

```

*****
44932 Fri May 8 18:04:17 2015
new/usr/src/uts/common/vm/seg_kmem.c
no need for bad-op segment op functions
The segment drivers have a number of bad-op functions that simply panic.
Keeping the function pointer NULL will accomplish the same thing in most
cases. In other cases, keeping the function pointer NULL will result in
proper error code being returned.
*****
_____unchanged_portion_omitted_____

431 static void
432 segkmem_badop()
433 {
434     panic("segkmem_badop");
435 }

437 #define SEGMEM_BADOP(t)      (t(*)())segkmem_badop

431 /*ARGSUSED*/
432 static faultcode_t
433 segkmem_fault(struct hat *hat, struct seg *seg, caddr_t addr, size_t size,
434               enum fault_type type, enum seg_rw rw)
435 {
436     pgcnt_t npages;
437     spgcnt_t pg;
438     page_t *pp;
439     struct vnode *vp = seg->s_data;

441     ASSERT(RW_READ_HELD(&seg->s_as->a_lock));

443     if (seg->s_as != &kas || size > seg->s_size ||
444         addr < seg->s_base || addr + size > seg->s_base + seg->s_size)
445         panic("segkmem_fault: bad args");

447     /*
448      * If it is one of segkp pages, call segkp_fault.
449      */
450     if (segkp_bitmap && seg == &kvseg &&
451         BT_TEST(segkp_bitmap, btop((uintptr_t)(addr - seg->s_base))))
452         return (segop_fault(hat, segkp, addr, size, type, rw));

454     if (rw != S_READ && rw != S_WRITE && rw != S_OTHER)
455         return (FC_NOSUPPORT);

457     npages = btopr(size);

459     switch (type) {
460     case F_SOFTLOCK: /* lock down already-loaded translations */
461         for (pg = 0; pg < npages; pg++) {
462             pp = page_lookup(vp, (u_offset_t)(uintptr_t)addr,
463                             SE_SHARED);
464             if (pp != NULL) {
465                 /*
466                  * Hmm, no page. Does a kernel mapping
467                  * exist for it?
468                  */
469                 if (!hat_probe(kas.a_hat, addr)) {
470                     addr -= PAGE_SIZE;
471                     while (--pg >= 0) {
472                         pp = page_find(vp, (u_offset_t)
473                                         (uintptr_t)addr);
474                         if (pp)
475                             page_unlock(pp);
476                         addr -= PAGE_SIZE;
477                     }

```

```

478         return (FC_NOMAP);
479     }
480     }
481     addr += PAGE_SIZE;
482 }
483 if (rw == S_OTHER)
484     hat_reserve(seg->s_as, addr, size);
485 return (0);
486 case F_SOFTUNLOCK:
487     while (npages-- > 0) {
488         pp = page_find(vp, (u_offset_t)(uintptr_t)addr);
489         if (pp)
490             page_unlock(pp);
491         addr += PAGE_SIZE;
492     }
493     return (0);
494 default:
495     return (FC_NOSUPPORT);
496 }
497 /*NOTREACHED*/
498 }
_____unchanged_portion_omitted_____

523 /*
524 * This is a dummy segkmem function overloaded to call segkp
525 * when segkp is under the heap.
526 */
527 /* ARGSUSED */
528 static int
529 segkmem_checkprot(struct seg *seg, caddr_t addr, size_t size, uint_t prot)
530 {
531     ASSERT(RW_LOCK_HELD(&seg->s_as->a_lock));

533     if (seg->s_as != &kas)
534         panic("segkmem_badop");
535     segkmem_badop();

536     /*
537      * If it is one of segkp pages, call into segkp.
538      */
539     if (segkp_bitmap && seg == &kvseg &&
540         BT_TEST(segkp_bitmap, btop((uintptr_t)(addr - seg->s_base))))
541         return (segop_checkprot(segkp, addr, size, prot));

543     panic("segkmem_badop");
544     segkmem_badop();
545     return (0);
546 }

547 /*
548 * This is a dummy segkmem function overloaded to call segkp
549 * when segkp is under the heap.
550 */
551 /* ARGSUSED */
552 static int
553 segkmem_kluster(struct seg *seg, caddr_t addr, ssize_t delta)
554 {
555     ASSERT(RW_LOCK_HELD(&seg->s_as->a_lock));

557     if (seg->s_as != &kas)
558         panic("segkmem_badop");
559     segkmem_badop();

560     /*
561      * If it is one of segkp pages, call into segkp.
562      */

```

```

563     if (segkp_bitmap && seg == &kvseg &&
564         BT_TEST(segkp_bitmap, btop((uintptr_t)(addr - seg->s_base))))
565         return (segop_kluster(segkp, addr, delta));

```

```

567     panic("segkmem badop");
575     segkmem_badop();
568     return (0);
569 }

