled (bit PCD). Thus, memory accesses go to main memory, bypassing the caches.

On success K_OK is returned, otherwise a negative error code is returned:

K_ENOMEM The system is out of memory.

K_ESIZE The *psize* argument is equal to zero.

svPhysUnmap

svPhysUnmap unmaps the physical memory chunk previously mapped by svPhysMap . The *chunk* argument points to the KnPhMemChunk structure which specifies the virtual and physical chunk start addresses and the chunk size. Note that the *paddr* and *psize* fields must have the values previously specified in svPhysMap . The *vaddr* field must have the value previously returned by svPhysMap (typically, a driver uses the same KnPhMemChunk structure for the svPhysMap and svPhysUnmap calls and the structure fields are not modified by the driver once svPhysMap is done).

vmMapToPhys

vmMapToPhys maps a given physical memory chunk to the target actor address space.

The *actor* argument specifies the target actor capability. If *actor* is K_MYACTOR , the address space of the current actor is used. If *actor* is K_SVACTOR , the supervisor address space is used.

The *chunk* argument points to the KnPhMemChunk structure which specifies the physical and virtual chunk start addresses and the chunk size. The specified virtual address range must be allocated using the K_RESERVED option (see *rgnAllocate* (2K)) prior to the invocation of vmMapToPhys . The specified actor can be a supervisor actor as well as a user actor. The mapping produced by vmMapToPhys can only be invalidated by rgnFree .

vaddr , *paddr* and *psize* must be page-aligned.

For Intel x86 family processors, the *cntlBits* is defined in the same way as for the svPhysMap routine above.

On success K_OK is returned, otherwise a negative error code is returned:

K_EFAULT The *actor* argument points to the outside of the caller's address space.

K_EINVAL An inconsistent actor capability was provided.

K_EUNKNOWN *actorcap* does not specify a reachable actor.

K_ENOMEM The system is out of memory.

K_ESIZE The *psize* argument is equal to zero.

K_EROUND *vaddr* or *paddr* or *psize* is not page-aligned.

K_EADDR Some or all addresses from the target virtual address range are out of a region allocated with the K_RESERVED option (for performance reasons the current implementation does not guarantee that any attempt to map a physical address to a virtual address out of a region allocated with K_RESERVED option would effectively produce a K_EADDR error).

Intel x86 Memory Mapping Allowed Calling Contexts

The table below specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysMap	+	+	-	+
svPhysUnmap	+	+	-	+
vmMapToPhys	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svPhysAlloc(9DKI) , *svDkiThreadCall(9DKI)* , *rgnAllocate(9DKI)*

NAME	svPhysMap, svPhysUnmap, vmMapToPhys – physical to virtual memory mapping				
FEATURES	DKI				
DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. This services should be used mainly by primary bus drivers to map bus I/O space, or DMA memory.</p> <p>See the architecture specific man pages:</p> <ul style="list-style-type: none"> ■ svPhysMap_usparc(9DKI) ■ svPhysUnmap_usparc(9DKI) ■ vmMapToPhys_usparc(9DKI) ■ svPhysMap_powerpc(9DKI) ■ svPhysUnmap_powerpc(9DKI) ■ vmMapToPhys_powerpc(9DKI) ■ svPhysMap_x86(9DKI) ■ svPhysUnmap_x86(9DKI) ■ vmMapToPhys_x86(9DKI) 				
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">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				

NAME	svPhysMap_powerpc, svPhysUnmap_powerpc, vmMapToPhys_powerpc – PowerPC physical to virtual memory mapping						
Synopsis	<pre>#include <dki/dki.h> KnError svPhysMap(KnPhMemChunk * chunk, PteCntlBits cntlBits); void * svPhysUnmap(KnPhMemChunk * chunk); KnError vmMapToPhys(KnCap * actor, KnPhMemChunk * chunk, PteCntlBits cntlBits); typedef struct { PhAddr paddr; /* physical start addree */ PhSize psize; /* size */ VmAddr vaddr; /* virtual start address */ } KnPhMemChunk;</pre>						
FEATURES	DKI						
DESCRIPTION	Provides physical to virtual memory mapping services.						
EXTENDED DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. These services should be used mainly by primary bus drivers to map bus I/O space or DMA memory.</p> <p>svPhysMap allocates a range of virtual addresses within the supervisor address space and maps it to a given range of physical addresses.</p> <p>The <i>chunk</i> argument points to the KnPhMemChunk structure.</p> <p>The <i>paddr</i> and <i>psize</i> fields specify the chunk of physical memory being mapped to.</p> <p>The <i>vaddr</i> field is used by svPhysMap to return the virtual address to which the physical one is mapped.</p> <p>For PowerPC family processors, the <i>cntlBits</i> argument should be constructed by 'oring' the following values:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-left: 2em;">PTE_READ_WRITE</td> <td>The mapped memory allows you to perform write accesses.</td> </tr> <tr> <td style="padding-left: 2em;">PTE_CACHE_WRITE_THROUGH</td> <td>The memory is mapped with Write-Through attribute (bit <i>w</i>). Thus a store operation updates the cache if necessary, and in addition the update is written to memory.</td> </tr> <tr> <td style="padding-left: 2em;">PTE_CACHE_DISABLE</td> <td>The memory is mapped with Cache-Inhibited attribute (bit <i>I</i>). Thus</td> </tr> </table>	PTE_READ_WRITE	The mapped memory allows you to perform write accesses.	PTE_CACHE_WRITE_THROUGH	The memory is mapped with Write-Through attribute (bit <i>w</i>). Thus a store operation updates the cache if necessary, and in addition the update is written to memory.	PTE_CACHE_DISABLE	The memory is mapped with Cache-Inhibited attribute (bit <i>I</i>). Thus
PTE_READ_WRITE	The mapped memory allows you to perform write accesses.						
PTE_CACHE_WRITE_THROUGH	The memory is mapped with Write-Through attribute (bit <i>w</i>). Thus a store operation updates the cache if necessary, and in addition the update is written to memory.						
PTE_CACHE_DISABLE	The memory is mapped with Cache-Inhibited attribute (bit <i>I</i>). Thus						

	accesses are made to main memory, bypassing the caches.
PTE_MEMORY_COHERENCY	The memory is mapped with Memory-Coherency attribute (bit M). This enforces coherency of memory shared between processors in a system. When performing an access to memory, there is a hardware indication to the rest of the system that the access is global. Other processors affected by the access must then respond to this global access. Typically, this is used for a snooping bus design.
PTE_MEMORY_GUARDED	The memory is mapped with Memory-Guarded attribute (bit G). This prevents the processor from making out-of-order access to that memory (that is, access not directly dictated by the program). This may be useful if there are holes in physical memory, or to prevent these accesses to certain peripheral devices.

If *cntlBits* is equal to zero, the memory is mapped as read-only, not guarded, with cache enabled in write-back mode, and no coherency is enforced.

Note - Any combination where PTE_CACHE_WRITE_THROUGH and PTE_CACHE_DISABLE are both set is not supported.

On success K_OK is returned, otherwise a negative error code is returned:

K_ENOMEM	The system is out of memory.
K_ESIZE	The <i>psize</i> argument is equal to zero.

`svPhysUnmap` unmaps the physical memory chunk previously mapped by `svPhysMap`. The *chunk* argument points to the `KnPhMemChunk` structure which specifies the virtual and physical chunk start addresses and the chunk size. Note that the *paddr* and *psize* fields must have the values previously specified in `svPhysMap`. The *vaddr* field must have the value previously returned by `svPhysMap`.

Note - Typically, a driver uses the same `KnPhMemChunk` structure for the `svPhysMap` and `svPhysUnmap` calls and the structure fields are not modified by the driver once `svPhysMap` has been done.

`vmMapToPhys` maps a given physical memory chunk to the target actor address space. The *actor* argument specifies the target actor capability.

If *actor* is `K_MYACTOR`, the address space of the current actor is used.

If *actor* is `K_SVACTOR`, the supervisor address space is used.

The *chunk* argument points to the `KnPhMemChunk` structure which specifies the physical and virtual chunk start addresses and the chunk size.

The specified virtual address range must be allocated using the `K_RESERVED` option (see `rgnAllocate(2)`) prior to the invocation of `vmMapToPhys`.

The specified actor can be a supervisor actor as well as a user actor.

The mapping produced by `vmMapToPhys` can only be invalidated by `rgnFree`. *vaddr*, *paddr* and *psize* must be page-aligned.

For PowerPC family processors, the *cntlBits* is defined in the same way as for the `svPhysMap` routine above.

On success `K_OK` is returned, otherwise a negative error code is returned:

<code>K_EFAULT</code>	The <i>actor</i> argument points outside the caller's address space.
<code>K_EINVAL</code>	An inconsistent actor capability was provided.
<code>K_EUNKNOWN</code>	<i>actorcap</i> does not specify a reachable actor.
<code>K_ENOMEM</code>	The system is out of memory.
<code>K_ESIZE</code>	The <i>psize</i> argument is equal to zero.
<code>K_EROUND</code>	<i>vaddr</i> or <i>paddr</i> or <i>psize</i> is not page-aligned.
<code>K_EADDR</code>	Some or all addresses from the target virtual address range are out of a region allocated with the <code>K_RESERVED</code> option.

Note - For performance reasons, the current implementation does not guarantee that any attempt to map a physical address to a virtual address out of a region allocated with `K_RESERVED` option would effectively produce a `K_EADDR` error.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysMap	+	+	-	+
svPhysUnmap	+	+	-	+
vmMapToPhys	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svPhysAlloc(9DKI)` , `svDkiThreadCall(9DKI)` , `rgnAllocate(2K)`

NAME	svPhysMap_usparc, svPhysUnmap_usparc, vmMapToPhys_usparc – UltraSPARC physical to virtual memory mapping
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svPhysMap(KnPhMemChunk * chunk, KnPteAttr attr); void svPhysUnmap(KnPhMemChunk * chunk); KnError vmMapToPhys(KnCap * actor, KnPhMemChunk * chunk, KnPteAttr attr);</pre>
FEATURES	DKI
DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. These services should be used mainly by a host bus driver to map bus I/O space or DMA memory.</p> <p>svPhysMap allocates a range of virtual addresses within the supervisor address space and maps it to a given range of physical addresses.</p> <p>The <i>chunk</i> argument points to the KnPhMemChunk structure shown below:</p> <pre>typedef struct { PhAddr paddr; /* physical start address */ PhSize psize; /* size */ VmAddr vaddr; /* virtual start address */ } KnPhMemChunk;</pre> <p>The <i>paddr</i> and <i>psize</i> fields specify the chunk of physical memory being mapped to. The <i>vaddr</i> field is used by svPhysMap to return the virtual address to which the physical one is mapped.</p> <p>The <i>attr</i> argument specifies the mapping attributes. <i>attr</i> is a bit-mask composed of two independent parts:</p> <ul style="list-style-type: none"> Data access attributes Instruction access attributes <p>Basically, each part of the attributes is a sub-set of bits defined by the translation table entry (TTE) of the UltraSPARC MMU.</p> <p>A combination of the following attributes may be specified for data access:</p> <p>PTE_DATTR_G — global The PTE_DATTR_G bit set allows the mapping to be shared among all (user and supervisor) contexts.</p> <p>PTE_DATTR_W — writable The PTE_DATTR_W bit set grants write permission for the mapping.</p> <p>PTE_DATTR_P — privileged</p>

The `PTE_DATTR_P` bit set restricts access to the mapping for the supervisor only.

`PTE_DATTR_E` — side-effect

The `PTE_DATTR_E` bit set makes noncacheable memory accesses to be strongly ordered against other E-bit accesses, and noncacheable stores are not merged. This bit should be set for I/O devices having side-effects. Note that the E-bit does not force an uncacheable access. It is expected that the `PTE_DATTR_CV` and `PTE_DATTR_CP` bits will be set to zero when the E-bit is set.

`PTE_DATTR_CV` — L1-cacheable

The `PTE_DATTR_CV` bit set allows data to be cached in the (L1) CPU data cache. Note that if the `PTE_DATTR_CV` bit is set, the `PTE_DATTR_CP` bit must also be set.

