

```

*****
18577 Thu Feb 12 20:32:16 2015
new/usr/src/uts/armv6/loader/fakeloader.c
loader: allow LMB device maps
There's no reason we shouldn't allow LMB PTEs for use on device memory.
loader: map as much as possible using LMB pages
Chances are that we never actually executed this bit of code since all the
maps we ever deal with are either very short or much larger than LMB.
unix: enable caches in locore
The loader should really be as simple as possible to be as small as
possible. It should configure the machine so that unix can make certain
assumptions but it should leave more complex initialization to unix.
*****
1 /*
2  * This file and its contents are supplied under the terms of the
3  * Common Development and Distribution License ("CDDL"), version 1.0.
4  * You may only use this file in accordance with the terms of version
5  * 1.0 of the CDDL.
6  *
7  * A full copy of the text of the CDDL should have accompanied this
8  * source. A copy of the CDDL is also available via the Internet at
9  * http://www.illumos.org/license/CDDL.
10 */

12 /*
13  * Copyright (c) 2014 Joyent, Inc. All rights reserved.
14  * Copyright (c) 2015 Josef 'Jeff' Sipek <jeffpc@josefsipek.net>
15  */

17 #include "fakeloader.h"

19 #include <sys/types.h>
20 #include <sys/param.h>
21 #include <sys/elf.h>
22 #include <sys/atag.h>
23 #include <sys/sysmacros.h>
24 #include <sys/machparam.h>

26 #include <vm/pte.h>

28 /*
29  * This is the stock ARM fake uniboot loader.
30  *
31  * Here's what we have to do:
32  *   o Read the atag header and find the combined archive header
33  *   o Determine the set of mappings we need to add for the following:
34  *       - unix
35  *       - boot_archive
36  *       - atags
37  *   o Enable unaligned access
38  *   o Enable virtual memory
39  *   o Enable the caches + virtual memory
40  *
41  * There are several important constraints that we have here:
42  *   o We cannot use any .data! Several loaders that come before us are broken
43  *     and only provide us with the ability to map our .text and potentially our
44  *     .bss. We should strive to avoid even that if we can.
45  */

47 #ifdef DEBUG
48 #define FAKELOAD_DPRINTF(x)    fakeload_puts(x)
49 #else
50 #define FAKELOAD_DPRINTF(x)
51 #endif /* DEBUG */

```

```

53 /*
54  * XXX ASSUMES WE HAVE Free memory following the boot archive
55  */
56 static uintptr_t freemem;
57 static uintptr_t pt_arena;
58 static uintptr_t pt_arena_max;
59 static uint32_t *pt_addr;
60 static int nl2pages;

62 /* Simple copy routines */
63 void
64 bcopy(const void *s, void *d, size_t n)
65 {
66     const char *src = s;
67     char *dest = d;

69     if (n == 0 || s == d)
70         return;

72     if (dest < src && dest + n < src) {
73         /* dest overlaps with the start of src, copy forward */
74         for (; n > 0; n--, src++, dest++)
75             *dest = *src;
76     } else {
77         /* src overlaps with start of dest or no overlap, copy rev */
78         src += n - 1;
79         dest += n - 1;
80         for (; n > 0; n--, src--, dest--)
81             *dest = *src;
82     }
83 }

unchanged_portion_omitted

152 static void
153 fakeload_map_lmb(uintptr_t pa, uintptr_t va, int prot)
154 {
155     int entry;
156     armpte_t *pte;
157     arm_lls_t *lle;

159     entry = ARMPT_VADDR_TO_L1E(va);
160     pte = &pt_addr[entry];
161     if (ARMPT_L1E_ISVALID(*pte))
162         fakeload_panic("armboot_mmu: asked to map a mapped region!\n");
163     lle = (arm_lls_t *)pte;
164     *pte = 0;
165     lle->al_type = ARMPT_L1_TYPE_SECT;

167     if (prot & PF_DEVICE) {
168         lle->al_bbit = 1;
169         lle->al_cbit = 0;
170         lle->al_tex = 0;
171         lle->al_sbit = 1;
172     } else {
173         /* Assume it's not device memory */
174         lle->al_bbit = 1;
175         lle->al_cbit = 1;
176         lle->al_tex = 1;
177         lle->al_sbit = 1;
178     }
179 #endif /* ! codereview */

180     if (!(prot & PF_X))
181         lle->al_xn = 1;
182     lle->al_domain = 0;

