

[PATCH] [28/34] i386: pte xchg optimization

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-04/msg12044.html>

- *From:* Andi Kleen <ak@xxxxxxx>
 - *Date:* Mon, 30 Apr 2007 17:50:01 +0200 (CEST)
-

From: Zachary Amsden <zach@xxxxxxxxxxx>

In situations where page table updates need only be made locally, and there is no cross-processor A/D bit races involved, we need not use the heavyweight xchg instruction to atomically fetch and clear page table entries. Instead, we can just read and clear them directly.

This introduces a neat optimization for non-SMP kernels; drop the atomic xchg operations from page table updates.

Thanks to Michel Lespinasse for noting this potential optimization.

Signed-off-by: Zachary Amsden <zach@xxxxxxxxxxx>
Signed-off-by: Andrew Morton <akpm@xxxxxxxxxxxxxxxxxxxxxx>
Signed-off-by: Andi Kleen <ak@xxxxxxx>

```
include/asm-i386/pgtable-2level.h | 14 +++++
include/asm-i386/pgtable-3level.h | 14 +++++
2 files changed, 28 insertions(+)
```

Index: linux/include/asm-i386/pgtable-2level.h

=====

```
--- linux.orig/include/asm-i386/pgtable-2level.h
+++ linux/include/asm-i386/pgtable-2level.h
@@ -41,10 +41,24 @@ static inline void native_pte_clear(stru
*xp = __pte(0);
}

+/* local pte updates need not use xchg for locking */
+static inline pte_t native_local_ptep_get_and_clear(pte_t *ptep)
+{
+ pte_t res;
+
+ res = *ptep;
+ native_pte_clear(NULL, 0, ptep);
+ return res;
```

[PATCH] [28/34] i386: pte xchg optimization

```
+}
+
+#ifdef CONFIG_SMP
static inline pte_t native_ptep_get_and_clear(pte_t *xp)
{
return __pte(xchg(&xp->pte_low, 0));
}
+#else
+#define native_ptep_get_and_clear(xp) native_local_ptep_get_and_clear(xp)
+#endif

#define pte_page(x) pfn_to_page(pte_pfn(x))
#define pte_none(x) (!(x).pte_low)
Index: linux/include/asm-i386/pgtable-3level.h
=====
--- linux.orig/include/asm-i386/pgtable-3level.h
+++ linux/include/asm-i386/pgtable-3level.h
@@ -139,6 +139,17 @@ static inline void pud_clear (pud_t * pu
#define pmd_offset(pud, address) ((pmd_t *) pud_page(*(pud)) + \
pmd_index(address))

+/* local pte updates need not use xchg for locking */
+static inline pte_t native_local_ptep_get_and_clear(pte_t *ptep)
+{
+ pte_t res;
+
+ res = *ptep;
+ native_pte_clear(NULL, 0, ptep);
+ return res;
+}
+
+#ifdef CONFIG_SMP
static inline pte_t native_ptep_get_and_clear(pte_t *ptep)
{
pte_t res;
@@ -150,6 +161,9 @@ static inline pte_t native_ptep_get_and_

return res;
}
+#else
+#define native_ptep_get_and_clear(xp) native_local_ptep_get_and_clear(xp)
+#endif

#define __HAVE_ARCH_PTE_SAME
static inline int pte_same(pte_t a, pte_t b)
-
```

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>