

```

*****
9266 Tue Nov 24 09:35:11 2015
new/usr/src/uts/common/vm/seg_kpm.c
6153 use NULL pagelock 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, 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_badop(void);
106 static int segkpm_notsup(void);

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

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

```

```
126     .advise      = SEGKPM_BADOP(int),
129     .pagelock   = SEGKPM_NOTSUP,
127     .setpagesize = SEGKPM_BADOP(int),
128     .getmemid   = SEGKPM_BADOP(int),
129     .getpolicy  = SEGKPM_BADOP(lgrp_mem_policy_info_t *),
130 };
```

unchanged_portion_omitted

```
300 /* ARGSUSED */
301 caddr_t segkpm_create_va(u_offset_t off) { return (NULL); }
```

```
303 /* ARGSUSED */
304 void segkpm_mapout_validkpme(struct kpme *kpme) {}
```

```
306 static void
307 segkpm_badop() {}
```

```
309 #endif /* SEGKPM_SUPPORT */
```

```
314 static int
315 segkpm_notsup()
316 {
317     return (ENOTSUP);
318 }
```

54634 Tue Nov 24 09:35:11 2015

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

6153 use NULL pagelock segop as a shorthand for ENOTSUP

_____unchanged_portion_omitted_____

```
1973 int
1974 segop_pagelock(struct seg *seg, caddr_t addr, size_t len, struct page ***page,
1975               enum lock_type type, enum seg_rw rw)
1976 {
1977     if (seg->s_ops->pagelock == NULL)
1978         return (ENOTSUP);
1980 #endif /* ! codereview */
1981     return (seg->s_ops->pagelock(seg, addr, len, page, type, rw));
1982 }
1984 int
1985 segop_setpagesize(struct seg *seg, caddr_t addr, size_t len, uint_t szc)
1986 {
1987     if (seg->s_ops->setpagesize == NULL)
1988         return (ENOTSUP);
1990     return (seg->s_ops->setpagesize(seg, addr, len, szc));
1991 }
1993 int
1994 segop_getmemid(struct seg *seg, caddr_t addr, memid_t *mp)
1995 {
1996     if (seg->s_ops->getmemid == NULL)
1997         return (ENODEV);
1999     return (seg->s_ops->getmemid(seg, addr, mp));
2000 }
2002 struct lgrp_mem_policy_info *
2003 segop_getpolicy(struct seg *seg, caddr_t addr)
2004 {
2005     if (seg->s_ops->getpolicy == NULL)
2006         return (NULL);
2008     return (seg->s_ops->getpolicy(seg, addr));
2009 }
2011 int
2012 segop_capable(struct seg *seg, segcapability_t cap)
2013 {
2014     if (seg->s_ops->capable == NULL)
2015         return (0);
2017     return (seg->s_ops->capable(seg, cap));
2018 }
2020 int
2021 segop_inherit(struct seg *seg, caddr_t addr, size_t len, uint_t op)
2022 {
2023     if (seg->s_ops->inherit == NULL)
2024         return (ENOTSUP);
2026     return (seg->s_ops->inherit(seg, addr, len, op));
2027 }
```