```

```

184     if (prot & PF_W) {
185         lle->al_ap2 = 1;
186         lle->al_ap = 1;
187     } else {
188         lle->al_ap2 = 0;
189         lle->al_ap = 1;
190     }
191     lle->al_ngbit = 0;
192     lle->al_issuper = 0;
193     lle->al_addr = ARMPT_PADDR_TO_L1SECT(pa);
194 }

196 /*
197  * Set freemem to be 1 MB aligned at the end of boot archive. While the L1 Page
198  * table only needs to be 16 KB aligned, we opt for 1 MB alignment so that way
199  * we can map it and all the other L2 page tables we might need. If we don't do
200  * this, it'll become problematic for unix to actually modify this.
201  */
202 static void
203 fakeload_pt_arena_init(const atag_initrd_t *aii)
204 {
205     int entry, i;
206     armpte_t *pte;
207     arm_lls_t *lle;

209     pt_arena = aii->ai_start + aii->ai_size;
210     if (pt_arena & MMU_PAGEOFFSET1M) {
211         pt_arena &= MMU_PAGEMASK1M;
212         pt_arena += MMU_PAGESIZE1M;
213     }
214     pt_arena_max = pt_arena + 4 * MMU_PAGESIZE1M;
215     freemem = pt_arena_max;

217     /* Set up the l1 page table by first invalidating it */
218     pt_addr = (armpte_t *)pt_arena;
219     pt_arena += ARMPT_L1_SIZE;
220     bzero(pt_addr, ARMPT_L1_SIZE);
221     for (i = 0; i < 4; i++)
222         fakeload_map_lmb((uintptr_t)pt_addr + i * MMU_PAGESIZE1M,
223             (uintptr_t)pt_addr + i * MMU_PAGESIZE1M,
224             PF_R | PF_W);
225 }

227 /*
228  * This is our generally entry point. We're passed in the entry point of the
229  * header.
230  */
231 static uintptr_t
232 fakeload_archive_mappings(atag_header_t *chain, const void *addr,
233     atag_illumos_status_t *aisp)
234 {
235     atag_illumos_mapping_t aim;
236     fakeloader_hdr_t *hdr;
237     Elf32_Ehdr *ehdr;
238     Elf32_Phdr *phdr;
239     int nhdrs, i;
240     uintptr_t ret;
241     uintptr_t text = 0, data = 0;
242     size_t textln = 0, dataln = 0;

244     hdr = (fakeloader_hdr_t *)addr;

246     if (hdr->fh_magic[0] != FH_MAGIC0)
247         fakeload_panic("fh_magic[0] is wrong!\n");
248     if (hdr->fh_magic[1] != FH_MAGIC1)
249         fakeload_panic("fh_magic[1] is wrong!\n");

```

```

250     if (hdr->fh_magic[2] != FH_MAGIC2)
251         fakeload_panic("fh_magic[2] is wrong!\n");
252     if (hdr->fh_magic[3] != FH_MAGIC3)
253         fakeload_panic("fh_magic[3] is wrong!\n");

255     if (hdr->fh_unix_size == 0)
256         fakeload_panic("hdr unix size is zero\n");
257     if (hdr->fh_unix_offset == 0)
258         fakeload_panic("hdr unix offset is zero\n");
259     if (hdr->fh_archive_size == 0)
260         fakeload_panic("hdr archive size is zero\n");
261     if (hdr->fh_archive_offset == 0)
262         fakeload_panic("hdr archive_offset is zero\n");

264     ehdr = (Elf32_Ehdr *)((uintptr_t)addr + hdr->fh_unix_offset);

266     if (ehdr->e_ident[EI_MAG0] != ELF_MAGIC0)
267         fakeload_panic("magic[0] wrong");
268     if (ehdr->e_ident[EI_MAG1] != ELF_MAGIC1)
269         fakeload_panic("magic[1] wrong");
270     if (ehdr->e_ident[EI_MAG2] != ELF_MAGIC2)
271         fakeload_panic("magic[2] wrong");
272     if (ehdr->e_ident[EI_MAG3] != ELF_MAGIC3)
273         fakeload_panic("magic[3] wrong");
274     if (ehdr->e_ident[EI_CLASS] != ELFCLASS32)
275         fakeload_panic("wrong elfclass");
276     if (ehdr->e_ident[EI_DATA] != ELFDATA2LSB)
277         fakeload_panic("wrong encoding");
278     if (ehdr->e_ident[EI_OSABI] != ELFOSABI_SOLARIS)
279         fakeload_panic("wrong os abi");
280     if (ehdr->e_ident[EI_ABIVERSION] != EAV_SUNW_CURRENT)
281         fakeload_panic("wrong abi version");
282     if (ehdr->e_type != ET_EXEC)
283         fakeload_panic("unix is not an executable");
284     if (ehdr->e_machine != EM_ARM)
285         fakeload_panic("unix is not an ARM Executable");
286     if (ehdr->e_version != EV_CURRENT)
287         fakeload_panic("wrong version");
288     if (ehdr->e_phnum == 0)
289         fakeload_panic("no program headers");
290     ret = ehdr->e_entry;

292     FAKeload_DPRINTF("validated unix's headers\n");

294     nhdrs = ehdr->e_phnum;
295     phdr = (Elf32_Phdr *)((uintptr_t)addr + hdr->fh_unix_offset +
296         ehdr->e_phoff);
297     for (i = 0; i < nhdrs; i++, phdr++) {
298         if (phdr->p_type != PT_LOAD) {
299             fakeload_puts("skipping non-PT_LOAD header\n");
300             continue;
301         }

303         if (phdr->p_filesz == 0 || phdr->p_memsz == 0) {
304             fakeload_puts("skipping PT_LOAD with 0 file/mem\n");
305             continue;
306         }

308         /*
309          * Create a mapping record for this in the atags.
310          */
311         aim.aim_header.ah_size = ATAG_ILLUMOS_MAPPING_SIZE;
312         aim.aim_header.ah_tag = ATAG_ILLUMOS_MAPPING;
313         aim.aim_paddr = (uintptr_t)addr + hdr->fh_unix_offset +
314             phdr->p_offset;
315         aim.aim_plen = phdr->p_filesz;