```

unchanged_portion_omitted

```

728 /*
729 * This is a dummy segkmem function overloaded to call segkp
730 * when segkp is under the heap.
731 */
732 /* ARGSUSED */
733 static int
734 segkmem_getmemid(struct seg *seg, caddr_t addr, memid_t *memidp)
735 {
736     ASSERT(RW_LOCK_HELD(&seg->s_as->a_lock));

```

```

738     if (seg->s_as != &kas)
739         panic("segkmem badop");
747     segkmem_badop();

```

```

741     /*
742     * If it is one of segkp pages, call into segkp.
743     */
744     if (segkp_bitmap && seg == &kvseg &&
745         BT_TEST(segkp_bitmap, btop((uintptr_t)(addr - seg->s_base))))
746         return (segop_getmemid(segkp, addr, memidp));

```

```

748     panic("segkmem badop");
756     segkmem_badop();
749     return (0);
750 }

```

unchanged_portion_omitted

```

768 static struct seg_ops segkmem_ops = {
777     .dup           = SEGKMEM_BADOP(int),
778     .unmap        = SEGKMEM_BADOP(int),
779     .free         = SEGKMEM_BADOP(void),
769     .fault        = segkmem_fault,
781     .faulta       = SEGKMEM_BADOP(faultcode_t),
770     .setprot      = segkmem_setprot,
771     .checkprot    = segkmem_checkprot,
772     .kluster      = segkmem_kluster,
785     .sync         = SEGKMEM_BADOP(int),
786     .incore        = SEGKMEM_BADOP(size_t),
787     .lockop       = SEGKMEM_BADOP(int),
788     .getprot      = SEGKMEM_BADOP(int),
789     .getoffset    = SEGKMEM_BADOP(u_offset_t),
790     .gettype      = SEGKMEM_BADOP(int),
791     .getvp        = SEGKMEM_BADOP(int),
792     .advise       = SEGKMEM_BADOP(int),
773     .dump         = segkmem_dump,
774     .pagelock     = segkmem_pagelock,
795     .setpagesize  = SEGKMEM_BADOP(int),
775     .getmemid     = segkmem_getmemid,
776     .getpolicy    = segkmem_getpolicy,
777     .capable      = segkmem_capable,
778     .inherit      = seg_inherit_notsup,
779 };

```

unchanged_portion_omitted

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

1

```
*****
36313 Fri May 8 18:04:17 2015
new/usr/src/uts/common/vm/seg_kp.c
no need for bad-op segment op functions
The segment drivers have a number of bad-op functions that simply panic.
Keeping the function pointer NULL will accomplish the same thing in most
cases. In other cases, keeping the function pointer NULL will result in
proper error code being returned.
*****
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 (c) 1991, 2010, Oracle and/or its affiliates. All rights reserved.
23 */

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

28 /*
29  * Portions of this source code were derived from Berkeley 4.3 BSD
30  * under license from the Regents of the University of California.
31 */

33 /*
34  * segkp is a segment driver that administers the allocation and deallocation
35  * of pageable variable size chunks of kernel virtual address space. Each
36  * allocated resource is page-aligned.
37  *
38  * The user may specify whether the resource should be initialized to 0,
39  * include a redzone, or locked in memory.
40  */

42 #include <sys/types.h>
43 #include <sys/t_lock.h>
44 #include <sys/thread.h>
45 #include <sys/param.h>
46 #include <sys/errno.h>
47 #include <sys/sysmacros.h>
48 #include <sys/system.h>
49 #include <sys/buf.h>
50 #include <sys/mman.h>
51 #include <sys/vnode.h>
52 #include <sys/cmn_err.h>
53 #include <sys/swap.h>
54 #include <sys/tuneable.h>
55 #include <sys/kmem.h>
56 #include <sys/vmem.h>
57 #include <sys/cred.h>
```

