

```

*****
113260 Tue Nov 24 09:35:08 2015
new/usr/src/uts/common/vm/seg_dev.c
6152 use NULL dump segop as a shorthand for no-op
*****
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/snoder.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 snoder */

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_getmemid(struct seg *, caddr_t, memid_t *);

185 /*
186 * XXX this struct is used by rootnex_map_fault to identify
187 * the segment it has been passed. So if you make it
188 * "static" you'll need to fix rootnex_map_fault.
189 */
190 struct seg_ops segdev_ops = {
191     .dup = segdev_dup,
192     .unmap = segdev_unmap,

```

```

193     .free = segdev_free,
194     .fault = segdev_fault,
195     .faulta = segdev_faulta,
196     .setprot = segdev_setprot,
197     .checkprot = segdev_checkprot,
198     .kluster = (int (*)())segdev_badop,
199     .sync = segdev_sync,
200     .incore = segdev_incore,
201     .lockop = segdev_lockop,
202     .getprot = segdev_getprot,
203     .getoffset = segdev_getoffset,
204     .gettype = segdev_gettype,
205     .getvp = segdev_getvp,
206     .advise = segdev_advise,
207     .dump = segdev_dump,
208     .pagelock = segdev_pagelock,
209     .getmemid = segdev_getmemid,
210 };
211 #endif
212
213 unchanged portion omitted
214
215 /*
216 * segdev pages are not dumped, so we just return
217 */
218 /*ARGSUSED*/
219 static void
220 segdev_dump(struct seg *seg)
221 {}

222
223 /*
224 * ddi_segmap_setup: Used by drivers who wish specify mapping attributes
225 * for a segment. Called from a drivers segmap(9E)
226 * routine.
227 */
228 /*ARGSUSED*/
229 int
230 ddi_segmap_setup(dev_t dev, off_t offset, struct as *as, caddr_t *addrp,
231     off_t len, uint_t prot, uint_t maxprot, uint_t flags, cred_t *cred,
232     ddi_device_acc_attr_t *accattrp, uint_t rnumber)
233 {
234     struct segdev_crargs dev_a;
235     int (*mapfunc)(dev_t dev, off_t off, int prot);
236     uint_t hat_attr;
237     pfn_t pfn;
238     int error, i;

239     TRACE_0(TR_FAC_DEVMAP, TR_DEVMAP_SEGMAP_SETUP,
240         "ddi_segmap_setup:start");

241     if ((mapfunc = devopsp[getmajor(dev)]->devo_cb_ops->cb_mmap) == nodev)
242         return (ENODEV);

243     /*
244      * Character devices that support the d_mmap
245      * interface can only be mmap'ed shared.
246      */
247     if ((flags & MAP_TYPE) != MAP_SHARED)
248         return (EINVAL);

249     /*
250      * Check that this region is indeed mappable on this platform.
251      * Use the mapping function.
252      */
253     if (ddi_device_mapping_check(dev, accattrp, rnumber, &hat_attr) == -1)
254         return (ENXIO);

255     /*