```

```

316     aim.aim_vaddr = phdr->p_vaddr;
317     aim.aim_vlen = phdr->p_memsz;
318     /* Round up vlen to be a multiple of 4k */
319     if (aim.aim_vlen & 0xfff) {
320         aim.aim_vlen &= ~0xfff;
321         aim.aim_vlen += 0x1000;
322     }
323     aim.aim_mapflags = phdr->p_flags;
324     atag_append(chain, &aim.aim_header);
325
326     /*
327     * When built with highvecs we need to account for the fact that
328     * _edata, _etext and _end are built assuming that the highvecs
329     * are normally part of our segments. ld is not doing anything
330     * wrong, but this breaks the assumptions that krtld currently
331     * has. As such, unix will use this information to overwrite the
332     * normal entry points that krtld uses in a similar style to
333     * SPARC.
334     */
335     if (aim.aim_vaddr != 0xffff0000) {
336         if ((phdr->p_flags & PF_W) != 0) {
337             data = aim.aim_vaddr;
338             dataln = aim.aim_vlen;
339         } else {
340             text = aim.aim_vaddr;
341             textln = aim.aim_vlen;
342         }
343     }
344 }
345
346 aisp->ais_stext = text;
347 aisp->ais_etext = text + textln;
348 aisp->ais_sdata = data;
349 aisp->ais_edata = data + dataln;
350
351 /* 1:1 map the boot archive */
352 aim.aim_header.ah_size = ATAG_ILLUMOS_MAPPING_SIZE;
353 aim.aim_header.ah_tag = ATAG_ILLUMOS_MAPPING;
354 aim.aim_paddr = (uintptr_t)addr + hdr->fh_archive_offset;
355 aim.aim_plen = hdr->fh_archive_size;
356 aim.aim_vaddr = aim.aim_paddr;
357 aim.aim_vlen = aim.aim_plen;
358 aim.aim_mapflags = PF_R | PF_W | PF_X;
359 atag_append(chain, &aim.aim_header);
360 aisp->ais_archive = aim.aim_paddr;
361 aisp->ais_archivelen = aim.aim_plen;
362
363 return (ret);
364 }
365
366 static void
367 fakeload_mktags(atag_header_t *chain)
368 {
369     atag_illumos_status_t ais;
370     atag_illumos_mapping_t aim;
371
372     bzero(&ais, sizeof(ais));
373     bzero(&aim, sizeof(aim));
374
375     ais.ais_header.ah_size = ATAG_ILLUMOS_STATUS_SIZE;
376     ais.ais_header.ah_tag = ATAG_ILLUMOS_STATUS;
377     atag_append(chain, &ais.ais_header);
378     aim.aim_header.ah_size = ATAG_ILLUMOS_MAPPING_SIZE;
379     aim.aim_header.ah_tag = ATAG_ILLUMOS_MAPPING;
380     atag_append(chain, &aim.aim_header);
381 }