`PTE_DATTR_CP` — L2-cacheable

The `PTE_DATTR_CP` bit set allows data to be cached in the (L2) external cache.

`PTE_DATTR_IE` — invert endianness

The `PTE_DATTR_IE` bit set causes data accesses to the mapping to be processed with inverse endianness from that specified by the instruction.

`PTE_DATTR_V` — valid

The `PTE_DATTR_V` bit set enables data accesses to the mapping. If this bit is not set, all other bits (described above) are ignored and a data access to the mapping will result in a data access exception.

Combinations of the following attributes may be specified for instruction access:

`PTE_IATTR_G` — global

The `PTE_IATTR_G` bit set allows the mapping to be shared among all (user and supervisor) contexts.

`PTE_IATTR_P` — privileged

The `PTE_IATTR_P` bit set restricts access to the mapping to the supervisor only.

`PTE_IATTR_CV` — L1-cacheable

The `PTE_IATTR_CV` bit set allows instructions to be cached in the (L1) CPU instruction cache. Note that if the `PTE_IATTR_CV` bit is set, the `PTE_IATTR_CP` bit must also be set.

`PTE_IATTR_CP` — L2-cacheable

The `PTE_IATTR_CP` bit set allows instructions to be cached in the (L2) external cache.

PTE_IATTR_V — valid

The PTE_IATTR_V bit set enables instructions to be obtained from the mapping. If this bit is not set, all other bits (described above) are ignored and an instruction to fetch from the mapping will result in an instruction access exception.

On success svPhysMap returns K_OK , otherwise a negative error code is returned as follows:

K_ENOMEM The system is out of memory.

K_ESIZE The *psize* argument is equal to zero.

svPhysUnmap

svPhysUnmap unmaps the physical memory chunk previously mapped by svPhysMap . The *chunk* argument points to the KnPhMemChunk structure which specifies the virtual and physical chunk start addresses and the chunk size. Note that the *paddr* and *psize* fields must have the values previously specified in svPhysMap . The *vaddr* fields must have the value previously returned by svPhysMap . Typically, a driver uses the same KnPhMemChunk structure for the svPhysMap and svPhysUnmap calls and the structure fields are not modified by the driver once svPhysMap has been performed.

vmMapToPhys

vmMapToPhys maps a given physical memory chunk to the target actor address space.

The *actor* argument specifies the target actor capability. If *actor* is K_MYACTOR , the address space of the current actor is used. If *actor* is K_SVACTOR , the supervisor address space is used.

The *chunk* argument points to the KnPhMemChunk structure which specifies the physical and virtual chunk start addresses and the chunk size. The specified virtual address range must be allocated using the K_RESERVED option (see *rgnAllocate(2K)* prior to the invocation of vmMapToPhys . The specified actor can be a supervisor actor as well as an user actor. The mapping produced by vmMapToPhys can only be invalidated by *rgnFree* .

vaddr , *paddr* and *psize* must be page-aligned.

For UltraSPARC family processors, the *attr* argument is defined in the same way as for the svPhysMap routine above.

On success K_OK is returned, otherwise a negative error code is returned as follows:

K_EFAULT The *actor* argument points outside of the caller's address space.

K_EINVAL	An inconsistent actor capability was specified.
K_EUNKNOWN	<i>actorcap</i> does not specify a reachable actor.
K_ENOMEM	The system is out of memory.
K_ESIZE	The <i>psize</i> argument is equal to zero.
K_EROUND	<i>vaddr</i> or <i>paddr</i> or <i>psize</i> is not page-aligned.
K_EADDR	Some or all addresses from the target virtual address range are out of a region allocated with the K_RESERVED option. For performance reasons, the current implementation does not guarantee that any attempt to map a physical address to a virtual address outside a region allocated with K_RESERVED option would effectively produce a K_EADDR error.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysMap	+	+	-	+
svPhysUnmap	+	+	-	+
vmMapToPhys	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svPhysAlloc(9DKI)` , `svDkiThreadCall(9DKI)` , `rgnAllocate(2K)`

NAME svPhysMap_x86, svPhysUnmap_x86, vmMapToPhys_x86 – Intel x86 physical to virtual memory mapping

SYNOPSIS

```
#include <dki/f_dki.h>
KnError svPhysMap(KnPhMemChunk * chunk, PteCntlBits cntlBits);

void svPhysUnmap(KnPhMemChunk * chunk);

KnError vmMapToPhys(KnCap * actor, KnPhMemChunk * chunk, PteCntlBits cntlBits);

typedef struct {
    PhAddr paddr; /* physical start address */
    PhSize psize; /* size */
    VmAddr vaddr; /* virtual start address */
} KnPhMemChunk;
```

FEATURES DKI

EXTENDED DESCRIPTION The microkernel provides services to allow device drivers to map physical space to virtual memory space. These services should be used mainly by primary bus drivers to map bus I/O space, or DMA memory.

svPhysMap

svPhysMap allocates a range of virtual addresses within the supervisor address space and maps it to a given range of physical addresses.

The *chunk* argument points to the KnPhMemChunk structure.

The *paddr* and *psize* fields specify the chunk of physical memory being mapped to.

The *vaddr* field is used by svPhysMap to return the virtual address to which the physical one is mapped.

For Intel ix86 family processor, the cntlBits argument is a direct mapping of the Page Control bits of Page Table Entries, as described in the Intel 386/486 Programming Reference Manual. Therefore, the cntlBits argument should be constructed by “oring” the following values:

- PTE_PRESENT The mapped memory must be present in physical memory.
- PTE_READ_WRITE The mapped memory allows write accesses to be performed.
- PTE_USER_SUPERVISOR The mapped memory can be used in user space.
- PTE_WRITE_TRANSPARENT The memory is mapped with the Write-Through attribute (bit PWT). Thus, a store operation updates the cache if necessary, and in addition the update is written to memory.

PTE_CACHE_DISABLE The memory is mapped with caching disabled (bit PCD). Thus, memory accesses go to main memory, bypassing the caches.

On success K_OK is returned, otherwise a negative error code is returned:

K_ENOMEM The system is out of memory.

K_ESIZE The *psize* argument is equal to zero.

svPhysUnmap

svPhysUnmap unmaps the physical memory chunk previously mapped by svPhysMap . The *chunk* argument points to the KnPhMemChunk structure which specifies the virtual and physical chunk start addresses and the chunk size. Note that the *paddr* and *psize* fields must have the values previously specified in svPhysMap . The *vaddr* field must have the value previously returned by svPhysMap (typically, a driver uses the same KnPhMemChunk structure for the svPhysMap and svPhysUnmap calls and the structure fields are not modified by the driver once svPhysMap is done).

vmMapToPhys

vmMapToPhys maps a given physical memory chunk to the target actor address space.

The *actor* argument specifies the target actor capability. If *actor* is K_MYACTOR , the address space of the current actor is used. If *actor* is K_SVACTOR , the supervisor address space is used.

The *chunk* argument points to the KnPhMemChunk structure which specifies the physical and virtual chunk start addresses and the chunk size. The specified virtual address range must be allocated using the K_RESERVED option (see *rgnAllocate* (2K)) prior to the invocation of vmMapToPhys . The specified actor can be a supervisor actor as well as a user actor. The mapping produced by vmMapToPhys can only be invalidated by rgnFree .

vaddr , *paddr* and *psize* must be page-aligned.

For Intel x86 family processors, the *cntlBits* is defined in the same way as for the svPhysMap routine above.

On success K_OK is returned, otherwise a negative error code is returned:

K_EFAULT The *actor* argument points to the outside of the caller's address space.

K_EINVAL An inconsistent actor capability was provided.

K_EUNKNOWN	<i>actorcap</i> does not specify a reachable actor.
K_ENOMEM	The system is out of memory.
K_ESIZE	The <i>psize</i> argument is equal to zero.
K_EROUND	<i>vaddr</i> or <i>paddr</i> or <i>psize</i> is not page-aligned.
K_EADDR	Some or all addresses from the target virtual address range are out of a region allocated with the K_RESERVED option (for performance reasons the current implementation does not guarantee that any attempt to map a physical address to a virtual address out of a region allocated with K_RESERVED option would effectively produce a K_EADDR error).

Intel x86 Memory Mapping Allowed Calling Contexts

The table below specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysMap	+	+	-	+
svPhysUnmap	+	+	-	+
vmMapToPhys	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svPhysAlloc(9DKI) , *svDkiThreadCall(9DKI)* , *rgnAllocate(9DKI)*

NAME	svIntrAttach_usparc, svIntrDetach_usparc, svSoftIntrAttach_usparc, svSoftIntrDetach_usparc, svTimerIntrAttach_uparc, svTimerIntrDetach_uparc, svIntrCtxGet_usparc – UltraSPARC interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(unsigned int intrNumb, unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrIdp); void svIntrDetach(CpuIntrId intrId); KnError svSoftIntrAttach(unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svSoftIntrDetach(CpuIntrId intrId); KnError svTimerIntrAttach(CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svTimerIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx);</pre>
FEATURES	DKI
DESCRIPTION	<p>Provides UltraSPARC interrupts management services.</p> <p>An UltraSPARC processor is an implementation of the SPARC-V9 CPU architecture. As specified by the SPARC-V9 architecture, there are fifteen CPU interrupt sources assigned to the fifteen CPU interrupt levels from 1 up to 15. However, the UltraSPARC processor uses these fifteen interrupt levels only for software generated interrupts. The hardware interrupts are delivered to the processor using the "Mondo" interrupt transfer mechanism. The hardware interrupt source is designated by the interrupt number. Typically, the interrupt number is 11-bit width and composed of the interrupt group number (5 MSB) and the interrupt offset number (6 LSB).</p> <p>The "Mondo" interrupt dispatch handler is built into the microkernel. The microkernel handles a mapping between the hardware interrupt numbers and software interrupt levels. When a mondo interrupt packet is received by the microkernel, the interrupt request descriptor is queued and an associated software interrupt is triggered. The microkernel software interrupt handler then dequeues the interrupt request descriptor and invokes a handler associated to the given interrupt number. In this way, a driver interrupt handler is always invoked in the UltraSPARC software interrupt context. This provides the interrupt handler with an environment analogous to the SPARC-V8 one (interrupt levels).</p> <p>The microkernel provides services which allow device drivers to manage UltraSPARC interrupts, mainly to attach/detach handlers to the CPU interrupts.</p> <pre>typedef void (*CpuIntrHandler)(void*);</pre>

svIntrAttach

svIntrAttach attaches a given handler to a given interrupt number at a given processor interrupt level.

The *intrNumb* argument specifies the interrupt number to which to attach.

The *intrLevel* argument specifies the interrupt level to which to attach.

The *intrHandler* argument specifies the handler to call back when the given interrupt occurs.

The *intrCookie* argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.

On success, K_OK is returned and the services defined on the attached interrupt are returned in the *intrOps* parameter. An identifier for the attached interrupt is also returned in *intrId*. This identifier must be used as first argument to subsequent calls to *intrOps* services.

On failure, an error code is returned as follows:

K_EINVAL	The specified interrupt number or interrupt level are invalid.
K_BUSY	Another handler is already attached to the given <i>intrNumb</i> .
K_ENOMEM	The system is out of memory.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure as follows:

```
typedef struct CpuIntrOps {
    void
    (*mask) (CpuIntrId intrId);

    void
    (*unmask) (CpuIntrId intrId);

    void
    (*enable) (CpuIntrId intrId);

    void
    (*disable) (CpuIntrId intrId);

    void
    (*trigger) (CpuIntrId intrId);
} CpuIntrOps;
```

The `CpuIntrOps.mask` routine disables the interrupt at CPU level identified by `intrId`. In other words, the PIL register is set to the level corresponding to `intrId`. Note that the original value of the PIL register is saved by DKI in order to be restored later by the `CpuIntrOps.unmask` routine.

The `CpuIntrOps.unmask` routine enables the interrupt at CPU level identified by `intrId`.

In other words, the PIL register is restored to the original value saved by the previously called `CpuIntrOps.mask` routine.

