

```

*****
113437 Tue Nov 24 09:35:06 2015
new/usr/src/uts/common/vm/seg_dev.c
6151 use NULL setpagesize segop as a shorthand for ENOTSUP
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */

22 /*
23  * Copyright 2010 Sun Microsystems, Inc. All rights reserved.
24  * Use is subject to license terms.
25  */

27 /*      Copyright (c) 1984, 1986, 1987, 1988, 1989 AT&T */
28 /*      All Rights Reserved      */

30 /*
31  * University Copyright- Copyright (c) 1982, 1986, 1988
32  * The Regents of the University of California
33  * All Rights Reserved
34  *
35  * University Acknowledgment- Portions of this document are derived from
36  * software developed by the University of California, Berkeley, and its
37  * contributors.
38  */

40 /*
41  * VM - segment of a mapped device.
42  *
43  * This segment driver is used when mapping character special devices.
44  */

46 #include <sys/types.h>
47 #include <sys/t_lock.h>
48 #include <sys/sysmacros.h>
49 #include <sys/vtrace.h>
50 #include <sys/system.h>
51 #include <sys/vmsystem.h>
52 #include <sys/mman.h>
53 #include <sys/errno.h>
54 #include <sys/kmem.h>
55 #include <sys/cmn_err.h>
56 #include <sys/vnode.h>
57 #include <sys/proc.h>
58 #include <sys/conf.h>
59 #include <sys/debug.h>
60 #include <sys/ddidevmap.h>
61 #include <sys/ddi_implfuncs.h>

```

```

62 #include <sys/lgrp.h>

64 #include <vm/page.h>
65 #include <vm/hat.h>
66 #include <vm/as.h>
67 #include <vm/seg.h>
68 #include <vm/seg_dev.h>
69 #include <vm/seg_kp.h>
70 #include <vm/seg_kmem.h>
71 #include <vm/vpage.h>

73 #include <sys/sunddi.h>
74 #include <sys/esunddi.h>
75 #include <sys/fs/snodel.h>

78 #if DEBUG
79 int segdev_debug;
80 #define DEBUGF(level, args) { if (segdev_debug >= (level)) cmn_err args; }
81 #else
82 #define DEBUGF(level, args)
83 #endif

85 /* Default timeout for devmap context management */
86 #define CTX_TIMEOUT_VALUE 0

88 #define HOLD_DHP_LOCK(dhp) if (dhp->dh_flags & DEVMAP_ALLOW_REMAP) \
89     { mutex_enter(&dhp->dh_lock); }

91 #define RELE_DHP_LOCK(dhp) if (dhp->dh_flags & DEVMAP_ALLOW_REMAP) \
92     { mutex_exit(&dhp->dh_lock); }

94 #define round_down_p2(a, s)    ((a) & ~((s) - 1))
95 #define round_up_p2(a, s)    (((a) + (s) - 1) & ~((s) - 1))

97 /*
98  * VA_PA_ALIGNED checks to see if both VA and PA are on pgsz boundary
99  * VA_PA_PGSIZE_ALIGNED check to see if VA is aligned with PA w.r.t. pgsz
100 */
101 #define VA_PA_ALIGNED(uvaddr, paddr, pgsz) \
102     (((uvaddr | paddr) & (pgsz - 1)) == 0)
103 #define VA_PA_PGSIZE_ALIGNED(uvaddr, paddr, pgsz) \
104     (((uvaddr ^ paddr) & (pgsz - 1)) == 0)

106 #define vpgtob(n)            ((n) * sizeof(struct vpage)) /* For brevity */

108 #define VTOCVP(vp)          (VTOS(vp)->s_commonvp) /* we "know" it's an snode */

110 static struct devmap_ctx *devmapctx_list = NULL;
111 static struct devmap_softlock *devmap_slist = NULL;

113 /*
114  * mutex, vnode and page for the page of zeros we use for the trash mappings.
115  * One trash page is allocated on the first ddi_umem_setup call that uses it
116  * XXX Eventually, we may want to combine this with what segnf does when all
117  * hat layers implement HAT_NOFAULT.
118  *
119  * The trash page is used when the backing store for a userland mapping is
120  * removed but the application semantics do not take kindly to a SIGBUS.
121  * In that scenario, the applications pages are mapped to some dummy page
122  * which returns garbage on read and writes go into a common place.
123  * (Perfect for NO_FAULT semantics)
124  * The device driver is responsible to communicating to the app with some
125  * other mechanism that such remapping has happened and the app should take
126  * corrective action.
127  * We can also use an anonymous memory page as there is no requirement to