```

```

383 static uintptr_t
384 fakeload_alloc_l2pt(void)
385 {
386     uintptr_t ret;
387
388     if (pt_arena & ARMPT_L2_MASK) {
389         ret = pt_arena;
390         ret &= ~ARMPT_L2_MASK;
391         ret += ARMPT_L2_SIZE;
392         pt_arena = ret + ARMPT_L2_SIZE;
393     } else {
394         ret = pt_arena;
395         pt_arena = ret + ARMPT_L2_SIZE;
396     }
397     if (pt_arena >= pt_arena_max) {
398         fakeload_puts("pt_arena, max\n");
399         fakeload_ultrstr(pt_arena);
400         fakeload_puts("\n");
401         fakeload_ultrstr(pt_arena_max);
402         fakeload_puts("\n");
403         fakeload_puts("l2pts allocated\n");
404         fakeload_ultrstr(nl2pages);
405         fakeload_puts("\n");
406         fakeload_panic("ran out of page tables!");
407     }
408
409     bzero((void *)ret, ARMPT_L2_SIZE);
410     nl2pages++;
411     return (ret);
412 }
413
414 /*
415 * Finally, do all the dirty work. Let's create some page tables. The L1 page
416 * table is full of 1 MB mappings by default. The L2 Page table is 1k in size
417 * and covers that 1 MB. We're going to always create L2 page tables for now
418 * which will use 4k and 64k pages.
419 */
420 static void
421 fakeload_map(armpte_t *pt, uintptr_t pstart, uintptr_t vstart, size_t len,
422             uint32_t prot)
423 {
424     int entry, chunksize;
425     armpte_t *pte, *l2pt;
426     arm_llpt_t *llpt;
427
428     /*
429     * Make sure both pstart + vstart are 4k aligned, along with len.
430     */
431     if (pstart & MMU_PAGEOFFSET)
432         fakeload_panic("pstart is not 4k aligned");
433     if (vstart & MMU_PAGEOFFSET)
434         fakeload_panic("vstart is not 4k aligned");
435     if (len & MMU_PAGEOFFSET)
436         fakeload_panic("len is not 4k aligned");
437
438     /*
439     * We're going to logically deal with each 1 MB chunk at a time.
440     */
441     while (len > 0) {
442         if (vstart & MMU_PAGEOFFSET1M) {
443             chunksize = MIN(len, MMU_PAGESIZE1M -
444                             (vstart & MMU_PAGEOFFSET1M));
445         } else {
446             chunksize = MIN(len, MMU_PAGESIZE1M);
447         }

```

```

449     entry = ARMPT_VADDR_TO_L1E(vstart);
450     pte = &pt[entry];

452     if (!ARMPT_L1E_ISVALID(*pte)) {
453         uintptr_t l2table;

455         if (!(vstart & MMU_PAGEOFFSET1M) &&
456             !(pstart & MMU_PAGEOFFSET1M) &&
457             len >= MMU_PAGESIZE1M) {
458             len == MMU_PAGESIZE1M) {
459                 fakeload_map_lmb(pstart, vstart, prot);
460                 vstart += MMU_PAGESIZE1M;
461                 pstart += MMU_PAGESIZE1M;
462                 len -= MMU_PAGESIZE1M;
463                 continue;
464             }

465             l2table = fakeload_alloc_l2pt();
466             *pte = 0;
467             llpt = (arm_llpt_t *)pte;
468             llpt->al_type = ARMPT_L1_TYPE_L2PT;
469             llpt->al_ptaddr = ARMPT_ADDR_TO_L1PTADDR(l2table);
470         } else if ((*pte & ARMPT_L1_TYPE_MASK) != ARMPT_L1_TYPE_L2PT) {
471             fakeload_panic("encountered l1 entry that's not a "
472                 "pointer to a level 2 table\n");
473         } else {
474             llpt = (arm_llpt_t *)pte;
475         }

477         /* Now that we have the llpt fill in l2 entries */
478         l2pt = (void *) (llpt->al_ptaddr << ARMPT_L1PT_TO_L2_SHIFT);
479         len -= chunksize;
480         while (chunksize > 0) {
481             arm_l2e_t *l2pte;

483             entry = ARMPT_VADDR_TO_L2E(vstart);
484             pte = &l2pt[entry];

486 #ifdef MAP_DEBUG
487             fakeload_puts("4k page pa->va, l2root, entry\n");
488             fakeload_ultostr(pstart);
489             fakeload_puts("->");
490             fakeload_ultostr(vstart);
491             fakeload_puts(", ");
492             fakeload_ultostr((uintptr_t)l2pt);
493             fakeload_puts(", ");
494             fakeload_ultostr(entry);
495             fakeload_puts("\n");
496 #endif

498             if ((*pte & ARMPT_L2_TYPE_MASK) !=
499                 ARMPT_L2_TYPE_INVALID)
500                 fakeload_panic("found existing l2 page table, "
501                     "overlap in requested mappings detected!");
502             /* Map vaddr to our paddr! */
503             l2pte = ((arm_l2e_t *)pte);
504             *pte = 0;
505             if (!(prot & PF_X))
506                 l2pte->ale_xn = 1;
507             l2pte->ale_ident = 1;
508             if (prot & PF_DEVICE) {
509                 l2pte->ale_bbit = 1;
510                 l2pte->ale_cbit = 0;
511                 l2pte->ale_tex = 0;
512                 l2pte->ale_sbit = 1;