The `mask / unmask` pair may be called from base level only and must not be nested. The `mask/unmask` pair is typically used to implement a critical section of code which needs to be protected against the interrupt. Note that with respect to the SPARC-V9 architecture, when an interrupt level `N` is masked, all interrupts with a level less than `N` are also masked. Thus, there is no way to `mask` only one CPU interrupt level except the lowest one.

The `enable / disable` pair may only be called from the attached interrupt handler. When an interrupt occurs, the attached `CpuIntrHandler` is invoked with the interrupt masked at processor level. This behaves in exactly the same way as if `CpuIntrOps.disable` was called just prior to the handler invocation. In other words, the PIL register is set to the interrupt level and the original interrupt processor level (which was when the interrupt occurred) is saved by DKI. Note that the interrupt handler must return to DKI in the same context as it was called, that is with the interrupt disabled at processor level.

On the other hand, the called interrupt handler may use the `enable/disable` pair to allow the interrupt to be nested. This feature is typically used by a host bus driver when the bus interrupts are multiplexed, that is, multiple bus interrupts are reported at the same CPU interrupt level. Typically, an interrupt handler of this type of host bus driver would take the following actions:

- Identify the bus interrupt source (through a PIC or special cycle).
- Disable the bus interrupt source at bus level (through PIC).
- Enable interrupt at processor level (`enable`).
- Call handlers attached to the identified bus interrupt source.
- Disable interrupt at processor level (`disable`).
- Acknowledge (if needed) and enable the bus interrupt source at bus level (through PIC).
- Return to the DKI.

The `CpuIntrOps.trigger` routine allows the interrupt to be triggered by software. Basically, this routine acts like the mondo interrupt dispatcher except the interrupt number is obtained from `intrId` rather than from the mondo interrupt packet. The `CpuIntrOps.trigger` routine is mainly dedicated to the

	software interrupts attached by <code>svSoftIntrAttach</code> . However, it may be also used for hardware interrupts, for instance, for debugging or diagnostic purposes.				
svIntrDetach	<p><code>svIntrDetach</code> detaches an interrupt handler previously connected by <code>svIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svIntrAttach</code> .</p>				
svSoftIntrDetach	<p><code>svSoftIntrAttach</code> attaches a given software interrupt handler to a given processor interrupt level.</p> <p>The <i>intrLevel</i> argument specifies the interrupt level to which to attach.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On success, <code>K_OK</code> is returned and the services defined on the attached interrupt are returned in the <i>intrOps</i> parameter.</p> <p>An identifier for the attached interrupt is also returned in <i>intrId</i> . This identifier must be used as the first argument to subsequent calls to <i>intrOps</i> services.</p> <p>On failure, an error code is returned as follows:</p> <table border="0"> <tr> <td><code>K_EINVAL</code></td> <td>The specified interrupt level is invalid.</td> </tr> <tr> <td><code>K_ENOMEM</code></td> <td>The system is out of memory.</td> </tr> </table> <p>Services available on an attached interrupt are defined by the <code>CpuIntrOps</code> structure described above.</p>	<code>K_EINVAL</code>	The specified interrupt level is invalid.	<code>K_ENOMEM</code>	The system is out of memory.
<code>K_EINVAL</code>	The specified interrupt level is invalid.				
<code>K_ENOMEM</code>	The system is out of memory.				
svSoftIntrDetach	<p><code>svSoftIntrDetach</code> detaches a software interrupt handler previously connected by <code>svSoftIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svSoftIntrAttach</code> .</p>				
svTimerIntrAttach	<p><code>svTimerIntrAttach</code> attaches a given interrupt handler to the UltraSPARC tick-counter interrupt.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque the microkernel.</p>				

On success, `K_OK` is returned and the services defined on the attached interrupt are returned in the `intrOps` parameter. An identifier for the attached interrupt is also returned in `intrId`. This identifier must be used as the first argument to subsequent calls to `intrOps` services.

On failure, an error code is returned as follows:

- `K_BUSY` Another handler is already attached to the tick-counter interrupt.
- `K_ENOMEM` The system is out of memory.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure described above.

svTimerIntrDetach

`svTimerIntrDetach` detaches the interrupt handler previously connected by `svTimerIntrAttach`.

The `intrId` argument identifies the attached interrupt handler, previously returned by `svTimerIntrAttach`.

svIntrCtxGet

`svIntrCtxGet` retrieves the current level interrupt context. It is typically used for profiling purposes.

On success, `K_OK` is returned and a pointer to the recently saved interrupt context is returned in the `intrCtx` argument. The CPU context saved on interrupt has the same structure as the thread context saved on exception or trap. It contains the global registers `%g1-%g7`, the output registers of the interrupted window `%o0-%o7` and the following processor registers: `%tstate`, `%pc`, `%npc`, `%tt`, `%y`. In addition, the thread context contains the number of outstanding windows and the pointer to the outstanding windows buffer if the number of windows is greater than zero.

On failure, an error code is returned as follows:

- `K_EINVAL` Interrupt level is zero (not called from an interrupt handler).
- `K_ENOTAVAILABLE` There is no context available for currently handled interrupt.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each interrupt management service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svIntrAttach</code>	+	+	-	+

svIntrDetach	+	+	-	+
svIntrSoftAttach	+	+	-	+
svIntrSoftDetach	+	+	-	+
svTimerIntrAttach	+	+	-	+
svTimerIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-
CpuIntrOps.mask	+	+	-	-
CpuIntrOps.unmask	+	+	-	-
CpuIntrOps.enable	-	-	+	-
CpuIntrOps.disable	-	-	+	-
CpuIntrOps.trigger	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME	svIntrAttach_usparc, svIntrDetach_usparc, svSoftIntrAttach_usparc, svSoftIntrDetach_usparc, svTimerIntrAttach_uparc, svTimerIntrDetach_uparc, svIntrCtxGet_usparc – UltraSPARC interrupts management
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svIntrAttach(unsigned int intrNumb, unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrIdp); void svIntrDetach(CpuIntrId intrId); KnError svSoftIntrAttach(unsigned int intrLevel, CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svSoftIntrDetach(CpuIntrId intrId); KnError svTimerIntrAttach(CpuIntrHandler intrHandler, void * intrCookie, CpuIntrOps ** intrOps, CpuIntrId * intrId); void svTimerIntrDetach(CpuIntrId intrId); KnError svIntrCtxGet(KnIntrCtx ** intrCtx);</pre>
FEATURES	DKI
DESCRIPTION	<p>Provides UltraSPARC interrupts management services.</p> <p>An UltraSPARC processor is an implementation of the SPARC-V9 CPU architecture. As specified by the SPARC-V9 architecture, there are fifteen CPU interrupt sources assigned to the fifteen CPU interrupt levels from 1 up to 15. However, the UltraSPARC processor uses these fifteen interrupt levels only for software generated interrupts. The hardware interrupts are delivered to the processor using the "Mondo" interrupt transfer mechanism. The hardware interrupt source is designated by the interrupt number. Typically, the interrupt number is 11-bit width and composed of the interrupt group number (5 MSB) and the interrupt offset number (6 LSB).</p> <p>The "Mondo" interrupt dispatch handler is built into the microkernel. The microkernel handles a mapping between the hardware interrupt numbers and software interrupt levels. When a mondo interrupt packet is received by the microkernel, the interrupt request descriptor is queued and an associated software interrupt is triggered. The microkernel software interrupt handler then dequeues the interrupt request descriptor and invokes a handler associated to the given interrupt number. In this way, a driver interrupt handler is always invoked in the UltraSPARC software interrupt context. This provides the interrupt handler with an environment analogous to the SPARC-V8 one (interrupt levels).</p> <p>The microkernel provides services which allow device drivers to manage UltraSPARC interrupts, mainly to attach/detach handlers to the CPU interrupts.</p> <pre>typedef void (*CpuIntrHandler)(void*);</pre>

svIntrAttach

`svIntrAttach` attaches a given handler to a given interrupt number at a given processor interrupt level.

The *intrNumb* argument specifies the interrupt number to which to attach.

The *intrLevel* argument specifies the interrupt level to which to attach.

The *intrHandler* argument specifies the handler to call back when the given interrupt occurs.

The *intrCookie* argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.

On success, `K_OK` is returned and the services defined on the attached interrupt are returned in the *intrOps* parameter. An identifier for the attached interrupt is also returned in *intrId*. This identifier must be used as first argument to subsequent calls to *intrOps* services.

On failure, an error code is returned as follows:

<code>K_EINVAL</code>	The specified interrupt number or interrupt level are invalid.
<code>K_BUSY</code>	Another handler is already attached to the given <i>intrNumb</i> .
<code>K_ENOMEM</code>	The system is out of memory.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure as follows:

```
typedef struct CpuIntrOps {
    void
    (*mask) (CpuIntrId intrId);

    void
    (*unmask) (CpuIntrId intrId);

    void
    (*enable) (CpuIntrId intrId);

    void
    (*disable) (CpuIntrId intrId);

    void
    (*trigger) (CpuIntrId intrId);
} CpuIntrOps;
```

The `CpuIntrOps.mask` routine disables the interrupt at CPU level identified by `intrId`. In other words, the PIL register is set to the level corresponding to `intrId`. Note that the original value of the PIL register is saved by DKI in order to be restored later by the `CpuIntrOps.unmask` routine.

The `CpuIntrOps.unmask` routine enables the interrupt at CPU level identified by `intrId`.

In other words, the PIL register is restored to the original value saved by the previously called `CpuIntrOps.mask` routine.

The `mask / unmask` pair may be called from base level only and must not be nested. The `mask/unmask` pair is typically used to implement a critical section of code which needs to be protected against the interrupt. Note that with respect to the SPARC-V9 architecture, when an interrupt level `N` is masked, all interrupts with a level less than `N` are also masked. Thus, there is no way to `mask` only one CPU interrupt level except the lowest one.

The `enable / disable` pair may only be called from the attached interrupt handler. When an interrupt occurs, the attached `CpuIntrHandler` is invoked with the interrupt masked at processor level. This behaves in exactly the same way as if `CpuIntrOps.disable` was called just prior to the handler invocation. In other words, the PIL register is set to the interrupt level and the original interrupt processor level (which was when the interrupt occurred) is saved by DKI. Note that the interrupt handler must return to DKI in the same context as it was called, that is with the interrupt disabled at processor level.

On the other hand, the called interrupt handler may use the `enable/disable` pair to allow the interrupt to be nested. This feature is typically used by a host bus driver when the bus interrupts are multiplexed, that is, multiple bus interrupts are reported at the same CPU interrupt level. Typically, an interrupt handler of this type of host bus driver would take the following actions:

- Identify the bus interrupt source (through a PIC or special cycle).
- Disable the bus interrupt source at bus level (through PIC).
- Enable interrupt at processor level (`enable`).
- Call handlers attached to the identified bus interrupt source.
- Disable interrupt at processor level (`disable`).
- Acknowledge (if needed) and enable the bus interrupt source at bus level (through PIC).
- Return to the DKI.