```

```
2411     * Check to ensure that the entire range is
2412     * legal and we are not trying to map in
2413     * more than the device will let us.
2414     */
2415     for (i = 0; i < len; i += PAGE_SIZE) {
2416         if (i == 0) {
2417             /*
2418              * Save the pfn at offset here. This pfn will be
2419              * used later to get user address.
2420              */
2421             if ((pfn = (pfn_t)cdev_mmap(mapfunc, dev, offset,
2422                                     maxprot)) == PFN_INVALID)
2423                 return (ENXIO);
2424         } else {
2425             if (cdev_mmap(mapfunc, dev, offset + i, maxprot) ==
2426                 PFN_INVALID)
2427                 return (ENXIO);
2428         }
2429     }
2431     as_rangelock(as);
2432     /* Pick an address w/o worrying about any vac alignment constraints. */
2433     error = choose_addr(as, addrp, len, ptob(pfn), ADDR_NOVACALIGN, flags);
2434     if (error != 0) {
2435         as_rangeunlock(as);
2436         return (error);
2437     }
2439     dev_a.mapfunc = mapfunc;
2440     dev_a.dev = dev;
2441     dev_a.offset = (offset_t)offset;
2442     dev_a.type = flags & MAP_TYPE;
2443     dev_a.prot = (uchar_t)prot;
2444     dev_a.maxprot = (uchar_t)maxprot;
2445     dev_a.hat_attr = hat_attr;
2446     dev_a.hat_flags = 0;
2447     dev_a.devmap_data = NULL;
2449     error = as_map(as, *addrp, len, segdev_create, &dev_a);
2450     as_rangeunlock(as);
2451     return (error);
2453 }
_____unchanged_portion_omitted_____
```

```

*****
9423 Tue Nov 24 09:35:08 2015
new/usr/src/uts/common/vm/seg_kpm.c
6152 use NULL dump segop as a shorthand for no-op
*****
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, Version 1.0 only
6  * (the "License"). You may not use this file except in compliance
7  * with the License.
8  *
9  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
10 * or http://www.opensolaris.org/os/licensing.
11 * See the License for the specific language governing permissions
12 * and limitations under the License.
13 *
14 * When distributing Covered Code, include this CDDL HEADER in each
15 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
16 * If applicable, add the following below this CDDL HEADER, with the
17 * fields enclosed by brackets "[]" replaced with your own identifying
18 * information: Portions Copyright [yyyy] [name of copyright owner]
19 *
20 * CDDL HEADER END
21 */
22 /*
23 * Copyright 2006 Sun Microsystems, Inc. All rights reserved.
24 * Use is subject to license terms.
25 */

27 /*
28 * Kernel Physical Mapping (kpm) segment driver (segkpm).
29 *
30 * This driver delivers along with the hat_kpm* interfaces an alternative
31 * mechanism for kernel mappings within the 64-bit Solaris operating system,
32 * which allows the mapping of all physical memory into the kernel address
33 * space at once. This is feasible in 64 bit kernels, e.g. for Ultrasparc II
34 * and beyond processors, since the available VA range is much larger than
35 * possible physical memory. Momentarily all physical memory is supported,
36 * that is represented by the list of memory segments (memsegs).
37 *
38 * Segkpm mappings have also very low overhead and large pages are used
39 * (when possible) to minimize the TLB and TSB footprint. It is also
40 * extensible for other than Sparc architectures (e.g. AMD64). Main
41 * advantage is the avoidance of the TLB-shutdown X-calls, which are
42 * normally needed when a kernel (global) mapping has to be removed.
43 *
44 * First example of a kernel facility that uses the segkpm mapping scheme
45 * is seg_map, where it is used as an alternative to hat_memload().
46 * See also hat layer for more information about the hat_kpm* routines.
47 * The kpm facility can be turned off at boot time (e.g. /etc/system).
48 */

50 #include <sys/types.h>
51 #include <sys/param.h>
52 #include <sys/sysmacros.h>
53 #include <sys/system.h>
54 #include <sys/vnode.h>
55 #include <sys/cmn_err.h>
56 #include <sys/debug.h>
57 #include <sys/thread.h>
58 #include <sys/cpuvar.h>
59 #include <sys/bitmap.h>
60 #include <sys/atomic.h>
61 #include <sys/lgrp.h>

```

```

63 #include <vm/seg_kmem.h>
64 #include <vm/seg_kpm.h>
65 #include <vm/hat.h>
66 #include <vm/as.h>
67 #include <vm/seg.h>
68 #include <vm/page.h>

70 /*
71 * Global kpm controls.
72 * See also platform and mmu specific controls.
73 *
74 * kpm_enable -- global on/off switch for segkpm.
75 * . Set by default on 64bit platforms that have kpm support.
76 * . Will be disabled from platform layer if not supported.
77 * . Can be disabled via /etc/system.
78 *
79 * kpm_smallpages -- use only regular/system pagesize for kpm mappings.
80 * . Can be useful for critical debugging of kpm clients.
81 * . Set to zero by default for platforms that support kpm large pages.
82 * . The use of kpm large pages reduces the footprint of kpm meta data
83 * and has all the other advantages of using large pages (e.g TLB
84 * miss reduction).
85 * . Set by default for platforms that don't support kpm large pages or
86 * where large pages cannot be used for other reasons (e.g. there are
87 * only few full associative TLB entries available for large pages).
88 *
89 * segmap_kpm -- separate on/off switch for segmap using segkpm:
90 * . Set by default.
91 * . Will be disabled when kpm_enable is zero.
92 * . Will be disabled when MAXBSIZE != PAGESIZE.
93 * . Can be disabled via /etc/system.
94 *
95 */
96 int kpm_enable = 1;
97 int kpm_smallpages = 0;
98 int segmap_kpm = 1;

100 /*
101 * Private seg op routines.
102 */
103 faultcode_t segkpm_fault(struct hat *hat, struct seg *seg, caddr_t addr,
104                          size_t len, enum fault_type type, enum seg_rw rw);
105 static void segkpm_dump(struct seg *);
105 static void segkpm_badop(void);
106 static int segkpm_notsup(void);

108 #define SEGKPM_BADOP(t) (t(*)())segkpm_badop
109 #define SEGKPM_NOTSUP (int(*)())segkpm_notsup

111 static struct seg_ops segkpm_ops = {
112     .dup = SEGKPM_BADOP(int),
113     .unmap = SEGKPM_BADOP(int),
114     .free = SEGKPM_BADOP(void),
115     .fault = segkpm_fault,
116     .faulta = SEGKPM_BADOP(int),
117     .setprot = SEGKPM_BADOP(int),
118     .checkprot = SEGKPM_BADOP(int),
119     .kluster = SEGKPM_BADOP(int),
120     .swapout = SEGKPM_BADOP(size_t),
121     .sync = SEGKPM_BADOP(int),
122     .incore = SEGKPM_BADOP(size_t),
123     .lockop = SEGKPM_BADOP(int),
124     .getprot = SEGKPM_BADOP(int),
125     .getoffset = SEGKPM_BADOP(u_offset_t),
126     .gettype = SEGKPM_BADOP(int),