```

```

513         } else {
514             l2pte->ale_bbit = 1;
515             l2pte->ale_cbit = 1;
516             l2pte->ale_tex = 1;
517             l2pte->ale_sbit = 1;
518         }
519         if (prot & PF_W) {
520             l2pte->ale_ap2 = 1;
521             l2pte->ale_ap = 1;
522         } else {
523             l2pte->ale_ap2 = 0;
524             l2pte->ale_ap = 1;
525         }
526         l2pte->ale_ngbit = 0;
527         l2pte->ale_addr = ARMPT_PADDR_TO_L2ADDR(pstart);

529         chunksize -= MMU_PAGESIZE;
530         vstart += MMU_PAGESIZE;
531         pstart += MMU_PAGESIZE;
532     }
533 }
534 }

```

unchanged portion omitted

```

597 void
598 fakeload_init(void *ident, void *ident2, void *atag)
599 {
600     atag_header_t *hdr;
601     atag_header_t *chain = (atag_header_t *)atag;
602     const atag_initrd_t *initrd;
603     atag_illumos_status_t *aisp;
604     atag_illumos_mapping_t *aimp;
605     uintptr_t unix_start;

607     fakeload_backend_init();
608     fakeload_puts("Hello from the loader\n");
609     initrd = (atag_initrd_t *)atag_find(chain, ATAG_INITRD2);
610     if (initrd == NULL)
611         fakeload_panic("missing the initial ramdisk\n");

613     /*
614      * Create the status atag header and the initial mapping record for the
615      * atags. We'll hold onto both of these.
616      */
617     fakeload_mktag(chain);
618     aisp = (atag_illumos_status_t *)atag_find(chain, ATAG_ILLUMOS_STATUS);
619     if (aisp == NULL)
620         fakeload_panic("can't find ATAG_ILLUMOS_STATUS");
621     aimp = (atag_illumos_mapping_t *)atag_find(chain, ATAG_ILLUMOS_MAPPING);
622     if (aimp == NULL)
623         fakeload_panic("can't find ATAG_ILLUMOS_MAPPING");
624     FAKELoad_DPRINTF("created proto atags\n");

626     fakeload_pt_arena_init(initrd);

628     fakeload_selfmap(chain);

630     /*
631      * Map the boot archive and all of unix
632      */
633     unix_start = fakeload_archive_mappings(chain,
634         (const void *) (uintptr_t) initrd->ai_start, aisp);
635     FAKELoad_DPRINTF("filled out unix and the archive's mappings\n");

637     /*
638      * Fill in the atag mapping header for the atags themselves. 1:1 map it.

```

```

639     */
640     aimp->aim_paddr = (uintptr_t)chain & ~0xfff;
641     aimp->aim_plen = atag_length(chain) & ~0xfff;
642     aimp->aim_plen += 0x1000;
643     aimp->aim_vaddr = aimp->aim_paddr;
644     aimp->aim_vlen = aimp->aim_plen;
645     aimp->aim_mapflags = PF_R | PF_W | PF_NOELOC;

647     /*
648     * Let the backend add mappings
649     */
650     fakeload_backend_addmaps(chain);

652     /*
653     * Turn on unaligned access
654     */
655     FAKELOAD_DPRINTF("turning on unaligned access\n");
656     fakeload_unaligned_enable();
657     FAKELOAD_DPRINTF("successfully enabled unaligned access\n");

659     /*
660     * To turn on the MMU we need to do the following:
661     * o Program all relevant CP15 registers
662     * o Program 1st and 2nd level page tables
663     * o Invalidate and Disable the I/D-cache
664     * o Fill in the last bits of the ATAG_ILLUMOS_STATUS atag
665     * o Turn on the MMU in SCTLR
666     * o Jump to unix
667     */

669     /* Last bits of the atag */
670     aisp->ais_freemem = freemem;
671     aisp->ais_version = 1;
672     aisp->ais_ptbase = (uintptr_t)pt_addr;

674     /*
675     * Our initial page table is a series of 1 MB sections. While we really
676     * should map 4k pages, for the moment we're just going to map 1 MB
677     * regions, yay team!
678     */
679     hdr = chain;
680     FAKELOAD_DPRINTF("creating mappings\n");
681     while (hdr != NULL) {
682         if (hdr->ah_tag == ATAG_ILLUMOS_MAPPING)
683             fakeload_create_map(pt_addr,
684                 (atag_illumos_mapping_t *)hdr);
685         hdr = atag_next(hdr);
686     }

688     /*
689     * Now that we've mapped everything, update the status atag.
690     */
691     aisp->ais_freeused = freemem - aisp->ais_freemem;
692     aisp->ais_pt_arena = pt_arena;
693     aisp->ais_pt_arena_max = pt_arena_max;

695     /* Cache disable */
696     FAKELOAD_DPRINTF("Flushing and disabling caches\n");
697     armv6_dcache_flush();
698     armv6_dcache_disable();
699     armv6_dcache_inval();
700     armv6_icache_disable();
701     armv6_icache_inval();

703     /* Program the page tables */
704     FAKELOAD_DPRINTF("programming cp15 regs\n");