The `CpuIntrOps.trigger` routine allows the interrupt to be triggered by software. Basically, this routine acts like the mondo interrupt dispatcher except the interrupt number is obtained from `intrId` rather than from the mondo interrupt packet. The `CpuIntrOps.trigger` routine is mainly dedicated to the

	software interrupts attached by <code>svSoftIntrAttach</code> . However, it may be also used for hardware interrupts, for instance, for debugging or diagnostic purposes.				
svIntrDetach	<p><code>svIntrDetach</code> detaches an interrupt handler previously connected by <code>svIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svIntrAttach</code> .</p>				
svSoftIntrDetach	<p><code>svSoftIntrAttach</code> attaches a given software interrupt handler to a given processor interrupt level.</p> <p>The <i>intrLevel</i> argument specifies the interrupt level to which to attach.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque to the microkernel.</p> <p>On success, <code>K_OK</code> is returned and the services defined on the attached interrupt are returned in the <i>intrOps</i> parameter.</p> <p>An identifier for the attached interrupt is also returned in <i>intrId</i> . This identifier must be used as the first argument to subsequent calls to <i>intrOps</i> services.</p> <p>On failure, an error code is returned as follows:</p> <table border="0"> <tr> <td><code>K_EINVAL</code></td> <td>The specified interrupt level is invalid.</td> </tr> <tr> <td><code>K_ENOMEM</code></td> <td>The system is out of memory.</td> </tr> </table> <p>Services available on an attached interrupt are defined by the <code>CpuIntrOps</code> structure described above.</p>	<code>K_EINVAL</code>	The specified interrupt level is invalid.	<code>K_ENOMEM</code>	The system is out of memory.
<code>K_EINVAL</code>	The specified interrupt level is invalid.				
<code>K_ENOMEM</code>	The system is out of memory.				
svSoftIntrDetach	<p><code>svSoftIntrDetach</code> detaches a software interrupt handler previously connected by <code>svSoftIntrAttach</code> .</p> <p>The <i>intrId</i> argument identifies the attached interrupt handler, previously returned by <code>svSoftIntrAttach</code> .</p>				
svTimerIntrAttach	<p><code>svTimerIntrAttach</code> attaches a given interrupt handler to the UltraSPARC tick-counter interrupt.</p> <p>The <i>intrHandler</i> argument specifies the handler to call back when the given interrupt is triggered.</p> <p>The <i>intrCookie</i> argument specifies a parameter to pass back to the handler when called. It is opaque the microkernel.</p>				

On success, `K_OK` is returned and the services defined on the attached interrupt are returned in the `intrOps` parameter. An identifier for the attached interrupt is also returned in `intrId`. This identifier must be used as the first argument to subsequent calls to `intrOps` services.

On failure, an error code is returned as follows:

- `K_BUSY` Another handler is already attached to the tick-counter interrupt.
- `K_ENOMEM` The system is out of memory.

Services available on an attached interrupt are defined by the `CpuIntrOps` structure described above.

svTimerIntrDetach

`svTimerIntrDetach` detaches the interrupt handler previously connected by `svTimerIntrAttach`.

The `intrId` argument identifies the attached interrupt handler, previously returned by `svTimerIntrAttach`.

svIntrCtxGet

`svIntrCtxGet` retrieves the current level interrupt context. It is typically used for profiling purposes.

On success, `K_OK` is returned and a pointer to the recently saved interrupt context is returned in the `intrCtx` argument. The CPU context saved on interrupt has the same structure as the thread context saved on exception or trap. It contains the global registers `%g1-%g7`, the output registers of the interrupted window `%o0-%o7` and the following processor registers: `%tstate`, `%pc`, `%npc`, `%tt`, `%y`. In addition, the thread context contains the number of outstanding windows and the pointer to the outstanding windows buffer if the number of windows is greater than zero.

On failure, an error code is returned as follows:

- `K_EINVAL` Interrupt level is zero (not called from an interrupt handler).
- `K_ENOTAVAILABLE` There is no context available for currently handled interrupt.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each interrupt management service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svIntrAttach</code>	+	+	-	+

svIntrDetach	+	+	-	+
svIntrSoftAttach	+	+	-	+
svIntrSoftDetach	+	+	-	+
svTimerIntrAttach	+	+	-	+
svTimerIntrDetach	+	+	-	+
svIntrCtxGet	-	-	+	-
CpuIntrOps.mask	+	+	-	-
CpuIntrOps.unmask	+	+	-	-
CpuIntrOps.enable	-	-	+	-
CpuIntrOps.disable	-	-	+	-
CpuIntrOps.trigger	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME	svTimeoutSet, svTimeoutCancel, svTimeoutGetRes – timeout operations
SYNOPSIS	<pre>#include <dki/dki.h> KnError svTimeoutSet(KnTimeout * timeout, KnToHdl handler, KnTimeVal * waitLimit, int flag); Bool svTimeoutCancel(KnTimeout * timeout); int svTimeoutGetRes(KnTimeVal * resolution);</pre>
FEATURES	DKI
DESCRIPTION	Provides timeout operations services.
EXTENDED DESCRIPTION	<p>Device drivers may need Timeout services to check whether there is any activity on a device, or to verify that a started action will terminate before a given time limit is reached. The svTimeout API is used for DKI purposes. Note that as these services should be implemented using drivers, they are not available and must not be used by drivers at initialization time.</p> <p>svTimeoutSet sets a timeout request. When the interval of time specified by waitLimit has expired, the handler is invoked with timeout passed as its only argument. Fields within the KnTimeout structure are initialized and modified solely within the nucleus and are inaccessible to an application.</p> <p>The handler argument points to a timeout handler routine. A timeout handler 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 the supervisor space. The set of kernel calls that can be used in a timeout handler are limited, as for any interrupt handler.</p> <p>The waitLimit pointer argument refers to a KnTimeVal structure whose members are defined in sysTime(2K).</p> <p>svTimeoutSet expects the KnTimeVal structure parameter to be set to a relative interval of time after which the handler is invoked. The flag argument is unused and should therefore be set to 0 .</p> <p>On success, svTimeoutSet returns K_OK .</p> <p>Otherwise, K_EINVAL is returned specifying that the waitLimit value is invalid or that no handler was specified.</p> <p>svTimeoutSet must not be called using a KnTimeout which has an existing timeout pending (use svTimeoutCancel to cancel the request first).</p>

svTimeoutCancel attempts to cancel a timeout request. It takes as an argument the address of the KnTimeout object used in the call to svTimeoutSet . If the timeout request is still pending, it is immediately cancelled and TRUE is returned signifying that the timer had not yet expired. Otherwise, the timeout interval had already passed and the timeout handler was called, in which case nothing occurs and FALSE is returned.

svTimeoutGetRes returns in resolution the smallest possible difference between two distinct waitLimit values. Note that Timeout resolution usually corresponds to the microkernel tick period (usually 10 milli-seconds).

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svTimeoutSet	+	+	+	-
svTimeoutCancel	+	+	+	-
svTimeoutGetRes	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

sysTime(2K) , svDkiThreadCall(9DKI)

NAME	svTimeoutSet, svTimeoutCancel, svTimeoutGetRes – timeout operations
SYNOPSIS	<pre>#include <dki/dki.h> KnError svTimeoutSet(KnTimeout * timeout, KnToHdl handler, KnTimeVal * waitLimit, int flag); Bool svTimeoutCancel(KnTimeout * timeout); int svTimeoutGetRes(KnTimeVal * resolution);</pre>
FEATURES	DKI
DESCRIPTION	Provides timeout operations services.
EXTENDED DESCRIPTION	<p>Device drivers may need Timeout services to check whether there is any activity on a device, or to verify that a started action will terminate before a given time limit is reached. The svTimeout API is used for DKI purposes. Note that as these services should be implemented using drivers, they are not available and must not be used by drivers at initialization time.</p> <p>svTimeoutSet sets a timeout request. When the interval of time specified by waitLimit has expired, the handler is invoked with timeout passed as its only argument. Fields within the KnTimeout structure are initialized and modified solely within the nucleus and are inaccessible to an application.</p> <p>The handler argument points to a timeout handler routine. A timeout handler 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 the supervisor space. The set of kernel calls that can be used in a timeout handler are limited, as for any interrupt handler.</p> <p>The waitLimit pointer argument refers to a KnTimeVal structure whose members are defined in sysTime(2K).</p> <p>svTimeoutSet expects the KnTimeVal structure parameter to be set to a relative interval of time after which the handler is invoked. The flag argument is unused and should therefore be set to 0 .</p> <p>On success, svTimeoutSet returns K_OK .</p> <p>Otherwise, K_EINVAL is returned specifying that the waitLimit value is invalid or that no handler was specified.</p> <p>svTimeoutSet must not be called using a KnTimeout which has an existing timeout pending (use svTimeoutCancel to cancel the request first).</p>

`svTimeoutCancel` attempts to cancel a timeout request. It takes as an argument the address of the `KnTimeout` object used in the call to `svTimeoutSet`. If the timeout request is still pending, it is immediately cancelled and `TRUE` is returned signifying that the timer had not yet expired. Otherwise, the timeout interval had already passed and the timeout handler was called, in which case nothing occurs and `FALSE` is returned.

`svTimeoutGetRes` returns in *resolution* the smallest possible difference between two distinct *waitLimit* values. Note that `Timeout` resolution usually corresponds to the microkernel tick period (usually 10 milli-seconds).

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
<code>svTimeoutSet</code>	+	+	+	-
<code>svTimeoutCancel</code>	+	+	+	-
<code>svTimeoutGetRes</code>	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`sysTime(2K)`, `svDkiThreadCall(9DKI)`

NAME	svTimeoutSet, svTimeoutCancel, svTimeoutGetRes – timeout operations
SYNOPSIS	<pre>#include <dki/dki.h> KnError svTimeoutSet(KnTimeout * timeout, KnToHdl handler, KnTimeVal * waitLimit, int flag); Bool svTimeoutCancel(KnTimeout * timeout); int svTimeoutGetRes(KnTimeVal * resolution);</pre>
FEATURES	DKI
DESCRIPTION	Provides timeout operations services.
EXTENDED DESCRIPTION	<p>Device drivers may need Timeout services to check whether there is any activity on a device, or to verify that a started action will terminate before a given time limit is reached. The svTimeout API is used for DKI purposes. Note that as these services should be implemented using drivers, they are not available and must not be used by drivers at initialization time.</p> <p>svTimeoutSet sets a timeout request. When the interval of time specified by waitLimit has expired, the handler is invoked with timeout passed as its only argument. Fields within the KnTimeout structure are initialized and modified solely within the nucleus and are inaccessible to an application.</p> <p>The handler argument points to a timeout handler routine. A timeout handler 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 the supervisor space. The set of kernel calls that can be used in a timeout handler are limited, as for any interrupt handler.</p> <p>The waitLimit pointer argument refers to a KnTimeVal structure whose members are defined in sysTime(2K).</p> <p>svTimeoutSet expects the KnTimeVal structure parameter to be set to a relative interval of time after which the handler is invoked. The flag argument is unused and should therefore be set to 0 .</p> <p>On success, svTimeoutSet returns K_OK .</p> <p>Otherwise, K_EINVAL is returned specifying that the waitLimit value is invalid or that no handler was specified.</p> <p>svTimeoutSet must not be called using a KnTimeout which has an existing timeout pending (use svTimeoutCancel to cancel the request first).</p>

svTimeoutCancel attempts to cancel a timeout request. It takes as an argument the address of the KnTimeout object used in the call to svTimeoutSet . If the timeout request is still pending, it is immediately cancelled and TRUE is returned signifying that the timer had not yet expired. Otherwise, the timeout interval had already passed and the timeout handler was called, in which case nothing occurs and FALSE is returned.

svTimeoutGetRes returns in resolution the smallest possible difference between two distinct waitLimit values. Note that Timeout resolution usually corresponds to the microkernel tick period (usually 10 milli-seconds).

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svTimeoutSet	+	+	+	-
svTimeoutCancel	+	+	+	-
svTimeoutGetRes	+	+	-	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

sysTime(2K) , svDkiThreadCall(9DKI)

NAME loadSwap_16, storeSwap_16, swap_16, loadSwap_32, storeSwap_32, swap_32, loadSwap_64, storeSwap_64, swap_64 – specific i/o services

SYNOPSIS

```
#include <dki/dki.h>
uint16_f loadSwap_16(uint16_f * addr);

void storeSwap_16(uint16_f * addr, uint16_f value);

void swap_16(uint16_f * addr);

uint32_f loadSwap_32(uint32_f * addr);

void storeSwap_32(uint32_f * addr, uint32_f value);

void swap_32(uint32_f * addr);

uint64_f loadSwap_64(uint64_f * addr);

void storeSwap_64(uint64_f * addr, uint64_f value);

void swap_64(uint64_f * addr);
```

FEATURES DKI

DESCRIPTION This function provides DKI routines to handle different byte ordering between the processor bus and the host bus.

EXTENDED DESCRIPTION Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus. Specific I/O services are defined below as sets of routines where the *_xx* suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- _16* for 16-bit data
- _32* for 32-bit data
- _64* for 64-bit data

loadSwap_xx loads data from a given address and returns the corresponding byte swapped value. The *addr* argument specifies the address to read from.

storeSwap_xx stores into a given address the *value* byte swapped. The *addr* argument specifies the address to write to.

swap_xx swap in place the bytes of the data stored at a given address. The *addr* argument specifies the address of the data to byte-swap.