```

```
127     .getvp      = SEGKPM_BADOP(int),
128     .advise     = SEGKPM_BADOP(int),
130     .dump       = segkpm_dump,
129     .pagelock   = SEGKPM_NOTSUP,
130     .setpagesize = SEGKPM_BADOP(int),
131     .getmemid   = SEGKPM_BADOP(int),
132     .getpolicy  = SEGKPM_BADOP(lgrp_mem_policy_info_t *),
133 };
```

unchanged_portion_omitted

```
322 /*
323  * segkpm pages are not dumped, so we just return
324  */
325 /*ARGSUSED*/
326 static void
327 segkpm_dump(struct seg *seg)
328 {}
```

```

*****
83531 Tue Nov 24 09:35:08 2015
new/usr/src/uts/common/vm/seg_spt.c
6152 use NULL dump segop as a shorthand for no-op
*****
_____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);
141 static int segspt_shmgetmemid(struct seg *, caddr_t, memid_t *);
142 static lgrp_mem_policy_info_t *segspt_shmgetpolicy(struct seg *, caddr_t);

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

2853 /*
2854 * We need to wait for pending IO to complete to a DISM segment in order for
2855 * pages to get kicked out of the seg_pcache. 120 seconds should be more

```

```

2856 * than enough time to wait.
2857 */
2858 static clock_t spt_pcache_wait = 120;

2860 /*ARGSUSED*/
2861 static int
2862 segspt_shmadvise(struct seg *seg, caddr_t addr, size_t len, uint_t behav)
2863 {
2864     struct shm_data *shmd = (struct shm_data *)seg->s_data;
2865     struct spt_data *sptd = (struct spt_data *)shmd->shm_sptseg->s_data;
2866     struct anon_map *amp;
2867     pgcnt_t pg_idx;
2868     ushort_t gen;
2869     clock_t end_lbolt;
2870     int writer;
2871     page_t **ppa;

2873     ASSERT(seg->s_as && AS_LOCK_HELD(seg->s_as, &seg->s_as->a_lock));

2875     if (behav == MADV_FREE) {
2876         if ((sptd->spt_flags & SHM_PAGEABLE) == 0)
2877             return (0);

2879         amp = sptd->spt_amp;
2880         pg_idx = seg_page(seg, addr);

2882         mutex_enter(&sptd->spt_lock);
2883         if ((ppa = sptd->spt_ppa) == NULL) {
2884             mutex_exit(&sptd->spt_lock);
2885             ANON_LOCK_ENTER(&amp->a_rwlock, RW_READER);
2886             anon_disclaim(amp, pg_idx, len);
2887             ANON_LOCK_EXIT(&amp->a_rwlock);
2888             return (0);
2889         }

2891         sptd->spt_flags |= DISM_PPA_CHANGED;
2892         gen = sptd->spt_gen;

2894         mutex_exit(&sptd->spt_lock);

2896         /*
2897          * Purge all DISM cached pages
2898          */
2899         seg_ppurge_wiredpp(ppa);

2901         /*
2902          * Drop the AS_LOCK so that other threads can grab it
2903          * in the as_pageunlock path and hopefully get the segment
2904          * kicked out of the seg_pcache. We bump the shm_softlockcnt
2905          * to keep this segment resident.
2906          */
2907         writer = AS_WRITE_HELD(seg->s_as, &seg->s_as->a_lock);
2908         atomic_inc_ulong_t(&(shmd->shm_softlockcnt));
2909         AS_LOCK_EXIT(seg->s_as, &seg->s_as->a_lock);

2911         mutex_enter(&sptd->spt_lock);

2913         end_lbolt = ddi_get_lbolt() + (hz * spt_pcache_wait);

2915         /*
2916          * Try to wait for pages to get kicked out of the seg_pcache.
2917          */
2918         while (sptd->spt_gen == gen &&
2919             (sptd->spt_flags & DISM_PPA_CHANGED) &&
2920             ddi_get_lbolt() < end_lbolt) {
2921             if (!cv_timedwait_sig(&sptd->spt_cv,