```

```

705     fakeload_pt_setup((uintptr_t)pt_addr);

708     /* MMU Enable */
709     FAKELOAD_DPRINTF("see you on the other side\n");
710     fakeload_mmu_enable();

712     FAKELOAD_DPRINTF("why helo thar\n");

428     /* Renable caches */
429     armv6_dcache_enable();
430     armv6_icache_enable();

714     /* we should never come back */
715     fakeload_exec(ident, ident2, chain, unix_start);
716     fakeload_panic("hit the end of the world\n");
717 }

```

unchanged portion omitted

new/usr/src/uts/armv6/loader/fakeloader_core.s

1

```
*****
2607 Thu Feb 12 20:32:17 2015
new/usr/src/uts/armv6/loader/fakeloader_core.s
loader: document the page table setup
loader: simplify MMU enabling code
*****
_____unchanged_portion_omitted_____

60 #endif /* __lint */

62 #if defined(__lint)

64 fakeload_pt_setup(uintptr_t pthead)
65 {}

67 #else /* __lint */

69 /*
70  * We need to set up the world for the first time. We'll do the
71  * following in order:
72  *
73  * o Set the TTBCR to always use TTBR0
74  * o Set domain 0 to manager mode
75  * o Program the Page table root
76  */
77 ENTRY(fakeload_pt_setup)
78 /* use TTBR0 only (should already be true) */
79 #endif /* ! codereview */
80 mov    r1, #0
81 mcr    p15, 0, r1, c2, c0, 2

83 /* set domain 0 to manager mode */
84 #endif /* ! codereview */
85 mov    r1, #3
86 mcr    p15, 0, r1, c3, c0, 0

88 /* set TTBR0 to page table root */
89 orr    r0, r0, #0x18 /* Outer WB, no WA Cachable */
90 orr    r0, r0, #0x2  /* Sharable */
91 orr    r0, r0, #0x1  /* Inner Cachable */
92 orr    r0, r0, #0x1b
93 mcr    p15, 0, r0, c2, c0, 0
94 bx     lr
95 SET_SIZE(fakeload_pt_setup)

96 #endif /* __lint */

98 #if defined(__lint)

100 /* ARGSUSED */
101 void
102 fakeload_mmu_enable(void)
103 {}

105 #else /* __lint */

107 /*
108  * We first make sure that the ARMv6 pages are enabled (bit 23) and then
109  * enable the MMU (bit 0).
110  */
111 ENTRY(fakeload_mmu_enable)
112 mrc    p15, 0, r0, c1, c0, 0
113 orr    r0, #0x800000 /* enable ARMv6 pages */
114 orr    r0, #0x1     /* enable MMU */
100 orr    r0, #0x800000
101 mcr    p15, 0, r0, c1, c0, 0
```