```

```

128 * keep the page locked, however this complicates the fault code. RFE.
129 */
130 static struct vnode trashvp;
131 static struct page *trashpp;

133 /* Non-pageable kernel memory is allocated from the umem_np_arena. */
134 static vmem_t *umem_np_arena;

136 /* Set the cookie to a value we know will never be a valid umem_cookie */
137 #define DEVMAP_DEVMEM_COOKIE ((ddi_umem_cookie_t)0x1)

139 /*
140 * Macros to check if type of devmap handle
141 */
142 #define cookie_is_devmem(c) \
143     ((c) == (struct ddi_umem_cookie *)DEVMAP_DEVMEM_COOKIE)

145 #define cookie_is_pmem(c) \
146     ((c) == (struct ddi_umem_cookie *)DEVMAP_PMEM_COOKIE)

148 #define cookie_is_kpmem(c) (!cookie_is_devmem(c) && !cookie_is_pmem(c) && \
149     ((c)->type == KMEM_PAGEABLE))

151 #define dhp_is_devmem(dhp) \
152     (cookie_is_devmem((struct ddi_umem_cookie *)((dhp)->dh_cookie)))

154 #define dhp_is_pmem(dhp) \
155     (cookie_is_pmem((struct ddi_umem_cookie *)((dhp)->dh_cookie)))

157 #define dhp_is_kpmem(dhp) \
158     (cookie_is_kpmem((struct ddi_umem_cookie *)((dhp)->dh_cookie)))

160 /*
161 * Private seg op routines.
162 */
163 static int segdev_dup(struct seg *, struct seg *);
164 static int segdev_unmap(struct seg *, caddr_t, size_t);
165 static void segdev_free(struct seg *);
166 static faultcode_t segdev_fault(struct hat *, struct seg *, caddr_t, size_t,
167     enum fault_type, enum seg_rw);
168 static faultcode_t segdev_faulta(struct seg *, caddr_t);
169 static int segdev_setprot(struct seg *, caddr_t, size_t, uint_t);
170 static int segdev_checkprot(struct seg *, caddr_t, size_t, uint_t);
171 static void segdev_badop(void);
172 static int segdev_sync(struct seg *, caddr_t, size_t, int, uint_t);
173 static size_t segdev_incore(struct seg *, caddr_t, size_t, char *);
174 static int segdev_lockop(struct seg *, caddr_t, size_t, int, int,
175     ulong_t *, size_t);
176 static int segdev_getprot(struct seg *, caddr_t, size_t, uint_t *);
177 static u_offset_t segdev_getoffset(struct seg *, caddr_t);
178 static int segdev_gettype(struct seg *, caddr_t);
179 static int segdev_getvp(struct seg *, caddr_t, struct vnode **);
180 static int segdev_advise(struct seg *, caddr_t, size_t, uint_t);
181 static void segdev_dump(struct seg *);
182 static int segdev_pagelock(struct seg *, caddr_t, size_t,
183     struct page ***, enum lock_type, enum seg_rw);
184 static int segdev_setpagesize(struct seg *, caddr_t, size_t, uint_t);
184 static int segdev_getmemid(struct seg *, caddr_t, memid_t *);

186 /*
187 * XXX this struct is used by rootnex_map_fault to identify
188 * the segment it has been passed. So if you make it
189 * "static" you'll need to fix rootnex_map_fault.
190 */
191 struct seg_ops segdev_ops = {
192     .dup = segdev_dup,

```

```

193     .unmap = segdev_unmap,
194     .free = segdev_free,
195     .fault = segdev_fault,
196     .faulta = segdev_faulta,
197     .setprot = segdev_setprot,
198     .checkprot = segdev_checkprot,
199     .kluster = (int (*)( ))segdev_badop,
200     .sync = segdev_sync,
201     .incore = segdev_incore,
202     .lockop = segdev_lockop,
203     .getprot = segdev_getprot,
204     .getoffset = segdev_getoffset,
205     .gettype = segdev_gettype,
206     .getvp = segdev_getvp,
207     .advise = segdev_advise,
208     .dump = segdev_dump,
209     .pagelock = segdev_pagelock,
211     .setpagesize = segdev_setpagesize,
210     .getmemid = segdev_getmemid,
211 };
    unchanged_portion_omitted

2465 /*ARGSUSED*/
2466 static int
2467 segdev_pagelock(struct seg *seg, caddr_t addr, size_t len,
2468     struct page ***ppp, enum lock_type type, enum seg_rw rw)
2469 {
2470     TRACE_0(TR_FAC_DEVMAP, TR_DEVMAP_PAGELOCK,
2471     "segdev_pagelock:start");
2474     return (ENOTSUP);
2475 }

2477 /*ARGSUSED*/
2478 static int
2479 segdev_setpagesize(struct seg *seg, caddr_t addr, size_t len,
2480     uint_t szc)
2481 {
2472     return (ENOTSUP);
2473 }
    unchanged_portion_omitted