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

2

```
58 #include <sys/dumphdr.h>
59 #include <sys/debug.h>
60 #include <sys/vtrace.h>
61 #include <sys/stack.h>
62 #include <sys/atomic.h>
63 #include <sys/archsystem.h>
64 #include <sys/lgrp.h>

66 #include <vm/as.h>
67 #include <vm/seg.h>
68 #include <vm/seg_kp.h>
69 #include <vm/seg_kmem.h>
70 #include <vm/anon.h>
71 #include <vm/page.h>
72 #include <vm/hat.h>
73 #include <sys/bitmap.h>

75 /*
76  * Private seg op routines
77  */
78 static void segkp_badop(void);
78 static void segkp_dump(struct seg *seg);
79 static int segkp_checkprot(struct seg *seg, caddr_t addr, size_t len,
80                          uint_t prot);
81 static int segkp_kluster(struct seg *seg, caddr_t addr, ssize_t delta);
82 static int segkp_pagelock(struct seg *seg, caddr_t addr, size_t len,
83                          struct page ***page, enum lock_type type,
84                          enum seg_rw rw);
85 static void segkp_insert(struct seg *seg, struct segkp_data *kpd);
86 static void segkp_delete(struct seg *seg, struct segkp_data *kpd);
87 static caddr_t segkp_get_internal(struct seg *seg, size_t len, uint_t flags,
88                                  struct segkp_data **tkpd, struct anon_map *amp);
89 static void segkp_release_internal(struct seg *seg,
90                                   struct segkp_data *kpd, size_t len);
91 static int segkp_unlock(struct hat *hat, struct seg *seg, caddr_t vaddr,
92                        size_t len, struct segkp_data *kpd, uint_t flags);
93 static int segkp_load(struct hat *hat, struct seg *seg, caddr_t vaddr,
94                      size_t len, struct segkp_data *kpd, uint_t flags);
95 static struct segkp_data *segkp_find(struct seg *seg, caddr_t vaddr);
96 static int segkp_getmemid(struct seg *seg, caddr_t addr, memid_t *memidp);
97 static lgrp_mem_policy_info_t *segkp_getpolicy(struct seg *seg,
98                                                caddr_t addr);
99 static int segkp_capable(struct seg *seg, segcapability_t capability);

101 /*
102  * Lock used to protect the hash table(s) and caches.
103  */
104 static kmutex_t segkp_lock;

106 /*
107  * The segkp caches
108  */
109 static struct segkp_cache segkp_cache[SEGKP_MAX_CACHE];

112 #define SEGKP_BADOP(t) (t(*)())segkp_badop

111 /*
112  * When there are fewer than red_minavail bytes left on the stack,
113  * segkp_map_red() will map in the redzone (if called). 5000 seems
114  * to work reasonably well...
115  */
116 long red_minavail = 5000;

118 /*
119  * will be set to 1 for 32 bit x86 systems only, in startup.c
120  */
```

```

121 int      segkp_fromheap = 0;
122 ulong_t  *segkp_bitmap;

124 /*
125 * If segkp_map_red() is called with the redzone already mapped and
126 * with less than RED_DEEP_THRESHOLD bytes available on the stack,
127 * then the stack situation has become quite serious; if much more stack
128 * is consumed, we have the potential of scrogging the next thread/LWP
129 * structure. To help debug the "can't happen" panics which may
130 * result from this condition, we record hrestime and the calling thread
131 * in red_deep_hires and red_deep_thread respectively.
132 */
133 #define RED_DEEP_THRESHOLD      2000

135 hrtime_t      red_deep_hires;
136 kthread_t     *red_deep_thread;

138 uint32_t      red_nmapped;
139 uint32_t      red_closest = UINT_MAX;
140 uint32_t      red_ndoubles;

142 pgcnt_t anon_segkp_pages_locked;      /* See vm/anon.h */
143 pgcnt_t anon_segkp_pages_resv;      /* anon reserved by seg_kp */

145 static struct seg_ops segkp_ops = {
149     .dup          = SEGKP_BADOP(int),
150     .unmap       = SEGKP_BADOP(int),
151     .free        = SEGKP_BADOP(void),
152     .fault       = segkp_fault,
153     .faulta     = SEGKP_BADOP(faultcode_t),
154     .setprot     = SEGKP_BADOP(int),
155     .checkprot  = segkp_checkprot,
156     .kluster    = segkp_kluster,
157     .sync       = SEGKP_BADOP(int),
158     .incore     = SEGKP_BADOP(size_t),
159     .lockop     = SEGKP_BADOP(int),
160     .getprot    = SEGKP_BADOP(int),
161     .getoffset  = SEGKP_BADOP(u_offset_t),
162     .gettype    = SEGKP_BADOP(int),
163     .getvp     = SEGKP_BADOP(int),
164     .advise     = SEGKP_BADOP(int),
165     .dump       = segkp_dump,
166     .pagelock   = segkp_pagelock,
167     .setpagesize = SEGKP_BADOP(int),
168     .getmemid   = segkp_getmemid,
169     .getpolicy  = segkp_getpolicy,
170     .capable    = segkp_capable,
171     .inherit    = seg_inherit_notsup,
172 };

175 static void
176 segkp_badop(void)
177 {
178     panic("segkp_badop");
179     /*NOTREACHED*/
180 }

182 static void segkpinit_mem_config(struct seg *);

184 static uint32_t segkp_indel;

186 /*
187 * Allocate the segment specific private data struct and fill it in
188 * with the per kp segment mutex, anon ptr. array and hash table.
189 */