Allowed Calling Contexts The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
loadSwap_xx	+	+	+	-

storeSwap_xx	+	+	+	-
swap_xx	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME loadSwap_16, storeSwap_16, swap_16, loadSwap_32, storeSwap_32, swap_32, loadSwap_64, storeSwap_64, swap_64 – specific i/o services

SYNOPSIS

```
#include <dki/dki.h>
uint16_f loadSwap_16(uint16_f * addr);

void storeSwap_16(uint16_f * addr, uint16_f value);

void swap_16(uint16_f * addr);

uint32_f loadSwap_32(uint32_f * addr);

void storeSwap_32(uint32_f * addr, uint32_f value);

void swap_32(uint32_f * addr);

uint64_f loadSwap_64(uint64_f * addr);

void storeSwap_64(uint64_f * addr, uint64_f value);

void swap_64(uint64_f * addr);
```

FEATURES DKI

DESCRIPTION This function provides DKI routines to handle different byte ordering between the processor bus and the host bus.

EXTENDED DESCRIPTION Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus. Specific I/O services are defined below as sets of routines where the *_xx* suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- _16* for 16-bit data
- _32* for 32-bit data
- _64* for 64-bit data

loadSwap_xx loads data from a given address and returns the corresponding byte swapped value. The *addr* argument specifies the address to read from.

storeSwap_xx stores into a given address the *value* byte swapped. The *addr* argument specifies the address to write to.

swap_xx swap in place the bytes of the data stored at a given address. The *addr* argument specifies the address of the data to byte-swap.

Allowed Calling Contexts The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
loadSwap_xx	+	+	+	-

storeSwap_xx	+	+	+	-
swap_xx	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svDkiThreadCall(9DKI)`

NAME loadSwap_16, storeSwap_16, swap_16, loadSwap_32, storeSwap_32, swap_32, loadSwap_64, storeSwap_64, swap_64 – specific i/o services

SYNOPSIS

```
#include <dki/dki.h>
uint16_f loadSwap_16(uint16_f * addr);

void storeSwap_16(uint16_f * addr, uint16_f value);

void swap_16(uint16_f * addr);

uint32_f loadSwap_32(uint32_f * addr);

void storeSwap_32(uint32_f * addr, uint32_f value);

void swap_32(uint32_f * addr);

uint64_f loadSwap_64(uint64_f * addr);

void storeSwap_64(uint64_f * addr, uint64_f value);

void swap_64(uint64_f * addr);
```

FEATURES DKI

DESCRIPTION This function provides DKI routines to handle different byte ordering between the processor bus and the host bus.

EXTENDED DESCRIPTION Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the processor bus and the host bus. Specific I/O services are defined below as sets of routines where the *_xx* suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:

- _16* for 16-bit data
- _32* for 32-bit data
- _64* for 64-bit data

loadSwap_xx loads data from a given address and returns the corresponding byte swapped value. The *addr* argument specifies the address to read from.

storeSwap_xx stores into a given address the *value* byte swapped. The *addr* argument specifies the address to write to.

swap_xx swap in place the bytes of the data stored at a given address. The *addr* argument specifies the address of the data to byte-swap.

Allowed Calling Contexts The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
loadSwap_xx	+	+	+	-

storeSwap_xx	+	+	+	-
swap_xx	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svDkiThreadCall(9DKI)

NAME | loadSwapEieio_16, storeSwapEieio_16, swapEieio_16, loadSwapEieio_32, storeSwapEieio_32, swapEieio_32, eieio - i/o services

FEATURES | DKI

DESCRIPTION | Provides access to DKI i/o routines, optimized to facilitate byte swapping and synchronization of input/output.
 See the architecture specific man pages:

- loadSwapEieio_16_powerpc(9DKI)
- storeSwapEieio_16_powerpc(9DKI)
- swapEieio_16_powerpc(9DKI)
- loadSwapEieio_32_powerpc(9DKI)
- storeSwapEieio_32_powerpc(9DKI)
- swapEieio_32_powerpc(9DKI)
- eieio(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	loadSwapEieio_16_powerpc, loadSwapSync_16_powerpc, storeSwapEieio_16_powerpc, storeSwapSync_16_powerpc, swapEieio_16_powerpc, loadSwapEieio_32_powerpc, loadSwapSync_32_powerpc, storeSwapEieio_32_powerpc, storeSwapSync_32_powerpc, swapEieio_32_powerpc, eieio_powerpc, ioSync_powerpc – PowerPC specific i/o services				
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwapEieio_16(uint16_f * addr); unit16_f loadSwapSync_16(uint16_f * addr); void storeSwapEieio_16(uint16_f * addr, uint16_f value); void storeSwapSync_16(uint16_f * addr, uint16_f value); void swapEieio_16(uint16_f * addr); uint32_f loadSwapEieio_32(uint32_f * addr); unit32_f loadSwapSync_32(uint32_f * addr); void storeSwapEieio_32(uint32_f * addr, uint32_f value); void storeSwapSync_32(uint32_f * addr, uint32_f value); void swapEieio_32 (uint32_f * addr); void ioSync(void); void eieio(void);</pre>				
FEATURES	DKI				
DESCRIPTION	Provides access to PowerPC DKI I/O routines, optimized to facilitate byte swapping and synchronization of input/output.				
EXTENDED DESCRIPTION	<p>The PowerPC DKI provides specific I/O routines to handle byte swapping and synchronization of I/O. These services are based on PowerPC specific instructions to store/load with byte swapping, and to enforce in-order execution of I/O.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the PowerPC bus and the host bus, or to enforce sequential execution of I/O. PowerPC specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <table border="0" style="margin-left: 20px;"> <tr> <td><code>_16</code></td> <td>for 16-bit data</td> </tr> <tr> <td><code>_32</code></td> <td>for 32-bit data</td> </tr> </table>	<code>_16</code>	for 16-bit data	<code>_32</code>	for 32-bit data
<code>_16</code>	for 16-bit data				
<code>_32</code>	for 32-bit data				

`loadSwapEieio_xx` loads data from a given address and returns the corresponding byte swapped value. A PowerPC `eieio` instruction is issued after the data is read, to enforce execution of subsequent I/O sequentially with respect to the current load. The *addr* argument specifies the address to read from.

`loadSwapSync_xx` behaves like `loadSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`storeSwapEieio_xx` stores, at a given address, the *value* byte swapped. A PowerPC `eieio` instruction is issued after the data is written, to enforce execution of subsequent I/O sequentially with respect to the current store. The *addr* argument specifies the address to write to.

`storeSwapSync_xx` behaves like `storeSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`swapEieio_xx` swaps in place the bytes of the data stored at a given address. A PowerPC `eieio` instruction is issued between reading the data and re-writing the byte-swapped data, to enforce execution in order. The *addr* argument specifies the address of the data to byte-swap.

`ioSync` enforces context synchronization, completion of all instructions, and completion of all exceptions before returning.

`eieio` enforces execution in order of I/O and is mapped to the corresponding PowerPC assembler instruction. This routine is intended for use in managing shared data structures, in doing memory-mapped I/O, and in preventing load/store combining operations in main memory. Refer to PowerPC programming manuals for more details.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwapEieio_xx</code>	+	+	+	-
<code>loadSwapSync_xx</code>	+	+	+	-
<code>storeSwapEieio_xx</code>	+	+	+	-
<code>storeSwapSync_xx</code>	+	+	+	-
<code>swapEieio_xx</code>	+	+	+	-
<code>ioSync</code>	+	+	+	-
<code>eieio</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap_16(9DKI)` , `svDkiThreadCall(9DKI)`

NAME loadSwapEieio_16, storeSwapEieio_16, swapEieio_16, loadSwapEieio_32, storeSwapEieio_32, swapEieio_32, eieio – i/o services

FEATURES DKI

DESCRIPTION Provides access to DKI i/o routines, optimized to facilitate byte swapping and synchronization of input/output.

See the architecture specific man pages:

- loadSwapEieio_16_powerpc(9DKI)
- storeSwapEieio_16_powerpc(9DKI)
- swapEieio_16_powerpc(9DKI)
- loadSwapEieio_32_powerpc(9DKI)
- storeSwapEieio_32_powerpc(9DKI)
- swapEieio_32_powerpc(9DKI)
- eieio(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	loadSwapEieio_16_powerpc, loadSwapSync_16_powerpc, storeSwapEieio_16_powerpc, storeSwapSync_16_powerpc, swapEieio_16_powerpc, loadSwapEieio_32_powerpc, loadSwapSync_32_powerpc, storeSwapEieio_32_powerpc, storeSwapSync_32_powerpc, swapEieio_32_powerpc, eieio_powerpc, ioSync_powerpc – PowerPC specific i/o services
SYNOPSIS	<pre>#include <dki/dki.h> uint16_f loadSwapEieio_16(uint16_f * addr); unit16_f loadSwapSync_16(uint16_f * addr); void storeSwapEieio_16(uint16_f * addr, uint16_f value); void storeSwapSync_16(uint16_f * addr, uint16_f value); void swapEieio_16(uint16_f * addr); uint32_f loadSwapEieio_32(uint32_f * addr); unit32_f loadSwapSync_32(uint32_f * addr); void storeSwapEieio_32(uint32_f * addr, uint32_f value); void storeSwapSync_32(uint32_f * addr, uint32_f value); void swapEieio_32 (uint32_f * addr); void ioSync(void); void eieio(void);</pre>
FEATURES	DKI
DESCRIPTION	Provides access to PowerPC DKI I/O routines, optimized to facilitate byte swapping and synchronization of input/output.
EXTENDED DESCRIPTION	<p>The PowerPC DKI provides specific I/O routines to handle byte swapping and synchronization of I/O. These services are based on PowerPC specific instructions to store/load with byte swapping, and to enforce in-order execution of I/O.</p> <p>Typically, these services are intended to be used by a host bus driver to handle different byte ordering between the PowerPC bus and the host bus, or to enforce sequential execution of I/O. PowerPC specific I/O services are defined below as sets of routines where the <code>_xx</code> suffix indicates the bit length of the data on which the service applies. This suffix may take one of the following values:</p> <pre>_16 for 16-bit data _32 for 32-bit data</pre>

`loadSwapEieio_xx` loads data from a given address and returns the corresponding byte swapped value. A PowerPC `eieio` instruction is issued after the data is read, to enforce execution of subsequent I/O sequentially with respect to the current load. The *addr* argument specifies the address to read from.

`loadSwapSync_xx` behaves like `loadSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`storeSwapEieio_xx` stores, at a given address, the *value* byte swapped. A PowerPC `eieio` instruction is issued after the data is written, to enforce execution of subsequent I/O sequentially with respect to the current store. The *addr* argument specifies the address to write to.

`storeSwapSync_xx` behaves like `storeSwapEieio_xx` but issues a PowerPC `isync` instruction after that data is read. This ensures that all possible exceptions are taken into account before returning.

`swapEieio_xx` swaps in place the bytes of the data stored at a given address. A PowerPC `eieio` instruction is issued between reading the data and re-writing the byte-swapped data, to enforce execution in order. The *addr* argument specifies the address of the data to byte-swap.

`ioSync` enforces context synchronization, completion of all instructions, and completion of all exceptions before returning.