```

```

2922         &sptd->spt_lock, end_lbolt)) {
2923             break;
2924         }
2925     }
2926
2927     mutex_exit(&sptd->spt_lock);
2928
2929     /* Regrab the AS_LOCK and release our hold on the segment */
2930     AS_LOCK_ENTER(seg->s_as, &seg->s_as->a_lock,
2931         writer ? RW_WRITER : RW_READER);
2932     atomic_dec_ulong((ulong_t *)&(shmd->shm_softlockcnt));
2933     if (shmd->shm_softlockcnt <= 0) {
2934         if (AS_ISUNMAPWAIT(seg->s_as)) {
2935             mutex_enter(&seg->s_as->a_contents);
2936             if (AS_ISUNMAPWAIT(seg->s_as)) {
2937                 AS_CLRUNMAPWAIT(seg->s_as);
2938                 cv_broadcast(&seg->s_as->a_cv);
2939             }
2940             mutex_exit(&seg->s_as->a_contents);
2941         }
2942     }
2943
2944     ANON_LOCK_ENTER(&amp->a_rwlock, RW_READER);
2945     anon_disclaim(amp, pg_idx, len);
2946     ANON_LOCK_EXIT(&amp->a_rwlock);
2947 } else if (lgrp_optimizations() && (behav == MADV_ACCESS_LWP ||
2948     behav == MADV_ACCESS_MANY || behav == MADV_ACCESS_DEFAULT)) {
2949     int         already_set;
2950     ulong_t     anon_index;
2951     lgrp_mem_policy_t  policy;
2952     caddr_t     shm_addr;
2953     size_t     share_size;
2954     size_t     size;
2955     struct seg  *sptseg = shmd->shm_sptseg;
2956     caddr_t     sptseg_addr;
2957
2958     /*
2959      * Align address and length to page size of underlying segment
2960      */
2961     share_size = page_get_pagesize(shmd->shm_sptseg->s_szc);
2962     shm_addr = (caddr_t)P2ALIGN((uintptr_t)(addr), share_size);
2963     size = P2ROUNDUP(((uintptr_t)((addr + len) - shm_addr)),
2964         share_size);
2965
2966     amp = shmd->shm_amp;
2967     anon_index = seg_page(seg, shm_addr);
2968
2969     /*
2970      * And now we may have to adjust size downward if we have
2971      * exceeded the realsize of the segment or initial anon
2972      * allocations.
2973      */
2974     sptseg_addr = sptseg->s_base + ptob(anon_index);
2975     if ((sptseg_addr + size) >
2976         (sptseg->s_base + sptd->spt_realsize))
2977         size = (sptseg->s_base + sptd->spt_realsize) -
2978             sptseg_addr;
2979
2980     /*
2981      * Set memory allocation policy for this segment
2982      */
2983     policy = lgrp_madv_to_policy(behav, len, MAP_SHARED);
2984     already_set = lgrp_shm_policy_set(policy, amp, anon_index,
2985         NULL, 0, len);
2986
2987     /*

```

```

2988         * If random memory allocation policy set already,
2989         * don't bother reapplying it.
2990         */
2991     if (already_set && !LGRP_MEM_POLICY_REAPPLICABLE(policy))
2992         return (0);
2993
2994     /*
2995      * Mark any existing pages in the given range for
2996      * migration, flushing the I/O page cache, and using
2997      * underlying segment to calculate anon index and get
2998      * anonmap and vnode pointer from
2999      */
3000     if (shmd->shm_softlockcnt > 0)
3001         segspt_purge(seg);
3002
3003     page_mark_migrate(seg, shm_addr, size, amp, 0, NULL, 0, 0);
3004 }
3005
3006     return (0);
3007 }
3008
3009 }
3010
3011 /*ARGSUSED*/
3012 void
3013 segspt_shmdump(struct seg *seg)
3014 {
3015     /* no-op for ISM segment */
3016 }
3017
3018 _____unchanged_portion_omitted_____