```

```

166 int
167 segkp_create(struct seg *seg)
168 {
169     struct segkp_segdata *kpsd;
170     size_t np;

172     ASSERT(seg != NULL && seg->s_as == &kas);
173     ASSERT(RW_WRITE_HELD(&seg->s_as->a_lock));

175     if (seg->s_size & PAGEOFFSET) {
176         panic("Bad segkp size");
177         /*NOTREACHED*/
178     }

180     kpsd = kmem_zalloc(sizeof (struct segkp_segdata), KM_SLEEP);

182     /*
183     * Allocate the virtual memory for segkp and initialize it
184     */
185     if (segkp_fromheap) {
186         np = btop(kvseg.s_size);
187         segkp_bitmap = kmem_zalloc(BT_SIZEOFMAP(np), KM_SLEEP);
188         kpsd->kpsd_arena = vmem_create("segkp", NULL, 0, PAGESIZE,
189             vmem_alloc, vmem_free, heap_arena, 5 * PAGESIZE, VM_SLEEP);
190     } else {
191         segkp_bitmap = NULL;
192         np = btop(seg->s_size);
193         kpsd->kpsd_arena = vmem_create("segkp", seg->s_base,
194             seg->s_size, PAGESIZE, NULL, NULL, NULL, 5 * PAGESIZE,
195             VM_SLEEP);
196     }

198     kpsd->kpsd_anon = anon_create(np, ANON_SLEEP | ANON_ALLOC_FORCE);

200     kpsd->kpsd_hash = kmem_zalloc(SEGKP_HASHSZ * sizeof (struct segkp *),
201         KM_SLEEP);
202     seg->s_data = (void *)kpsd;
203     seg->s_ops = &segkp_ops;
204     segkpinit_mem_config(seg);
205     return (0);
206 }

```

unchanged portion omitted

```

*****
9619 Fri May 8 18:04:17 2015
new/usr/src/uts/common/vm/seg_kpm.c
no need for bad-op segment op functions
The segment drivers have a number of bad-op functions that simply panic.
Keeping the function pointer NULL will accomplish the same thing in most
cases. In other cases, keeping the function pointer NULL will result in
proper error code being returned.
*****
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 *);
106 static int segkpm_pagelock(struct seg *seg, caddr_t addr, size_t len,
107                            struct page ***page, enum lock_type type,
108                            enum seg_rw rw);
109 static void segkpm_badop(void);
110 static int segkpm_notsup(void);
111 static int segkpm_capable(struct seg *, segcapability_t);

112 #define SEGKPM_BADOP(t) (t(*)())segkpm_badop
113 #define SEGKPM_NOTSUP (int(*)())segkpm_notsup

114 static struct seg_ops segkpm_ops = {
115     .dup = SEGKPM_BADOP(int),
116     .unmap = SEGKPM_BADOP(int),
117     .free = SEGKPM_BADOP(void),
118     .fault = segkpm_fault,
119     .faulta = SEGKPM_BADOP(int),
120     .setprot = SEGKPM_BADOP(int),
121     .checkprot = SEGKPM_BADOP(int),

```

```

121     .kluster      = SEGKPM_BADOP(int),
122     .sync         = SEGKPM_BADOP(int),
123     .incore       = SEGKPM_BADOP(size_t),
124     .lockop       = SEGKPM_BADOP(int),
125     .getprot      = SEGKPM_BADOP(int),
126     .getoffset    = SEGKPM_BADOP(u_offset_t),
127     .gettype      = SEGKPM_BADOP(int),
128     .getvp        = SEGKPM_BADOP(int),
129     .advise       = SEGKPM_BADOP(int),
130     .dump         = segkpm_dump,
131     .pagelock     = segkpm_pagelock,
132     .setpagesize  = SEGKPM_BADOP(int),
133     .getmemid     = SEGKPM_BADOP(int),
134     .getpolicy    = SEGKPM_BADOP(lgrp_mem_policy_info_t *),
135     .capable      = segkpm_capable,
136     .inherit      = seg_inherit_notsup,
137 // #ifndef SEGKPM_SUPPORT
138 #if 0
139 #error FIXME: define nop
140 .dup             = nop,
141 .unmap          = nop,
142 .free           = nop,
143 .faulta        = nop,
144 .setprot       = nop,
145 .checkprot     = nop,
146 .kluster       = nop,
147 .sync         = nop,
148 .incore        = nop,
149 .lockop        = nop,
150 .getprot       = nop,
151 .getoffset     = nop,
152 .gettype       = nop,
153 .getvp         = nop,
154 .advise        = nop,
155 .setpagesize   = nop,
156 .getmemid     = nop,
157 .getpolicy     = nop,
158 #endif
159 #endif /* ! codereview */
160 };

161 /*
162 * kpm_pgsz and kpm_pgshft are set by platform layer.
163 */
164 size_t      kpm_pgsz;      /* kpm page size */
165 uint_t      kpm_pgshft;   /* kpm page shift */
166 u_offset_t  kpm_pgoff;    /* kpm page offset mask */
167 uint_t      kpm_p2pshft; /* kpm page to page shift */
168 pgcnt_t     kpm_pnpgs;   /* how many pages per kpm page */

```

```

152 #ifdef SEGKPM_SUPPORT

```

```

154 int
155 segkpm_create(struct seg *seg, void *argsp)
156 {
157     struct segkpm_data *skd;
158     struct segkpm_crargs *b = (struct segkpm_crargs *)argsp;
159     ushort_t *p;
160     int i, j;

161     ASSERT(seg->s_as && RW_WRITE_HELD(&seg->s_as->a_lock));
162     ASSERT(btokpmp(seg->s_size) >= 1 &&
163            kmpmpageoff((uintptr_t)seg->s_base) == 0 &&
164            kmpmpageoff((uintptr_t)seg->s_base + seg->s_size) == 0);