```

83686 Tue Nov 24 09:35:06 2015

new/usr/src/uts/common/vm/seg_spt.c

6151 use NULL setpagesize segop as a shorthand for ENOTSUP

unchanged portion omitted

```

114 static int segspt_shmdup(struct seg *seg, struct seg *newseg);
115 static int segspt_shmunmap(struct seg *seg, caddr_t raddr, size_t ssize);
116 static void segspt_shmfree(struct seg *seg);
117 static faultcode_t segspt_shmfault(struct hat *hat, struct seg *seg,
118     caddr_t addr, size_t len, enum fault_type type, enum seg_rw rw);
119 static faultcode_t segspt_shmfaultra(struct seg *seg, caddr_t addr);
120 static int segspt_shmsetprot(register struct seg *seg, register caddr_t addr,
121     register size_t len, register uint_t prot);
122 static int segspt_shmcheckprot(struct seg *seg, caddr_t addr, size_t size,
123     uint_t prot);
124 static int segspt_shmkluster(struct seg *seg, caddr_t addr, ssize_t delta);
125 static size_t segspt_shmswapout(struct seg *seg);
126 static size_t segspt_shmincore(struct seg *seg, caddr_t addr, size_t len,
127     register char *vec);
128 static int segspt_shmsync(struct seg *seg, register caddr_t addr, size_t len,
129     int attr, uint_t flags);
130 static int segspt_shmlockop(struct seg *seg, caddr_t addr, size_t len,
131     int attr, int op, ulong_t *lockmap, size_t pos);
132 static int segspt_shmgetprot(struct seg *seg, caddr_t addr, size_t len,
133     uint_t *protv);
134 static u_offset_t segspt_shmgetoffset(struct seg *seg, caddr_t addr);
135 static int segspt_shmgettype(struct seg *seg, caddr_t addr);
136 static int segspt_shmgetvp(struct seg *seg, caddr_t addr, struct vnode **vpp);
137 static int segspt_shmadvise(struct seg *seg, caddr_t addr, size_t len,
138     uint_t behav);
139 static void segspt_shmdump(struct seg *seg);
140 static int segspt_shmpagelock(struct seg *, caddr_t, size_t,
141     struct page ***, enum lock_type, enum seg_rw);
142 static int segspt_shmsetpgsz(struct seg *, caddr_t, size_t, uint_t);
143 static int segspt_shmgetmemid(struct seg *, caddr_t, memid_t *);
144 static lgrp_mem_policy_info_t *segspt_shmgetpolicy(struct seg *, caddr_t);

```

```

145 struct seg_ops segspt_shmops = {
146     .dup          = segspt_shmdup,
147     .unmap       = segspt_shmunmap,
148     .free        = segspt_shmfree,
149     .fault       = segspt_shmfault,
150     .faultra    = segspt_shmfaultra,
151     .setprot     = segspt_shmsetprot,
152     .checkprot  = segspt_shmcheckprot,
153     .kluster    = segspt_shmkluster,
154     .swapout    = segspt_shmswapout,
155     .sync       = segspt_shmsync,
156     .incore     = segspt_shmincore,
157     .lockop     = segspt_shmlockop,
158     .getprot    = segspt_shmgetprot,
159     .getoffset  = segspt_shmgetoffset,
160     .gettype    = segspt_shmgettype,
161     .getvp      = segspt_shmgetvp,
162     .advise     = segspt_shmadvise,
163     .dump       = segspt_shmdump,
164     .pagelock   = segspt_shmpagelock,
165     .setpagesize = segspt_shmsetpgsz,
166     .getmemid  = segspt_shmgetmemid,
167     .getpolicy  = segspt_shmgetpolicy,
168 };