new/usr/src/uts/armv6/loader/fakeloader_core.s

2

```
102 mrc    p15, 0, r0, c1, c0, 0
103 orr    r0, #0x1
115 mcr    p15, 0, r0, c1, c0, 0
116 bx     lr
117 SET_SIZE(fakeload_mmu_enable)
_____unchanged_portion_omitted_____
```

```

*****
5444 Thu Feb 12 20:32:17 2015
new/usr/src/uts/armv6/ml/cache.s
armv6: p15 cache functions say that value passed in should be zero
*****
_____unchanged_portion_omitted_____

191     ENTRY(armv6_icache_inval)
192     mov     r0, #0
193 #endif /* ! codereview */
194     mcr     p15, 0, r0, c7, c5, 0      @ Invalidate i-cache
195     bx     lr
196     SET_SIZE(armv6_icache_inval)

198     ENTRY(armv6_dcache_inval)
199     mov     r0, #0
200 #endif /* ! codereview */
201     mcr     p15, 0, r0, c7, c6, 0      @ Invalidate d-cache
202     ARM_DSB_INSTR(r2)
203     bx     lr
204     SET_SIZE(armv6_dcache_inval)

206     ENTRY(armv6_dcache_flush)
207     mov     r0, #0
208 #endif /* ! codereview */
209     mcr     p15, 0, r0, c7, c10, 4     @ Flush d-cache
210     ARM_DSB_INSTR(r2)
211     bx     lr
212     SET_SIZE(armv6_dcache_flush)

214     ENTRY(armv6_text_flush_range)
215     add     r1, r1, r0
216     sub     r1, r1, r0
217     mcrr   p15, 0, r1, r0, c5         @ Invalidate i-cache range
218     mcrr   p15, 0, r1, r0, c12       @ Flush d-cache range
219     ARM_DSB_INSTR(r2)
220     ARM_ISB_INSTR(r2)
221     bx     lr
222     SET_SIZE(armv6_text_flush_range)

224     ENTRY(armv6_text_flush)
225     mov     r0, #0
226 #endif /* ! codereview */
227     mcr     p15, 0, r0, c7, c5, 0      @ Invalidate i-cache
228     mcr     p15, 0, r0, c7, c10, 4     @ Flush d-cache
229     ARM_DSB_INSTR(r2)
230     ARM_ISB_INSTR(r2)
231     bx     lr
232     SET_SIZE(armv6_text_flush)

234 #endif

236 #ifdef __lint

238 /*
239  * Perform all of the operations necessary for tlb maintenance after an update
240  * to the page tables.
241  */
242 void
243 armv6_tlb_sync(void)
244 {}

246 #else /* __lint */

248     ENTRY(armv6_tlb_sync)
249     mov     r0, #0

```

```

250     mcr     p15, 0, r0, c7, c10, 4     @ Flush d-cache
251     ARM_DSB_INSTR(r0)
252     mcr     p15, 0, r0, c8, c7, 0      @ invalidate tlb
253     mcr     p15, 0, r0, c8, c5, 0      @ Invalidate I-cache + btc
254     ARM_DSB_INSTR(r0)
255     ARM_ISB_INSTR(r0)
256     bx     lr
257     SET_SIZE(armv6_tlb_sync)

259 #endif /* __lint */

```

```

*****
5123 Thu Feb 12 20:32:17 2015
new/usr/src/uts/armv6/ml/glocore.s
unix: enable caches in locore
The loader should really be as simple as possible to be as small as
possible. It should configure the machine so that unix can make certain
assumptions but it should leave more complex initialization to unix.
*****
1 /*
2  * This file and its contents are supplied under the terms of the
3  * Common Development and Distribution License ("CDDL"), version 1.0.
4  * You may only use this file in accordance with the terms of version
5  * 1.0 of the CDDL.
6  *
7  * A full copy of the text of the CDDL should have accompanied this
8  * source. A copy of the CDDL is also available via the Internet at
9  * http://www.illumos.org/license/CDDL.
10 */

12 /*
13  * Copyright 2013 (c) Joyent, Inc. All rights reserved.
14  * Copyright (c) 2015 Josef 'Jeff' Sipek <jeffpc@josefsipek.net>
15 */

17 #include <sys/asm_linkage.h>
18 #include <sys/machparam.h>
19 #include <sys/cpu_asm.h>

21 #include "assym.h"

23 /*
24  * Every story needs a beginning. This is ours.
25 */

27 /*
28  * Each of the different machines has its own locore.s to take care of getting
29  * the machine specific setup done. Just before jumping into fakebop the
30  * first time, we call this machine specific code.
31 */

33 /*
34  * We are in a primordial world here. The loader is going to come along and
35  * boot us at _start. As we've started the world, we also need to set up a
36  * few things about us, for example our stack pointer. To help us out, it's
37  * useful to remember what the loader set up for us:
38  *
39  * - unaligned access are allowed (A = 0, U = 1)
40  * - virtual memory is enabled
41  * - we (unix) are mapped right were we want to be
42  * - a UART has been enabled & any memory mapped registers have been 1:1
43  *   mapped
44  * - ATAGs have been updated to tell us what the mappings are
45  * - I/D L1 caches have may be disabled
46  * - I/D L1 caches have been enabled
47 */

48 /*
49  * External globals
50  */
51 .globl _locore_start
52 .globl mlsetup
53 .globl sysp
54 .globl bootops
55 .globl bootopsp
56 .globl t0