```

```

167     skd = kmem_zalloc(sizeof (struct segkpm_data), KM_SLEEP);
168
169     seg->s_data = (void *)skd;
170     seg->s_ops = &segkpm_ops;
171     skd->skd_prot = b->prot;

172 /*
173  * (1) Segkpm virtual addresses are based on physical addresses.
174  * From this and in opposite to other segment drivers it is
175  * often required to allocate a page first to be able to
176  * calculate the final segkpm virtual address.
177  * (2) Page allocation is done by calling page_create_va(),
178  * one important input argument is a virtual address (also
179  * expressed by the "va" in the function name). This function
180  * is highly optimized to select the right page for an optimal
181  * processor and platform support (e.g. virtual addressed
182  * caches (VAC), physical addressed caches, NUMA).
183  *
184  * Because of (1) the approach is to generate a faked virtual
185  * address for calling page_create_va(). In order to exploit
186  * the abilities of (2), especially to utilize the cache
187  * hierarchy (3) and to avoid VAC alias conflicts (4) the
188  * selection has to be done carefully. For each virtual color
189  * a separate counter is provided (4). The count values are
190  * used for the utilization of all cache lines (3) and are
191  * corresponding to the cache bins.
192  */
193     skd->skd_nvcolors = b->nvcolors;

194     p = skd->skd_va_select =
195         kmem_zalloc(NCPU * b->nvcolors * sizeof (ushort_t), KM_SLEEP);

196     for (i = 0; i < NCPU; i++)
197         for (j = 0; j < b->nvcolors; j++, p++)
198             *p = j;

199     return (0);
200 }

201 /*
202 * This routine is called via a machine specific fault handling
203 * routine.
204 */
205 /* ARGSUSED */
206 faultcode_t
207 segkpm_fault(struct hat *hat, struct seg *seg, caddr_t addr, size_t len,
208              enum fault_type type, enum seg_rw rw)
209 {
210     ASSERT(seg->s_as && AS_LOCK_HELD(seg->s_as, &seg->s_as->a_lock));

211     switch (type) {
212     case F_INVAL:
213         return (hat_kpm_fault(hat, addr));
214     case F_SOFTLOCK:
215     case F_SOFTUNLOCK:
216         return (0);
217     default:
218         return (FC_NOSUPPORT);
219     }
220     /*NOTREACHED*/
221 }

222 #define addr_to_vcolor(addr, vcolors) \
223     ((int)(((uintptr_t)(addr) & ((vcolors << PAGESHIFT) - 1)) >> PAGESHIFT))

```

```

232 /*
233  * Create a virtual address that can be used for invocations of
234  * page_create_va. Goal is to utilize the cache hierarchy (round
235  * robin bins) and to select the right color for virtual indexed
236  * caches. It isn't exact since we also increment the bin counter
237  * when the caller uses VOP_GETPAGE and gets a hit in the page
238  * cache, but we keep the bins turning for cache distribution
239  * (see also segkpm_create block comment).
240  */
241 caddr_t
242 segkpm_create_va(u_offset_t off)
243 {
244     int vcolor;
245     ushort_t *p;
246     struct segkpm_data *skd = (struct segkpm_data *)segkpm->s_data;
247     int nvcolors = skd->skd_nvcolors;
248     caddr_t va;
249
250     vcolor = (nvcolors > 1) ? addr_to_vcolor(off, nvcolors) : 0;
251     p = &skd->skd_va_select[(CPU->cpu_id * nvcolors) + vcolor];
252     va = (caddr_t)ptob(*p);
253
254     atomic_add_16(p, nvcolors);
255
256     return (va);
257 }
258
259 /*
260  * Unload mapping if the instance has an active kpm mapping.
261  */
262 void
263 segkpm_mapout_validkpm(struct kpme *kpme)
264 {
265     caddr_t vaddr;
266     page_t *pp;
267
268     retry:
269     if ((pp = kpme->kpe_page) == NULL) {
270         return;
271     }
272
273     if (page_lock(pp, SE_SHARED, (kmutex_t *)NULL, P_RECLAIM) == 0)
274         goto retry;
275
276     /*
277      * Check if segkpm mapping is not unloaded in the meantime
278      */
279     if (kpme->kpe_page == NULL) {
280         page_unlock(pp);
281         return;
282     }
283
284     vaddr = hat_kpm_page2va(pp, 1);
285     hat_kpm_mapout(pp, kpme, vaddr);
286     page_unlock(pp);
287 }
288
289 static void
290 segkpm_badop()
291 {
292     panic("segkpm_badop");
293 }
294
295 #else /* SEGKPM_SUPPORT */
296
297 /* segkpm stubs */