```

unchanged portion omitted

3011 /*ARGSUSED*/

```

3012 void
3013 segspt_shmdump(struct seg *seg)
3014 {
3015     /* no-op for ISM segment */
3016 }

3020 /*ARGSUSED*/
3021 static faultcode_t
3022 segspt_shmsetpgsz(struct seg *seg, caddr_t addr, size_t len, uint_t szc)
3023 {
3024     return (ENOTSUP);
3025 }

```

unchanged portion omitted

```
*****  
54536 Tue Nov 24 09:35:06 2015  
new/usr/src/uts/common/vm/vm_seg.c  
6151 use NULL setpagesize segop as a shorthand for ENOTSUP  
*****  
_____unchanged_portion_omitted_____
```

```
1977 int  
1978 segop_setpagesize(struct seg *seg, caddr_t addr, size_t len, uint_t szc)  
1979 {  
1980     if (seg->s_ops->setpagesize == NULL)  
1981         return (ENOTSUP);
```

```
1983 #endif /* ! codereview */  
1984     return (seg->s_ops->setpagesize(seg, addr, len, szc));  
1985 }
```

```
1987 int  
1988 segop_getmemid(struct seg *seg, caddr_t addr, memid_t *mp)  
1989 {  
1990     if (seg->s_ops->getmemid == NULL)  
1991         return (ENODEV);  
  
1993     return (seg->s_ops->getmemid(seg, addr, mp));  
1994 }
```

```
1996 struct lgrp_mem_policy_info *  
1997 segop_getpolicy(struct seg *seg, caddr_t addr)  
1998 {  
1999     if (seg->s_ops->getpolicy == NULL)  
2000         return (NULL);  
  
2002     return (seg->s_ops->getpolicy(seg, addr));  
2003 }
```

```
2005 int  
2006 segop_capable(struct seg *seg, segcapability_t cap)  
2007 {  
2008     if (seg->s_ops->capable == NULL)  
2009         return (0);  
  
2011     return (seg->s_ops->capable(seg, cap));  
2012 }
```

```
2014 int  
2015 segop_inherit(struct seg *seg, caddr_t addr, size_t len, uint_t op)  
2016 {  
2017     if (seg->s_ops->inherit == NULL)  
2018         return (ENOTSUP);  
  
2020     return (seg->s_ops->inherit(seg, addr, len, op));  
2021 }
```

```
*****
16536 Tue Nov 24 09:35:06 2015
new/usr/src/uts/i86xpv/vm/seg_mf.c
6151 use NULL setpagesize segop as a shorthand for ENOTSUP
*****
```

unchanged_portion_omitted_

```
485 /*ARGSUSED*/
486 static int
487 segmf_setpagesize(struct seg *seg, caddr_t addr, size_t len, uint_t szc)
488 {
489     return (ENOTSUP);
490 }
```

```
485 static int
486 segmf_getmemid(struct seg *seg, caddr_t addr, memid_t *memid)
487 {
488     struct segmf_data *data = seg->s_data;
490     memid->val[0] = (uintptr_t)VTOCVP(data->vp);
491     memid->val[1] = (uintptr_t)seg_page(seg, addr);
492     return (0);
493 }
```

unchanged_portion_omitted_

```
739 static struct seg_ops segmf_ops = {
740     .dup           = segmf_dup,
741     .unmap        = segmf_unmap,
742     .free         = segmf_free,
743     .fault        = segmf_fault,
744     .faulta       = segmf_faulta,
745     .setprot      = segmf_setprot,
746     .checkprot    = segmf_checkprot,
747     .kluster      = segmf_kluster,
748     .sync         = segmf_sync,
749     .incore       = segmf_inc core,
750     .lockop       = segmf_lockop,
751     .getprot      = segmf_getprot,
752     .getoffset    = segmf_getoffset,
753     .gettype      = segmf_gettype,
754     .getvp        = segmf_getvp,
755     .advise       = segmf_advise,
756     .dump         = segmf_dump,
757     .pagelock     = segmf_pagelock,
765     .setpagesize  = segmf_setpagesize,
758     .getmemid     = segmf_getmemid,
759 };
```