`eieio` enforces execution in order of I/O and is mapped to the corresponding PowerPC assembler instruction. This routine is intended for use in managing shared data structures, in doing memory-mapped I/O, and in preventing load/store combining operations in main memory. Refer to PowerPC programming manuals for more details.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service:

Services	Base level	DKI thread	Interrupt	Blocking
<code>loadSwapEieio_xx</code>	+	+	+	-
<code>loadSwapSync_xx</code>	+	+	+	-
<code>storeSwapEieio_xx</code>	+	+	+	-
<code>storeSwapSync_xx</code>	+	+	+	-
<code>swapEieio_xx</code>	+	+	+	-
<code>ioSync</code>	+	+	+	-
<code>eieio</code>	+	+	+	-

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`loadSwap_16(9DKI)` , `svDkiThreadCall(9DKI)`

NAME	usecBusyWait – precise busy wait service										
SYNOPSIS	<pre>#include <dki/dki.h> void usecBusyWait(unsigned int <i>micro</i>);</pre>										
FEATURES	DKI										
DESCRIPTION	Provides precise busy wait services.										
EXTENDED DESCRIPTION	<p>Device drivers may use precise busy wait services to wait for a very short time. Note that busy wait means that the caller waits without releasing the CPU, as if executing a busy loop. Typically, this service may be used by a driver to comply with a specific timing of its device (accessing a serial ROM), or to wait for a command to complete on the device before starting another one (resetting a device).</p> <p>usecBusyWait waits for at least <i>micro</i> micro-seconds, before returning. Note that <i>micro</i> must be in the range of 1 .. 1000. Behavior is unpredictable for values out of that range.</p>										
Allowed Calling Contexts	<p>The following table specifies the contexts in which a caller is allowed to invoke each service:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Services</th> <th style="text-align: center;">Base level</th> <th style="text-align: center;">DKI thread</th> <th style="text-align: center;">Interrupt</th> <th style="text-align: center;">Blocking</th> </tr> </thead> <tbody> <tr> <td>usecBusyWait</td> <td style="text-align: center;">+</td> <td style="text-align: center;">+</td> <td style="text-align: center;">+</td> <td style="text-align: center;">-</td> </tr> </tbody> </table>	Services	Base level	DKI thread	Interrupt	Blocking	usecBusyWait	+	+	+	-
Services	Base level	DKI thread	Interrupt	Blocking							
usecBusyWait	+	+	+	-							
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: 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										
SEE ALSO	svDkiThreadCall(9DKI)										

NAME svPhysMap, svPhysUnmap, vmMapToPhys – physical to virtual memory mapping

FEATURES DKI

DESCRIPTION The microkernel provides services to allow device drivers to map physical space to virtual memory space. This services should be used mainly by primary bus drivers to map bus I/O space, or DMA memory.

See the architecture specific man pages:

- svPhysMap_usparc(9DKI)
- svPhysUnmap_usparc(9DKI)
- vmMapToPhys_usparc(9DKI)
- svPhysMap_powerpc(9DKI)
- svPhysUnmap_powerpc(9DKI)
- vmMapToPhys_powerpc(9DKI)
- svPhysMap_x86(9DKI)
- svPhysUnmap_x86(9DKI)
- vmMapToPhys_x86(9DKI)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

NAME	svPhysMap_powerpc, svPhysUnmap_powerpc, vmMapToPhys_powerpc – PowerPC physical to virtual memory mapping						
Synopsis	<pre>#include <dki/dki.h> KnError svPhysMap(KnPhMemChunk * chunk, PteCntlBits cntlBits); void * svPhysUnmap(KnPhMemChunk * chunk); KnError vmMapToPhys(KnCap * actor, KnPhMemChunk * chunk, PteCntlBits cntlBits); typedef struct { PhAddr paddr; /* physical start addree */ PhSize psize; /* size */ VmAddr vaddr; /* virtual start address */ } KnPhMemChunk;</pre>						
FEATURES	DKI						
DESCRIPTION	Provides physical to virtual memory mapping services.						
EXTENDED DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. These services should be used mainly by primary bus drivers to map bus I/O space or DMA memory.</p> <p>svPhysMap allocates a range of virtual addresses within the supervisor address space and maps it to a given range of physical addresses.</p> <p>The <i>chunk</i> argument points to the KnPhMemChunk structure.</p> <p>The <i>paddr</i> and <i>psize</i> fields specify the chunk of physical memory being mapped to.</p> <p>The <i>vaddr</i> field is used by svPhysMap to return the virtual address to which the physical one is mapped.</p> <p>For PowerPC family processors, the <i>cntlBits</i> argument should be constructed by 'oring' the following values:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">PTE_READ_WRITE</td> <td>The mapped memory allows you to perform write accesses.</td> </tr> <tr> <td style="padding-right: 20px;">PTE_CACHE_WRITE_THROUGH</td> <td>The memory is mapped with Write-Through attribute (bit <i>w</i>). Thus a store operation updates the cache if necessary, and in addition the update is written to memory.</td> </tr> <tr> <td style="padding-right: 20px;">PTE_CACHE_DISABLE</td> <td>The memory is mapped with Cache-Inhibited attribute (bit <i>I</i>). Thus</td> </tr> </table>	PTE_READ_WRITE	The mapped memory allows you to perform write accesses.	PTE_CACHE_WRITE_THROUGH	The memory is mapped with Write-Through attribute (bit <i>w</i>). Thus a store operation updates the cache if necessary, and in addition the update is written to memory.	PTE_CACHE_DISABLE	The memory is mapped with Cache-Inhibited attribute (bit <i>I</i>). Thus
PTE_READ_WRITE	The mapped memory allows you to perform write accesses.						
PTE_CACHE_WRITE_THROUGH	The memory is mapped with Write-Through attribute (bit <i>w</i>). Thus a store operation updates the cache if necessary, and in addition the update is written to memory.						
PTE_CACHE_DISABLE	The memory is mapped with Cache-Inhibited attribute (bit <i>I</i>). Thus						

	accesses are made to main memory, bypassing the caches.
PTE_MEMORY_COHERENCY	The memory is mapped with Memory-Coherency attribute (bit M). This enforces coherency of memory shared between processors in a system. When performing an access to memory, there is a hardware indication to the rest of the system that the access is global. Other processors affected by the access must then respond to this global access. Typically, this is used for a snooping bus design.
PTE_MEMORY_GUARDED	The memory is mapped with Memory-Guarded attribute (bit G). This prevents the processor from making out-of-order access to that memory (that is, access not directly dictated by the program). This may be useful if there are holes in physical memory, or to prevent these accesses to certain peripheral devices.

If *cntlBits* is equal to zero, the memory is mapped as read-only, not guarded, with cache enabled in write-back mode, and no coherency is enforced.

Note - Any combination where PTE_CACHE_WRITE_THROUGH and PTE_CACHE_DISABLE are both set is not supported.

On success *K_OK* is returned, otherwise a negative error code is returned:

<i>K_ENOMEM</i>	The system is out of memory.
<i>K_ESIZE</i>	The <i>psize</i> argument is equal to zero.

svPhysUnmap unmaps the physical memory chunk previously mapped by *svPhysMap*. The *chunk* argument points to the *KnPhMemChunk* structure which specifies the virtual and physical chunk start addresses and the chunk size. Note that the *paddr* and *psize* fields must have the values previously specified in *svPhysMap*. The *vaddr* field must have the value previously returned by *svPhysMap*.

Note - Typically, a driver uses the same `KnPhMemChunk` structure for the `svPhysMap` and `svPhysUnmap` calls and the structure fields are not modified by the driver once `svPhysMap` has been done.

`vmMapToPhys` maps a given physical memory chunk to the target actor address space. The *actor* argument specifies the target actor capability.

If *actor* is `K_MYACTOR`, the address space of the current actor is used.

If *actor* is `K_SVACTOR`, the supervisor address space is used.

The *chunk* argument points to the `KnPhMemChunk` structure which specifies the physical and virtual chunk start addresses and the chunk size.

The specified virtual address range must be allocated using the `K_RESERVED` option (see `rgnAllocate(2)`) prior to the invocation of `vmMapToPhys`.

The specified actor can be a supervisor actor as well as a user actor.

The mapping produced by `vmMapToPhys` can only be invalidated by `rgnFree`. *vaddr*, *paddr* and *psize* must be page-aligned.

For PowerPC family processors, the *cntlBits* is defined in the same way as for the `svPhysMap` routine above.

On success `K_OK` is returned, otherwise a negative error code is returned:

<code>K_EFAULT</code>	The <i>actor</i> argument points outside the caller's address space.
<code>K_EINVAL</code>	An inconsistent actor capability was provided.
<code>K_EUNKNOWN</code>	<i>actorcap</i> does not specify a reachable actor.
<code>K_ENOMEM</code>	The system is out of memory.
<code>K_ESIZE</code>	The <i>psize</i> argument is equal to zero.
<code>K_EROUND</code>	<i>vaddr</i> or <i>paddr</i> or <i>psize</i> is not page-aligned.
<code>K_EADDR</code>	Some or all addresses from the target virtual address range are out of a region allocated with the <code>K_RESERVED</code> option.

Note - For performance reasons, the current implementation does not guarantee that any attempt to map a physical address to a virtual address out of a region allocated with `K_RESERVED` option would effectively produce a `K_EADDR` error.

Allowed Calling Contexts

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysMap	+	+	-	+
svPhysUnmap	+	+	-	+
vmMapToPhys	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

svPhysAlloc(9DKI) , svDkiThreadCall(9DKI) , rgnAllocate(2K)

NAME	svPhysMap_usparc, svPhysUnmap_usparc, vmMapToPhys_usparc – UltraSPARC physical to virtual memory mapping
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svPhysMap(KnPhMemChunk * chunk, KnPteAttr attr); void svPhysUnmap(KnPhMemChunk * chunk); KnError vmMapToPhys(KnCap * actor, KnPhMemChunk * chunk, KnPteAttr attr);</pre>
FEATURES	DKI
DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. These services should be used mainly by a host bus driver to map bus I/O space or DMA memory.</p> <p>svPhysMap allocates a range of virtual addresses within the supervisor address space and maps it to a given range of physical addresses.</p> <p>The <i>chunk</i> argument points to the KnPhMemChunk structure shown below:</p> <pre>typedef struct { PhAddr paddr; /* physical start address */ PhSize psize; /* size */ VmAddr vaddr; /* virtual start address */ } KnPhMemChunk;</pre> <p>The <i>paddr</i> and <i>psize</i> fields specify the chunk of physical memory being mapped to. The <i>vaddr</i> field is used by <code>svPhysMap</code> to return the virtual address to which the physical one is mapped.</p> <p>The <i>attr</i> argument specifies the mapping attributes. <i>attr</i> is a bit-mask composed of two independent parts:</p> <ul style="list-style-type: none"> Data access attributes Instruction access attributes <p>Basically, each part of the attributes is a sub-set of bits defined by the translation table entry (TTE) of the UltraSPARC MMU.</p> <p>A combination of the following attributes may be specified for data access:</p> <p>PTE_DATTR_G — global The PTE_DATTR_G bit set allows the mapping to be shared among all (user and supervisor) contexts.</p> <p>PTE_DATTR_W — writable The PTE_DATTR_W bit set grants write permission for the mapping.</p> <p>PTE_DATTR_P — privileged</p>

The `PTE_DATTR_P` bit set restricts access to the mapping for the supervisor only.

`PTE_DATTR_E` — side-effect

The `PTE_DATTR_E` bit set makes noncacheable memory accesses to be strongly ordered against other E-bit accesses, and noncacheable stores are not merged. This bit should be set for I/O devices having side-effects. Note that the E-bit does not force an uncacheable access. It is expected that the `PTE_DATTR_CV` and `PTE_DATTR_CP` bits will be set to zero when the E-bit is set.

`PTE_DATTR_CV` — L1-cacheable

The `PTE_DATTR_CV` bit set allows data to be cached in the (L1) CPU data cache. Note that if the `PTE_DATTR_CV` bit is set, the `PTE_DATTR_CP` bit must also be set.

`PTE_DATTR_CP` — L2-cacheable

The `PTE_DATTR_CP` bit set allows data to be cached in the (L2) external cache.

`PTE_DATTR_IE` — invert endianness

The `PTE_DATTR_IE` bit set causes data accesses to the mapping to be processed with inverse endianness from that specified by the instruction.

`PTE_DATTR_V` — valid

The `PTE_DATTR_V` bit set enables data accesses to the mapping. If this bit is not set, all other bits (described above) are ignored and a data access to the mapping will result in a data access exception.

Combinations of the following attributes may be specified for instruction access:

`PTE_IATTR_G` — global

The `PTE_IATTR_G` bit set allows the mapping to be shared among all (user and supervisor) contexts.

`PTE_IATTR_P` — privileged

The `PTE_IATTR_P` bit set restricts access to the mapping to the supervisor only.