```

```

293 /*ARGSUSED*/
294 int segkpm_create(struct seg *seg, void *argsp)
295 {
296     return (0);
297 }
298 int segkpm_create(struct seg *seg, void *argsp) { return (0); }
299
300 /* ARGSUSED */
301 faultcode_t
302 segkpm_fault(struct hat *hat, struct seg *seg, caddr_t addr, size_t len,
303              enum fault_type type, enum seg_rw rw)
304 {
305     return (0);
306     return ((faultcode_t)0);
307 }
308
309 /* ARGSUSED */
310 caddr_t segkpm_create_va(u_offset_t off)
311 {
312     return (NULL);
313 }
314 caddr_t segkpm_create_va(u_offset_t off) { return (NULL); }
315
316 /* ARGSUSED */
317 void segkpm_mapout_validkpm(struct kpme *kpme)
318 {
319 }
320 void segkpm_mapout_validkpm(struct kpme *kpme) {}
321
322 static void
323 segkpm_badop() {}
324
325 #endif /* SEGKPM_SUPPORT */
326
327 /* ARGSUSED */
328 #endif /* ! codereview */
329 static int
330 segkpm_pagelock(struct seg *seg, caddr_t addr, size_t len,
331                struct page ***page, enum lock_type type, enum seg_rw rw)
332 segkpm_notsup()
333 {
334     return (ENOTSUP);
335 }
336
337 /*
338  * segkpm pages are not dumped, so we just return
339  */
340 /*ARGSUSED*/
341 static void
342 segkpm_dump(struct seg *seg)
343 {
344 }
345
346 /*
347  * We claim to have no special capabilities.
348  */
349 /*ARGSUSED*/
350 static int
351 segkpm_capable(struct seg *seg, segcapability_t capability)
352 {
353     return (0);
354 }
355
356 unchanged_portion_omitted

```

```

*****
57729 Fri May 8 18:04:17 2015
new/usr/src/uts/common/vm/seg_map.c
no need for bad-op segment op functions
The segment drivers have a number of bad-op functions that simply panic.
Keeping the function pointer NULL will accomplish the same thing in most
cases. In other cases, keeping the function pointer NULL will result in
proper error code being returned.
*****
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 2009 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 - generic vnode mapping segment.
36  *
37  * The segmap driver is used only by the kernel to get faster (than seg_vn)
38  * mappings [lower routine overhead; more persistent cache] to random
39  * vnode/offsets. Note than the kernel may (and does) use seg_vn as well.
40 */

42 #include <sys/types.h>
43 #include <sys/t_lock.h>
44 #include <sys/param.h>
45 #include <sys/sysmacros.h>
46 #include <sys/buf.h>
47 #include <sys/system.h>
48 #include <sys/vnode.h>
49 #include <sys/mman.h>
50 #include <sys/errno.h>
51 #include <sys/cred.h>
52 #include <sys/kmem.h>
53 #include <sys/vtrace.h>
54 #include <sys/cmn_err.h>
55 #include <sys/bug.h>
56 #include <sys/thread.h>
57 #include <sys/dumphdr.h>

```

```

58 #include <sys/bitmap.h>
59 #include <sys/lgrp.h>

61 #include <vm/seg_kmem.h>
62 #include <vm/hat.h>
63 #include <vm/as.h>
64 #include <vm/seg.h>
65 #include <vm/seg_kpm.h>
66 #include <vm/seg_map.h>
67 #include <vm/page.h>
68 #include <vm/pvn.h>
69 #include <vm/rm.h>

71 /*
72  * Private seg op routines.
73 */
74 static void      segmap_free(struct seg *seg);
75 faultcode_t segmap_fault(struct hat *hat, struct seg *seg, caddr_t addr,
76                          size_t len, enum fault_type type, enum seg_rw rw);
77 static faultcode_t segmap_faulta(struct seg *seg, caddr_t addr);
78 static int      segmap_checkprot(struct seg *seg, caddr_t addr, size_t len,
79                                  uint_t prot);
80 static int      segmap_kluster(struct seg *seg, caddr_t addr, ssize_t);
81 static int      segmap_getprot(struct seg *seg, caddr_t addr, size_t len,
82                                 uint_t *protv);
83 static u_offset_t segmap_getoffset(struct seg *seg, caddr_t addr);
84 static int      segmap_gettype(struct seg *seg, caddr_t addr);
85 static int      segmap_getvp(struct seg *seg, caddr_t addr, struct vnode **vpp);
86 static void      segmap_dump(struct seg *seg);
87 static int      segmap_pagelock(struct seg *seg, caddr_t addr, size_t len,
88                                  struct page **ppp, enum lock_type type,
89                                  enum seg_rw rw);
90 static void      segmap_badop(void);
90 static int      segmap_getmemid(struct seg *seg, caddr_t addr, memid_t *memidp);
91 static lgrp_mem_policy_info_t *segmap_getpolicy(struct seg *seg,
92                                                  caddr_t addr);
93 static int      segmap_capable(struct seg *seg, segcapability_t capability);

95 /* segkpm support */
96 static caddr_t segmap_pagecreate_kpm(struct seg *, vnode_t *, u_offset_t,
97                                       struct smap *, enum seg_rw);
98 struct smap      *get_smap_kpm(caddr_t, page_t **);

101 #define SEGMAP_BADOP(t) (t(*)())segmap_badop

100 static struct seg_ops segmap_ops = {
104     .dup          = SEGMAP_BADOP(int),
105     .unmap        = SEGMAP_BADOP(int),
101     .free         = segmap_free,
102     .fault        = segmap_fault,
103     .faulta       = segmap_faulta,
109     .setprot      = SEGMAP_BADOP(int),
104     .checkprot    = segmap_checkprot,
105     .kluster      = segmap_kluster,
112     .sync         = SEGMAP_BADOP(int),
113     .incore       = SEGMAP_BADOP(size_t),
114     .lockop       = SEGMAP_BADOP(int),
106     .getprot      = segmap_getprot,
107     .getoffset    = segmap_getoffset,
108     .gettype      = segmap_gettype,
109     .getvp        = segmap_getvp,
119     .advise       = SEGMAP_BADOP(int),
110     .dump         = segmap_dump,
111     .pagelock     = segmap_pagelock,
122     .setpagesize  = SEGMAP_BADOP(int),
112     .getmemid     = segmap_getmemid,