```

```

*****
54578 Tue Nov 24 09:35:09 2015
new/usr/src/uts/common/vm/vm_seg.c
6152 use NULL dump segop as a shorthand for no-op
*****
_____unchanged_portion_omitted_____

1964 void
1965 segop_dump(struct seg *seg)
1966 {
1967     if (seg->s_ops->dump == NULL)
1968         return;

1970 #endif /* ! codereview */
1971     seg->s_ops->dump(seg);
1972 }

1974 int
1975 segop_pagelock(struct seg *seg, caddr_t addr, size_t len, struct page ***page,
1976               enum lock_type type, enum seg_rw rw)
1977 {
1978     return (seg->s_ops->pagelock(seg, addr, len, page, type, rw));
1979 }

1981 int
1982 segop_setpagesize(struct seg *seg, caddr_t addr, size_t len, uint_t szc)
1983 {
1984     if (seg->s_ops->setpagesize == NULL)
1985         return (ENOTSUP);

1987     return (seg->s_ops->setpagesize(seg, addr, len, szc));
1988 }

1990 int
1991 segop_getmemid(struct seg *seg, caddr_t addr, memid_t *mp)
1992 {
1993     if (seg->s_ops->getmemid == NULL)
1994         return (ENODEV);

1996     return (seg->s_ops->getmemid(seg, addr, mp));
1997 }

1999 struct lgrp_mem_policy_info *
2000 segop_getpolicy(struct seg *seg, caddr_t addr)
2001 {
2002     if (seg->s_ops->getpolicy == NULL)
2003         return (NULL);

2005     return (seg->s_ops->getpolicy(seg, addr));
2006 }

2008 int
2009 segop_capable(struct seg *seg, segcapability_t cap)
2010 {
2011     if (seg->s_ops->capable == NULL)
2012         return (0);

2014     return (seg->s_ops->capable(seg, cap));
2015 }

2017 int
2018 segop_inherit(struct seg *seg, caddr_t addr, size_t len, uint_t op)
2019 {
2020     if (seg->s_ops->inherit == NULL)
2021         return (ENOTSUP);

```

```

2023     return (seg->s_ops->inherit(seg, addr, len, op));
2024 }

```

16457 Tue Nov 24 09:35:09 2015

new/usr/src/uts/i86xpv/vm/seg_mf.c

6152 use NULL dump segop as a shorthand for no-op

unchanged portion omitted

```
472 /*ARGSUSED*/
473 static void
474 segmf_dump(struct seg *seg)
475 {}
```

```
477 /*ARGSUSED*/
478 static int
479 segmf_pagelock(struct seg *seg, caddr_t addr, size_t len,
480               struct page ***ppp, enum lock_type type, enum seg_rw rw)
481 {
482     return (ENOTSUP);
483 }
```

unchanged portion omitted

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

unchanged portion omitted

```

*****
11601 Tue Nov 24 09:35:09 2015
new/usr/src/uts/sparc/v9/vm/seg_nf.c
6152 use NULL dump segop as a shorthand for no-op
*****
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,
                                struct page ***ppp, enum lock_type type, enum seg_rw rw);

80 struct seg_ops segnf_ops = {
81     .dup          = segnf_dup,
82     .unmap       = segnf_unmap,
83     .free        = segnf_free,
84     .fault       = (faultcode_t (*)(struct hat *, struct seg *, caddr_t,
85                                   size_t, enum fault_type, enum seg_rw)) segnf_nomap,
86     .faulta     = (faultcode_t (*)(struct seg *, caddr_t)) segnf_nomap,
87     .setprot     = segnf_setprot,
88     .checkprot  = segnf_checkprot,
89     .kluster    = (int (*)()) segnf_badop,
90     .sync       = (int (*)(struct seg *, caddr_t, size_t, int, uint_t))
91                 segnf_nop,
92     .incore     = (size_t (*)(struct seg *, caddr_t, size_t, char *))
93                 segnf_nop,
94     .lockop     = (int (*)(struct seg *, caddr_t, size_t, int, int,
95                           ulong_t *, size_t)) segnf_nop,
96     .getprot    = segnf_getprot,
97     .getoffset  = segnf_getoffset,
98     .gettype    = segnf_gettype,
99     .getvp     = segnf_getvp,
100    .advise     = (int (*)(struct seg *, caddr_t, size_t, uint_t))
101                 segnf_nop,
102    .dump       = segnf_dump,
103    .pagelock   = segnf_pagelock,
104 };
105
106 unchanged_portion_omitted
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000

```