`PTE_IATTR_CV` — L1-cacheable

The `PTE_IATTR_CV` bit set allows instructions to be cached in the (L1) CPU instruction cache. Note that if the `PTE_IATTR_CV` bit is set, the `PTE_IATTR_CP` bit must also be set.

`PTE_IATTR_CP` — L2-cacheable

The `PTE_IATTR_CP` bit set allows instructions to be cached in the (L2) external cache.

PTE_IATTR_V — valid

The `PTE_IATTR_V` bit set enables instructions to be obtained from the mapping. If this bit is not set, all other bits (described above) are ignored and an instruction to fetch from the mapping will result in an instruction access exception.

On success `svPhysMap` returns `K_OK`, otherwise a negative error code is returned as follows:

<code>K_ENOMEM</code>	The system is out of memory.
<code>K_ESIZE</code>	The <i>psize</i> argument is equal to zero.

svPhysUnmap

`svPhysUnmap` unmaps the physical memory chunk previously mapped by `svPhysMap`. The *chunk* argument points to the `KnPhMemChunk` structure which specifies the virtual and physical chunk start addresses and the chunk size. Note that the *paddr* and *psize* fields must have the values previously specified in `svPhysMap`. The *vaddr* fields must have the value previously returned by `svPhysMap`. Typically, a driver uses the same `KnPhMemChunk` structure for the `svPhysMap` and `svPhysUnmap` calls and the structure fields are not modified by the driver once `svPhysMap` has been performed.

vmMapToPhys

`vmMapToPhys` maps a given physical memory chunk to the target actor address space.

The *actor* argument specifies the target actor capability. If *actor* is `K_MYACTOR`, the address space of the current actor is used. If *actor* is `K_SVACTOR`, the supervisor address space is used.

The *chunk* argument points to the `KnPhMemChunk` structure which specifies the physical and virtual chunk start addresses and the chunk size. The specified virtual address range must be allocated using the `K_RESERVED` option (see `rgnAllocate(2K)`) prior to the invocation of `vmMapToPhys`. The specified actor can be a supervisor actor as well as an user actor. The mapping produced by `vmMapToPhys` can only be invalidated by `rgnFree`.

vaddr, *paddr* and *psize* must be page-aligned.

For UltraSPARC family processors, the *attr* argument is defined in the same way as for the `svPhysMap` routine above.

On success `K_OK` is returned, otherwise a negative error code is returned as follows:

<code>K_EFAULT</code>	The <i>actor</i> argument points outside of the caller's address space.
-----------------------	---

K_EINVAL	An inconsistent actor capability was specified.
K_EUNKNOWN	<i>actorcap</i> does not specify a reachable actor.
K_ENOMEM	The system is out of memory.
K_ESIZE	The <i>psize</i> argument is equal to zero.
K_EROUND	<i>vaddr</i> or <i>paddr</i> or <i>psize</i> is not page-aligned.
K_EADDR	Some or all addresses from the target virtual address range are out of a region allocated with the <code>K_RESERVED</code> option. For performance reasons, the current implementation does not guarantee that any attempt to map a physical address to a virtual address outside a region allocated with <code>K_RESERVED</code> option would effectively produce a <code>K_EADDR</code> error.

ALLOWED CALLING CONTEXTS

The following table specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysMap	+	+	-	+
svPhysUnmap	+	+	-	+
vmMapToPhys	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svPhysAlloc(9DKI)` , `svDkiThreadCall(9DKI)` , `rgnAllocate(2K)`

NAME	svPhysMap_x86, svPhysUnmap_x86, vmMapToPhys_x86 – Intel x86 physical to virtual memory mapping								
SYNOPSIS	<pre>#include <dki/f_dki.h> KnError svPhysMap(KnPhMemChunk * chunk, PteCntlBits cntlBits); void svPhysUnmap(KnPhMemChunk * chunk); KnError vmMapToPhys(KnCap * actor, KnPhMemChunk * chunk, PteCntlBits cntlBits); typedef struct { PhAddr paddr; /* physical start address */ PhSize psize; /* size */ VmAddr vaddr; /* virtual start address */ } KnPhMemChunk;</pre>								
FEATURES	DKI								
EXTENDED DESCRIPTION	<p>The microkernel provides services to allow device drivers to map physical space to virtual memory space. These services should be used mainly by primary bus drivers to map bus I/O space, or DMA memory.</p> <p>svPhysMap</p> <p>svPhysMap allocates a range of virtual addresses within the supervisor address space and maps it to a given range of physical addresses.</p> <p>The <i>chunk</i> argument points to the KnPhMemChunk structure.</p> <p>The <i>paddr</i> and <i>psize</i> fields specify the chunk of physical memory being mapped to.</p> <p>The <i>vaddr</i> field is used by svPhysMap to return the virtual address to which the physical one is mapped.</p> <p>For Intel ix86 family processor, the cntlBits argument is a direct mapping of the Page Control bits of Page Table Entries, as described in the Intel 386/486 Programming Reference Manual. Therefore, the cntlBits argument should be constructed by “oring” the following values:</p> <table border="0"> <tr> <td>PTE_PRESENT</td> <td>The mapped memory must be present in physical memory.</td> </tr> <tr> <td>PTE_READ_WRITE</td> <td>The mapped memory allows write accesses to be performed.</td> </tr> <tr> <td>PTE_USER_SUPERVISOR</td> <td>The mapped memory can be used in user space.</td> </tr> <tr> <td>PTE_WRITE_TRANSPARENT</td> <td>The memory is mapped with the Write-Through attribute (bit PWT). Thus, a store operation updates the cache if necessary, and in addition the update is written to memory.</td> </tr> </table>	PTE_PRESENT	The mapped memory must be present in physical memory.	PTE_READ_WRITE	The mapped memory allows write accesses to be performed.	PTE_USER_SUPERVISOR	The mapped memory can be used in user space.	PTE_WRITE_TRANSPARENT	The memory is mapped with the Write-Through attribute (bit PWT). Thus, a store operation updates the cache if necessary, and in addition the update is written to memory.
PTE_PRESENT	The mapped memory must be present in physical memory.								
PTE_READ_WRITE	The mapped memory allows write accesses to be performed.								
PTE_USER_SUPERVISOR	The mapped memory can be used in user space.								
PTE_WRITE_TRANSPARENT	The memory is mapped with the Write-Through attribute (bit PWT). Thus, a store operation updates the cache if necessary, and in addition the update is written to memory.								

PTE_CACHE_DISABLE The memory is mapped with caching disabled (bit PCD). Thus, memory accesses go to main memory, bypassing the caches.

On success K_OK is returned, otherwise a negative error code is returned:

K_ENOMEM The system is out of memory.

K_ESIZE The *psize* argument is equal to zero.

svPhysUnmap

svPhysUnmap unmaps the physical memory chunk previously mapped by svPhysMap . The *chunk* argument points to the KnPhMemChunk structure which specifies the virtual and physical chunk start addresses and the chunk size. Note that the *paddr* and *psize* fields must have the values previously specified in svPhysMap . The *vaddr* field must have the value previously returned by svPhysMap (typically, a driver uses the same KnPhMemChunk structure for the svPhysMap and svPhysUnmap calls and the structure fields are not modified by the driver once svPhysMap is done).

vmMapToPhys

vmMapToPhys maps a given physical memory chunk to the target actor address space.

The *actor* argument specifies the target actor capability. If *actor* is K_MYACTOR , the address space of the current actor is used. If *actor* is K_SVACTOR , the supervisor address space is used.

The *chunk* argument points to the KnPhMemChunk structure which specifies the physical and virtual chunk start addresses and the chunk size. The specified virtual address range must be allocated using the K_RESERVED option (see *rgnAllocate* (2K)) prior to the invocation of vmMapToPhys . The specified actor can be a supervisor actor as well as a user actor. The mapping produced by vmMapToPhys can only be invalidated by rgnFree .

vaddr , *paddr* and *psize* must be page-aligned.

For Intel x86 family processors, the *cntlBits* is defined in the same way as for the svPhysMap routine above.

On success K_OK is returned, otherwise a negative error code is returned:

K_EFAULT The *actor* argument points to the outside of the caller's address space.

K_EINVAL An inconsistent actor capability was provided.

K_EUNKNOWN *actorcap* does not specify a reachable actor.

K_ENOMEM The system is out of memory.

K_ESIZE The *psize* argument is equal to zero.

K_EROUND *vaddr* or *paddr* or *psize* is not page-aligned.

K_EADDR Some or all addresses from the target virtual address range are out of a region allocated with the K_RESERVED option (for performance reasons the current implementation does not guarantee that any attempt to map a physical address to a virtual address out of a region allocated with K_RESERVED option would effectively produce a K_EADDR error).

Intel x86 Memory Mapping Allowed Calling Contexts

The table below specifies the contexts in which a caller is allowed to invoke each service.

Services	Base level	DKI thread	Interrupt	Blocking
svPhysMap	+	+	-	+
svPhysUnmap	+	+	-	+
vmMapToPhys	+	+	-	+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`svPhysAlloc(9DKI)` , `svDkiThreadCall(9DKI)` , `rgnAllocate(9DKI)`

Index

D

- dataCacheBlockFlush — cache management 71, 75, 79, 225, 229
- dataCacheBlockFlush_powerpc — PowerPC cache management 72, 76, 80, 226, 230
- dataCacheBlockInvalidate — cache management 71, 75, 79, 225, 229
- dataCacheBlockInvalidate_powerpc — PowerPC cache management 72, 76, 80, 226, 230
- dataCacheInvalidate — cache management 71, 75, 79, 225, 229
- dataCacheInvalidate_powerpc — PowerPC cache management 72, 76, 80, 226, 230
- dcacheBlockFlush — cache management 83, 87, 91, 211, 215, 219
- dcacheBlockFlush_usparc — UltraSPARC cache management 84, 88, 92, 212, 216, 220
- dcacheFlush — cache management 83, 87, 91, 211, 215, 219
- dcacheFlush_usparc — UltraSPARC cache management 84, 88, 92, 212, 216, 220
- dcacheLineFlush — cache management 83, 87, 91, 211, 215, 219
- dcacheLineFlush_usparc — UltraSPARC cache management 84, 88, 92, 212, 216, 220
- DISABLE_PREEMPT — thread preemption disabling 95, 209
- dtreeNodeAdd — device tree operations 97, 103, 109, 115, 121, 127, 133, 139, 145, 151, 157, 163, 169, 175, 181, 187, 193, 199
- dtreeNodeAlloc — device tree operations 97, 103, 109, 115, 121, 127, 133, 139, 145, 151, 157, 163, 169, 175, 181, 187, 193, 199
- dtreeNodeAttach — device tree operations 97, 103, 109, 115, 121, 127, 133, 139, 145, 151, 157, 163, 169, 175, 181, 187, 193, 199
- dtreeNodeChild — device tree operations 97, 103, 109, 115, 121, 127, 133, 139, 145, 151, 157, 163, 169, 175, 181, 187, 193, 199
- dtreeNodeDetach — device tree operations 97, 103, 109, 115, 121, 127, 133, 139, 145, 151, 157, 163, 169, 175, 181, 187, 193, 199
- dtreeNodeFind — device tree operations 97, 103, 109, 115, 121, 127, 133, 139, 145, 151, 157, 163, 169, 175, 181, 187, 193, 199
- dtreeNodeFree — device tree operations 97, 103, 109, 115, 121, 127, 133,

139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreeNodeParent` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreeNodePeer` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreeNodeRoot` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreePropAdd` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreePropAlloc` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreePropAttach` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreePropDetach` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreePropFind` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreePropFindNext` — device tree
operations 97, 103,
109, 115, 121, 127, 133, 139,
145, 151, 157, 163, 169, 175,
181, 187, 193, 199

`dtreePropFree` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreePropLength` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreePropName` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

`dtreePropValue` — device tree operations 97,
103, 109, 115, 121, 127, 133,
139, 145, 151, 157, 163, 169,
175, 181, 187, 193, 199

E

`eieio` — i/o services 205, 275, 279, 310, 314,
558, 562

`eieio_powerpc` — PowerPC specific i/o
services 206, 276, 280, 311,
315, 559, 563