```



```
113     .getpolicy      = segmap_getpolicy,  
114     .capable        = segmap_capable,  
115     .inherit        = seg_inherit_notsup,  
116 };
```

_____ unchanged_portion_omitted _____

```
886 /*  
887 * Check to see if it makes sense to do kluster/read ahead to  
888 * addr + delta relative to the mapping at addr. We assume here  
889 * that delta is a signed PAGE_SIZE'd multiple (which can be negative).  
890 *  
891 * For segmap we always "approve" of this action from our standpoint.  
892 */  
893 /*ARGSUSED*/  
894 static int  
895 segmap_kluster(struct seg *seg, caddr_t addr, ssize_t delta)  
896 {  
897     return (0);  
898 }  
899 }
```

```
911 static void  
912 segmap_badop()  
913 {  
914     panic("segmap_badop");  
915     /*NOTREACHED*/  
916 }
```

_____ unchanged_portion_omitted _____

```

*****
82793 Fri May 8 18:04:18 2015
new/usr/src/uts/common/vm/seg_spt.c
no need for bad-op segment op functions
The segment drivers have a number of bad-op functions that simply panic.
Keeping the function pointer NULL will accomplish the same thing in most
cases. In other cases, keeping the function pointer NULL will result in
proper error code being returned.
*****
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 (c) 1993, 2010, Oracle and/or its affiliates. All rights reserved.
23 */

25 #include <sys/param.h>
26 #include <sys/user.h>
27 #include <sys/mman.h>
28 #include <sys/kmem.h>
29 #include <sys/sysmacros.h>
30 #include <sys/cmn_err.h>
31 #include <sys/system.h>
32 #include <sys/tunable.h>
33 #include <vm/hat.h>
34 #include <vm/seg.h>
35 #include <vm/as.h>
36 #include <vm/anon.h>
37 #include <vm/page.h>
38 #include <sys/buf.h>
39 #include <sys/swap.h>
40 #include <sys/atomic.h>
41 #include <vm/seg_spt.h>
42 #include <sys/debug.h>
43 #include <sys/vtrace.h>
44 #include <sys/shm.h>
45 #include <sys/shm_impl.h>
46 #include <sys/lgrp.h>
47 #include <sys/vmsystem.h>
48 #include <sys/policy.h>
49 #include <sys/project.h>
50 #include <sys/tnf_probe.h>
51 #include <sys/zone.h>

53 #define SEGSPtADDR      (caddr_t)0x0

55 /*
56  * # pages used for spt
57 */