unchanged_portion_omitted_

```

*****
11779 Tue Nov 24 09:35:06 2015
new/usr/src/uts/sparc/v9/vm/seg_nf.c
6151 use NULL setpagesize segop as a shorthand for ENOTSUP
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */
21 /*
22 * Copyright 2006 Sun Microsystems, Inc. All rights reserved.
23 * Use is subject to license terms.
24 */

26 /* Copyright (c) 1983, 1984, 1985, 1986, 1987, 1988, 1989 AT&T */
27 /* All Rights Reserved */

29 /*
30 * Portions of this source code were derived from Berkeley 4.3 BSD
31 * under license from the Regents of the University of California.
32 */

34 /*
35 * VM - segment for non-faulting loads.
36 */

38 #include <sys/types.h>
39 #include <sys/t_lock.h>
40 #include <sys/param.h>
41 #include <sys/mman.h>
42 #include <sys/errno.h>
43 #include <sys/kmem.h>
44 #include <sys/cmn_err.h>
45 #include <sys/vnode.h>
46 #include <sys/proc.h>
47 #include <sys/conf.h>
48 #include <sys/debug.h>
49 #include <sys/archsystem.h>
50 #include <sys/lgrp.h>

52 #include <vm/page.h>
53 #include <vm/hat.h>
54 #include <vm/as.h>
55 #include <vm/seg.h>
56 #include <vm/vpage.h>

58 /*
59 * Private seg op routines.
60 */
61 static int      segnf_dup(struct seg *seg, struct seg *newseg);

```

```

62 static int      segnf_unmap(struct seg *seg, caddr_t addr, size_t len);
63 static void     segnf_free(struct seg *seg);
64 static faultcode_t segnf_nomap(void);
65 static int      segnf_setprot(struct seg *seg, caddr_t addr,
66                             size_t len, uint_t prot);
67 static int      segnf_checkprot(struct seg *seg, caddr_t addr,
68                                size_t len, uint_t prot);
69 static void     segnf_badop(void);
70 static int      segnf_nop(void);
71 static int      segnf_getprot(struct seg *seg, caddr_t addr,
72                               size_t len, uint_t *protv);
73 static u_offset_t segnf_getoffset(struct seg *seg, caddr_t addr);
74 static int      segnf_gettype(struct seg *seg, caddr_t addr);
75 static int      segnf_getvp(struct seg *seg, caddr_t addr, struct vnode **vpp);
76 static void     segnf_dump(struct seg *seg);
77 static int      segnf_pagelock(struct seg *seg, caddr_t addr, size_t len,
78                               struct page ***ppp, enum lock_type type, enum seg_rw rw);
79 static int      segnf_setpagesize(struct seg *seg, caddr_t addr, size_t len,
80                                  uint_t szc);

81 struct seg_ops segnf_ops = {
82     .dup          = segnf_dup,
83     .unmap        = segnf_unmap,
84     .free         = segnf_free,
85     .fault        = (faultcode_t (*)(struct hat *, struct seg *, caddr_t,
86                                     size_t, enum fault_type, enum seg_rw))segnf_nomap,
87     .faulta       = (faultcode_t (*)(struct seg *, caddr_t)) segnf_nomap,
88     .setprot      = segnf_setprot,
89     .checkprot    = segnf_checkprot,
90     .kluster      = (int (*)())segnf_badop,
91     .sync         = (int (*)(struct seg *, caddr_t, size_t, int, uint_t))
92                   segnf_nop,
93     .incore       = (size_t (*)(struct seg *, caddr_t, size_t, char *))
94                   segnf_nop,
95     .lockop       = (int (*)(struct seg *, caddr_t, size_t, int, int,
96                             ulong_t *, size_t))segnf_nop,
97     .getprot      = segnf_getprot,
98     .getoffset    = segnf_getoffset,
99     .gettype      = segnf_gettype,
100    .getvp         = segnf_getvp,
101    .advise        = (int (*)(struct seg *, caddr_t, size_t, uint_t))
102                   segnf_nop,
103    .dump          = segnf_dump,
104    .pagelock      = segnf_pagelock,
105    .setpagesize   = segnf_setpagesize,
106 };
107
108 unchanged_portion_omitted
109
149 /*
150 * segnf pages are not dumped, so we just return
151 */
152 /* ARGSUSED */
153 static void
154 segnf_dump(struct seg *seg)
155 {}

157 /*ARGSUSED*/
158 static int
159 segnf_pagelock(struct seg *seg, caddr_t addr, size_t len,
160               struct page ***ppp, enum lock_type type, enum seg_rw rw)
161 {
162     return (ENOTSUP);
163 }

164 /*ARGSUSED*/

```

```
469 static int
470 segnf_setpagesize(struct seg *seg, caddr_t addr, size_t len,
471                 uint_t szc)
461 {
462     return (ENOTSUP);
463 }
unchanged_portion_omitted
```