`ENABLE_PREEMPT` — thread preemption
enabling 95, 209

I

`icacheBlockInval` — cache management 83, 87,
91, 211, 215, 219

`icacheBlockInval_usparc` — UltraSPARC cache
management 84, 88, 92, 212,
216, 220

`icacheInval` — cache management 83, 87, 91,
211, 215, 219

`icacheInval_usparc` — UltraSPARC cache
management 84, 88, 92, 212,
216, 220

`icacheLineInval` — cache management 83, 87,
91, 211, 215, 219

`icacheLineInval_usparc` — UltraSPARC cache
management 84, 88, 92, 212,
216, 220

`imsIntrMask_f` — global interrupts
masking 223–224

`imsIntrUnmask_f` — global interrupts
masking 223–224

`instCacheBlockInvalidate` — cache
management 71, 75,
79, 225, 229

`instCacheBlockInvalidate_powerpc`
— PowerPC cache
management 72, 76, 80, 226,
230

instCacheInvalidate — cache management 71, 75, 79, 225, 229
 instCacheInvalidate_powerpc — PowerPC cache management 72, 76, 80, 226, 230
 intro — driver kernel interface introduction 63
 ioLoad16 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioLoad16_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioLoad32 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioLoad32_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioLoad8 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioLoad8_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioRead16 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioRead16_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioRead32 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioRead32_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioRead8 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioRead8_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioStore16 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioStore16_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioStore32 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioStore32_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioStore8 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioStore8_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioWrite16 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioWrite16_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioWrite32 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioWrite32_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267
 ioWrite8 — I/O services 233, 236, 239, 242, 245, 248, 251, 254, 257, 260, 263, 266
 ioWrite8_x86 — Intel x86 specific I/O services 234, 237, 240, 243, 246, 249, 252, 255, 258, 261, 264, 267

L

load_sync_16_usparc — UltraSparc specific i/o services 283, 286, 289, 292,

295, 298, 301, 318, 321, 324,
 327, 330, 333, 336
 load_sync_32_usparc — UltraSparc specific i/o
 services 283, 286, 289, 292,
 295, 298, 301, 318, 321, 324,
 327, 330, 333, 336
 load_sync_64_usparc — UltraSparc specific i/o
 services 283, 286, 289, 292,
 295, 298, 301, 318, 321, 324,
 327, 330, 333, 336
 load_sync_8_usparc — UltraSparc specific i/o
 services 283, 286, 289, 292,
 295, 298, 301, 318, 321, 324,
 327, 330, 333, 336
 loadSwap_16 — specific i/o services 269, 271,
 273, 304, 306, 308, 552, 554, 556
 loadSwap_32 — specific i/o services 269, 271,
 273, 304, 306, 308, 552, 554, 556
 loadSwap_64 — specific i/o services 269, 271,
 273, 304, 306, 308, 552, 554, 556
 loadSwap_sync_16_usparc — UltraSparc
 specific i/o services 283, 286,
 289, 292, 295, 298, 301, 318,
 321, 324, 327, 330, 333, 336
 loadSwap_sync_32_usparc — UltraSparc
 specific i/o services 283, 286,
 289, 292, 295, 298, 301, 318,
 321, 324, 327, 330, 333, 336
 loadSwap_sync_64_usparc — UltraSparc
 specific i/o services 283, 286,
 289, 292, 295, 298, 301, 318,
 321, 324, 327, 330, 333, 336
 loadSwapEieio_16 — i/o services 205, 275,
 279, 310, 314, 558, 562
 loadSwapEieio_16_powerpc — PowerPC
 specific i/o services 206, 276,
 280, 311, 315, 559, 563
 loadSwapEieio_32 — i/o services 205, 275,
 279, 310, 314, 558, 562
 loadSwapEieio_32_powerpc — PowerPC
 specific i/o services 206, 276,
 280, 311, 315, 559, 563
 loadSwapSync_16 — PowerPC specific i/o
 services 206, 276, 280, 311,
 315, 559, 563

loadSwapSync_32 — PowerPC specific i/o
 services 206, 276, 280, 311,
 315, 559, 563

S

store_sync_16_usparc — UltraSparc specific i/o
 services 283, 286, 289, 292,
 295, 298, 301, 318, 321, 324,
 327, 330, 333, 336
 store_sync_32_usparc — UltraSparc specific i/o
 services 283, 286, 289, 292,
 295, 298, 301, 318, 321, 324,
 327, 330, 333, 336
 store_sync_64_usparc — UltraSparc specific i/o
 services 283, 286, 289, 292,
 295, 298, 301, 318, 321, 324,
 327, 330, 333, 336
 store_sync_8_usparc — UltraSparc specific i/o
 services 283, 286, 289, 292,
 295, 298, 301, 318, 321, 324,
 327, 330, 333, 336
 storeSwap_16 — specific i/o services 269, 271,
 273, 304, 306, 308, 552, 554, 556
 storeSwap_32 — specific i/o services 269, 271,
 273, 304, 306, 308, 552, 554, 556
 storeSwap_64 — specific i/o services 269, 271,
 273, 304, 306, 308, 552, 554, 556
 storeSwap_sync_16_usparc — UltraSparc
 specific i/o services 283, 286,
 289, 292, 295, 298, 301, 318,
 321, 324, 327, 330, 333, 336
 storeSwap_sync_32_usparc — UltraSparc
 specific i/o services 283, 286,
 289, 292, 295, 298, 301, 318,
 321, 324, 327, 330, 333, 336
 storeSwap_sync_64_usparc — UltraSparc
 specific i/o services 283, 286,
 289, 292, 295, 298, 301, 318,
 321, 324, 327, 330, 333, 336
 storeSwapEieio_16 — i/o services 205, 275,
 279, 310, 314, 558, 562
 storeSwapEieio_16_powerpc — PowerPC
 specific i/o services 206, 276,
 280, 311, 315, 559, 563
 storeSwapEieio_32 — i/o services 205, 275,
 279, 310, 314, 558, 562

storeSwapEieio_32_powerpc — PowerPC specific i/o services 206, 276, 280, 311, 315, 559, 563

storeSwapSync_16 — PowerPC specific i/o services 206, 276, 280, 311, 315, 559, 563

storeSwapSync_32 — PowerPC specific i/o services 206, 276, 280, 311, 315, 559, 563

svAsyncExcepAttach — asynchronous exceptions management 339, 343

svAsyncExcepAttach_usparc — UltraSPARC asynchronous exceptions management 340, 344

svAsyncExcepDetach — asynchronous exceptions management 339, 343

svAsyncExcepDetach_usparc — UltraSPARC asynchronous exceptions management 340, 344

svDeviceAlloc — device registry operations 347, 354, 361, 368, 375, 382, 389, 396

svDeviceEntry — device registry operations 347, 354, 361, 368, 375, 382, 389, 396

svDeviceEvent — device registry operations 347, 354, 361, 368, 375, 382, 389, 396

svDeviceFree — device registry operations 347, 354, 361, 368, 375, 382, 389, 396

svDeviceLookup — device registry operations 347, 354, 361, 368, 375, 382, 389, 396

svDeviceRegister — device registry operations 347, 354, 361, 368, 375, 382, 389, 396

svDeviceRelease — device registry operations 347, 354, 361, 368, 375, 382, 389, 396

svDeviceUnregister — device registry operations 347, 354, 361, 368, 375, 382, 389, 396

svDkiClose — system event management 403, 406, 409

svDkiEvent — system event management 403, 406, 409

svDkiOpen — system event management 403, 406, 409

svDkiThreadCall — call a routine in the DKI thread context 412, 414

svDkiThreadTrigger — call a routine in the DKI thread context 412, 414

svDriverCap — driver registry operations 416, 422, 428, 434, 440, 446, 452

svDriverEntry — driver registry operations 416, 422, 428, 434, 440, 446, 452

svDriverLookupFirst — driver registry operations 416, 422, 428, 434, 440, 446, 452

svDriverLookupNext — driver registry operations 416, 422, 428, 434, 440, 446, 452

svDriverRegister — driver registry operations 416, 422, 428, 434, 440, 446, 452

svDriverRelease — driver registry operations 416, 422, 428, 434, 440, 446, 452

svDriverUnregister — driver registry operations 416, 422, 428, 434, 440, 446, 452

svIntrAttach — interrupts management 458, 472, 486

svIntrAttach_powerpc — PowerPC interrupts management 459, 473, 487

svIntrAttach_usparc — UltraSPARC interrupts management 462, 476, 490, 534, 540

svIntrAttach_x86 — Intel x86 interrupts management 468, 482, 496

svIntrCtxGet — interrupts management 458, 472, 486

svIntrCtxGet_powerpc — PowerPC interrupts management 459, 473, 487

svIntrCtxGet_usparc — UltraSPARC interrupts management 462, 476, 490, 534, 540

svIntrCtxGet_x86 — Intel x86 interrupts management 468, 482, 496

svIntrDetach — interrupts management 458, 472, 486

svIntrDetach_powerpc — PowerPC interrupts management 459, 473, 487

svIntrDetach_usparc — UltraSPARC interrupts management 462, 476, 490, 534, 540

svIntrDetach_x86 — Intel x86 interrupts management 468, 482, 496

svMemAlloc — A general purpose memory allocator 500, 502

svMemFree — A general purpose memory allocator 500, 502

svPhysAlloc — A special purpose physical memory allocator 504, 507

svPhysFree — A special purpose physical memory allocator 504, 507

svPhysMap — physical to virtual memory mapping 510, 522, 567

svPhysMap_powerpc — PowerPC physical to virtual memory mapping 511, 523, 568

svPhysMap_usparc — UltraSPARC physical to virtual memory mapping 515, 527, 572

svPhysMap_x86 — Intel x86 physical to virtual memory mapping 519, 531, 576

svPhysUnmap — physical to virtual memory mapping 510, 522, 567

svPhysUnmap_powerpc — PowerPC physical to virtual memory mapping 511, 523, 568

svPhysUnmap_usparc — UltraSPARC physical to virtual memory mapping 515, 527, 572

svPhysUnmap_x86 — Intel x86 physical to virtual memory mapping 519, 531, 576

svSoftIntrAttach — interrupts management 458, 472, 486

svSoftIntrAttach_usparc — UltraSPARC interrupts management 462, 476, 490, 534, 540

svSoftIntrDetach — interrupts management 458, 472, 486

svSoftIntrDetach_usparc — UltraSPARC interrupts management 462, 476, 490, 534, 540

svTimeoutCancel — timeout operations 546, 548, 550

svTimeoutGetRes — timeout operations 546, 548, 550

svTimeoutSet — timeout operations 546, 548, 550

svTimerIntrAttach — interrupts management 458, 472, 486

svTimerIntrAttach_uparc — UltraSPARC interrupts management 462, 476, 490, 534, 540

svTimerIntrDetach — interrupts management 458, 472, 486

svTimerIntrDetach_uparc — UltraSPARC interrupts management 462, 476, 490, 534, 540

swap_16 — specific i/o services 269, 271, 273, 304, 306, 308, 552, 554, 556

swap_32 — specific i/o services 269, 271, 273, 304, 306, 308, 552, 554, 556

swap_64 — specific i/o services 269, 271, 273, 304, 306, 308, 552, 554, 556

swapEieio_16 — i/o services 205, 275, 279, 310, 314, 558, 562

swapEieio_16_powerpc — PowerPC specific i/o services 206, 276, 280, 311, 315, 559, 563

swapEieio_32 — i/o services 205, 275, 279, 310, 314, 558, 562

swapEieio_32_powerpc — PowerPC specific i/o services 206, 276, 280, 311, 315, 559, 563

U

usecBusyWait — precise busy wait service 566

V

vmMapToPhys — physical to virtual memory
mapping 510, 522, 567
vmMapToPhys_powerpc — PowerPC
physical to virtual memory
mapping 511, 523, 568

vmMapToPhys_usparc — UltraSPARC
physical to virtual memory
mapping 515, 527, 572
vmMapToPhys_x86 — Intel x86 physical
to virtual memory
mapping 519, 531, 576