```

```

58 size_t  spt_used;

60 /*
61  * segspt_minfree is the memory left for system after ISM
62  * locked its pages; it is set up to 5% of availmem in
63  * sptcreate when ISM is created. ISM should not use more
64  * than ~90% of availmem; if it does, then the performance
65  * of the system may decrease. Machines with large memories may
66  * be able to use up more memory for ISM so we set the default
67  * segspt_minfree to 5% (which gives ISM max 95% of availmem.
68  * If somebody wants even more memory for ISM (risking hanging
69  * the system) they can patch the segspt_minfree to smaller number.
70 */
71 pgcnt_t segspt_minfree = 0;

73 static int segspt_create(struct seg *seg, caddr_t argsp);
74 static int segspt_unmap(struct seg *seg, caddr_t raddr, size_t ssize);
75 static void segspt_free(struct seg *seg);
76 static void segspt_free_pages(struct seg *seg, caddr_t addr, size_t len);
77 static lgrp_mem_policy_info_t *segspt_getpolicy(struct seg *seg, caddr_t addr);

79 static void
80 segspt_badop()
81 {
82     panic("segspt_badop called");
83     /*NOTREACHED*/
84 }

86 #define SEGSPt_BADOP(t) (t(*)())segspt_badop

79 struct seg_ops segspt_ops = {
89     .dup          = SEGSPt_BADOP(int),
90     .unmap        = segspt_unmap,
91     .free         = segspt_free,
92     .fault        = SEGSPt_BADOP(int),
93     .faulta       = SEGSPt_BADOP(faultcode_t),
94     .setprot      = SEGSPt_BADOP(int),
95     .checkprot    = SEGSPt_BADOP(int),
96     .kluster      = SEGSPt_BADOP(int),
97     .sync         = SEGSPt_BADOP(int),
98     .incore       = SEGSPt_BADOP(size_t),
99     .lockop       = SEGSPt_BADOP(int),
100    .getprot       = SEGSPt_BADOP(int),
101    .getoffset     = SEGSPt_BADOP(u_offset_t),
102    .gettype       = SEGSPt_BADOP(int),
103    .getvp         = SEGSPt_BADOP(int),
104    .advise        = SEGSPt_BADOP(int),
105    .dump          = SEGSPt_BADOP(void),
106    .pagelock      = SEGSPt_BADOP(int),
107    .setpagesize   = SEGSPt_BADOP(int),
108    .getmemid      = SEGSPt_BADOP(int),
82    .getpolicy     = segspt_getpolicy,
110    .capable       = SEGSPt_BADOP(int),
83    .inherit       = seg_inherit_notsup,
84 };
    unchanged_portion_omitted

```

```

*****
12330 Fri May 8 18:04:18 2015
new/usr/src/uts/sparc/v9/vm/seg_nf.c
no need for bad-op segment op functions
The segment drivers have a number of bad-op functions that simply panic.
Keeping the function pointer NULL will accomplish the same thing in most
cases. In other cases, keeping the function pointer NULL will result in
proper error code being returned.
*****
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);
69 static int     segnf_nop(void);
70 static int     segnf_getprot(struct seg *seg, caddr_t addr,
71                             size_t len, uint_t *protv);
72 static u_offset_t segnf_getoffset(struct seg *seg, caddr_t addr);
73 static int     segnf_gettype(struct seg *seg, caddr_t addr);
74 static int     segnf_getvp(struct seg *seg, caddr_t addr, struct vnode **vpp);
75 static void    segnf_dump(struct seg *seg);
76 static int     segnf_pagelock(struct seg *seg, caddr_t addr, size_t len,
77                              struct page ***ppp, enum lock_type type, enum seg_rw rw);
78 static int     segnf_setpagesize(struct seg *seg, caddr_t addr, size_t len,
79                                  uint_t szc);
80 static int     segnf_getmemid(struct seg *seg, caddr_t addr, memid_t *memidp);
81 static lgrp_mem_policy_info_t *segnf_getpolicy(struct seg *seg,
82                                                caddr_t addr);

85 struct seg_ops segnf_ops = {
86     .dup           = segnf_dup,
87     .unmap        = segnf_unmap,
88     .free         = segnf_free,
89     .fault        = (faultcode_t (*)(struct hat *, struct seg *, caddr_t,
90                                     size_t, enum fault_type, enum seg_rw)) segnf_nomap,
91     .faulta       = (faultcode_t (*)(struct seg *, caddr_t)) segnf_nomap,
92     .setprot      = segnf_setprot,
93     .checkprot    = segnf_checkprot,
94     .kluster      = (int (*)()) segnf_badop,
95     .sync         = (int (*)(struct seg *, caddr_t, size_t, int, uint_t))
96                     segnf_nop,
97     .incore       = (size_t (*)(struct seg *, caddr_t, size_t, char *))
98                     segnf_nop,
99     .lockop       = (int (*)(struct seg *, caddr_t, size_t, int, int,
100                             ulong_t *, size_t)) segnf_nop,
101     .getprot      = segnf_getprot,
102     .getoffset    = segnf_getoffset,
103     .gettype      = segnf_gettype,
104     .getvp        = segnf_getvp,
105     .advise       = (int (*)(struct seg *, caddr_t, size_t, uint_t))
106                     segnf_nop,
107     .dump         = segnf_dump,
108     .pagelock     = segnf_pagelock,
109     .setpagesize  = segnf_setpagesize,
110     .getmemid     = segnf_getmemid,
111     .getpolicy    = segnf_getpolicy,
112 };
113 #unchanged_portion_omitted

391 /* ARGSUSED */
392 static int
393 segnf_checkprot(struct seg *seg, caddr_t addr, size_t len, uint_t prot)
394 {
395     uint_t sprot;
396     ASSERT(seg->s_as && AS_LOCK_HELD(seg->s_as, &seg->s_as->a_lock));

398     sprot = seg->s_as == &kas ? PROT_READ : PROT_READ|PROT_USER;

```

```
399         return ((prot & sprot) == prot ? 0 : EACCES);  
402     }
```

```
404 static void  
405 segnf_badop(void)  
406 {  
407     panic("segnf_badop");  
408     /*NOTREACHED*/  
400 }
```

_____unchanged_portion_omitted_____