```

```

58 .data
59 .comm t0stack, DEFAULTSTKSZ, 32
60 .comm t0, 4094, 32

63 /*
64  * Recall that _start is the traditional entry point for an ELF binary.
65  */
66 ENTRY(_start)
67 ldr sp, =t0stack
68 ldr r4, =DEFAULTSTKSZ
69 add sp, r4
70 bic sp, sp, #0xff

72 /*
73  * establish bogus stacks for exceptional CPU states, our exception
74  * code should never make use of these, and we want loud and violent
75  * failure should we accidentally try.
76  */
77 cps #(CPU_MODE_UND)
78 mov sp, #-1
79 cps #(CPU_MODE_ABT)
80 mov sp, #-1
81 cps #(CPU_MODE_FIQ)
82 mov sp, #-1
83 cps #(CPU_MODE_IRQ)
84 mov sp, #-1
85 cps #(CPU_MODE_SVC)

87 /* Enable highvecs (moves the base of the exception vector) */
88 mrc p15, 0, r3, c1, c0, 0
89 orr r3, r3, #(1 << 13)
90 mcr p15, 0, r3, c1, c0, 0

92 /*
93  * Go ahead now and enable the L1 I/D caches. (Involves
94  * invalidating the caches and the TLB.)
95  */
96 mov r4, #0
97 mov r5, #0
98 mcr p15, 0, r4, c7, c7, 0 /* invalidate caches */
99 mcr p15, 0, r4, c8, c7, 0 /* invalidate tlb */
100 mcr p15, 0, r5, c7, c10, 4 /* DSB */
101 mrc p15, 0, r4, c1, c0, 0
102 orr r4, #0x04 /* D-cache */
103 orr r4, #0x1000 /* I-cache */
104 mcr p15, 0, r4, c1, c0, 0

106 #endif /* ! codereview */
107 /* invoke machine specific setup */
108 bl _mach_start

110 bl _fakebop_start
111 SET_SIZE(_start)

114 #if defined(__lint)

116 /* ARGSUSED */
117 void
118 _locore_start(struct boot_syscalls *sysp, struct bootops *bop)
119 {}

121 #else /* __lint */

123 /*

```



```

124 * We got here from _kobj_init() via exitto(). We have a few different
125 * tasks that we need to take care of before we hop into mlsetup and
126 * then main. We're never going back so we shouldn't feel compelled to
127 * preserve any registers.
128 *
129 *   o Enable our I/D-caches
130 *   o Save the boot syscalls and bootops for later
131 *   o Set up our stack to be the real stack of t0stack.
132 *   o Save t0 as curthread
133 *   o Set up a struct REGS for mlsetup
134 *   o Make sure that we're 8 byte aligned for the call
135 */
136 ENTRY(_locore_start)
137
138
139 /*
140 * We've been running in t0stack anyway, up to this point, but
141 * _locore_start represents what is in effect a fresh start in the
142 * real kernel -- We'll never return back through here.
143 *
144 * So reclaim those few bytes
145 */
146 ldr    sp, =t0stack
147 ldr    r4, =(DEFAULTSTKSZ - REGSIZE)
148 add    sp, r4
149 bic    sp, sp, #0xff
150
151 /*
152 * Save flags and arguments for potential debugging
153 */
154 str    r0, [sp, #REGOFF_R0]
155 str    r1, [sp, #REGOFF_R1]
156 str    r2, [sp, #REGOFF_R2]
157 str    r3, [sp, #REGOFF_R3]
158 mrs    r4, CPSR
159 str    r4, [sp, #REGOFF_CPSR]
160
161 /*
162 * Save back the bootops and boot_syscalls.
163 */
164 ldr    r2, =sysp
165 str    r0, [r2]
166 ldr    r2, =bootops
167 str    r1, [r2]
168 ldr    r2, =bootopsp
169 ldr    r2, [r2]
170 str    r1, [r2]
171
172 /*
173 * Set up our curthread pointer
174 */
175 ldr    r0, =t0
176 mcr    p15, 0, r0, c13, c0, 4
177
178 /*
179 * Go ahead now and enable the L1 I/D caches.
180 */
181 mrc    p15, 0, r0, c1, c0, 0
182 orr    r0, #0x04 /* D-cache */
183 orr    r0, #0x1000 /* I-cache */
184 mcr    p15, 0, r0, c1, c0, 0
185
186 /*
187 * mlsetup() takes the struct regs as an argument. main doesn't take
188 * any and should never return. Currently, we have an 8-byte aligned

```

```

181 * stack. We want to push a zero frame pointer to terminate any
182 * stack walking, but that would cause us to end up with only a
183 * 4-byte aligned stack. So, to keep things nice and correct, we
184 * push a zero value twice - it's similar to a typical function
185 * entry:
186 *   push { r9, lr }
187 */
188 mov    r9, #0
189 push  { r9 } /* link register */
190 push  { r9 } /* frame pointer */
191 mov    r0, sp
192 bl    mlsetup
193 bl    main
194 /* NOTREACHED */
195 ldr    r0, =__return_from_main
196 ldr    r0, [r0]
197 bl    panic
198 SET_SIZE(_locore_start)

```

unchanged portion omitted