
Subject: [PATCH] IA64,sparc: local DoS with corrupted ELF
Posted by [Kirill Korotaev](#) on Mon, 04 Sep 2006 12:14:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Linus,

this patch is already committed into -stable 2.6.17.y tree.

<http://www.kernel.org/git/?p=linux/kernel/git/stable/linux-2.6.17.y.git;a=commit;h=8833ebaa3f4325820fe3338ccf6fae04f6669254>

I only incorporated a small compilation fix from
Fernando Vazquez <fernando@oss.ntt.co.jp>.
please, commit it into next 2.6.18-rcX.

local DoS with corrupted ELF

This patch prevents cross-region mappings
on IA64 and SPARC which could lead to system crash.

davem@ confirmed: "This looks fine to me." :)

Signed-Off-By: Pavel Emelianov <xemul@openvz.org>
Signed-Off-By: Kirill Korotaev <dev@openvz.org>
Signed-off-by: Greg Kroah-Hartman <gregkh@suse.de>

```
diff --git a/arch/ia64/kernel/sys_ia64.c b/arch/ia64/kernel/sys_ia64.c
index 40722d8..9ef62a3 100644
--- a/arch/ia64/kernel/sys_ia64.c
+++ b/arch/ia64/kernel/sys_ia64.c
@@ -163,10 +163,25 @@ sys_pipe (void)
    return retval;
}

+int ia64_mmap_check(unsigned long addr, unsigned long len,
+ unsigned long flags)
+{
+ unsigned long roff;
+
+ /*
+ * Don't permit mappings into unmapped space, the virtual page table
+ * of a region, or across a region boundary. Note: RGN_MAP_LIMIT is
+ * equal to 2^n-PAGE_SIZE (for some integer n <= 61) and len > 0.
+ */
+ roff = REGION_OFFSET(addr);
```

```

+ if ((len > RGN_MAP_LIMIT) || (roff > (RGN_MAP_LIMIT - len)))
+ return -EINVAL;
+ return 0;
+}
+
static inline unsigned long
do_mmap2 (unsigned long addr, unsigned long len, int prot, int flags, int fd, unsigned long pgoff)
{
- unsigned long roff;
    struct file *file = NULL;

    flags &= ~(MAP_EXECUTABLE | MAP_DENYWRITE);
@@ -188,17 +203,6 @@ do_mmap2 (unsigned long addr, unsigned l
    goto out;
}

- /*
- * Don't permit mappings into unmapped space, the virtual page table of a region,
- * or across a region boundary. Note: RGN_MAP_LIMIT is equal to 2^n-PAGE_SIZE
- * (for some integer n <= 61) and len > 0.
- */
- roff = REGION_OFFSET(addr);
- if ((len > RGN_MAP_LIMIT) || (roff > (RGN_MAP_LIMIT - len))) {
-     addr = -EINVAL;
-     goto out;
- }
-
    down_write(&current->mm->mmap_sem);
    addr = do_mmap_pgoff(file, addr, len, prot, flags, pgoff);
    up_write(&current->mm->mmap_sem);
diff --git a/arch/sparc/kernel/sys_sparc.c b/arch/sparc/kernel/sys_sparc.c
index a41c8a5..94ff58c 100644
--- a/arch/sparc/kernel/sys_sparc.c
+++ b/arch/sparc/kernel/sys_sparc.c
@@ -219,6 +219,21 @@ out:
    return err;
}

+int sparc_mmap_check(unsigned long addr, unsigned long len, unsigned long flags)
+{
+ if (ARCH_SUN4C_SUN4 &&
+     (len > 0x20000000 || 
+      ((flags & MAP_FIXED) &&
+       (addr < 0xe0000000 && addr + len > 0x20000000)))
+ return -EINVAL;
+
+ /* See asm-sparc/uaccess.h */
+ if (len > TASK_SIZE - PAGE_SIZE || addr + len > TASK_SIZE - PAGE_SIZE)

```

```

+ return -EINVAL;
+
+ return 0;
+}
+
/* Linux version of mmap */
static unsigned long do_mmap2(unsigned long addr, unsigned long len,
    unsigned long prot, unsigned long flags, unsigned long fd,
@@ -233,25 +248,13 @@ static unsigned long do_mmap2(unsigned l
    goto out;
}

-retval = -EINVAL;
len = PAGE_ALIGN(len);
-if (ARCH_SUN4C_SUN4 &&
-    (len > 0x20000000 || ((flags & MAP_FIXED) &&
-    (addr < 0xe0000000 && addr + len > 0x20000000)))
-    goto out_putchar;
-
- /* See asm-sparc/uaccess.h */
- if (len > TASK_SIZE - PAGE_SIZE || addr + len > TASK_SIZE - PAGE_SIZE)
-    goto out_putchar;
-
flags &= ~(MAP_EXECUTABLE | MAP_DENYWRITE);

down_write(&current->mm->mmap_sem);
retval = do_mmap_pgoff(file, addr, len, prot, flags, pgoff);
up_write(&current->mm->mmap_sem);

-out_putchar:
if (file)
    fput(file);
out:
diff --git a/arch/sparc64/kernel/sys_sparc.c b/arch/sparc64/kernel/sys_sparc.c
index 054d0ab..bf5f14e 100644
--- a/arch/sparc64/kernel/sys_sparc.c
+++ b/arch/sparc64/kernel/sys_sparc.c
@@ -548,6 +548,26 @@ asmlinkage long sparc64_personality(unsi
    return ret;
}

+int sparc64_mmap_check(unsigned long addr, unsigned long len,
+    unsigned long flags)
+{
+if (test_thread_flag(TIF_32BIT)) {
+if (len >= STACK_TOP32)
+return -EINVAL;

```

```

+
+ if ((flags & MAP_FIXED) && addr > STACK_TOP32 - len)
+ return -EINVAL;
+ } else {
+ if (len >= VA_EXCLUDE_START)
+ return -EINVAL;
+
+ if ((flags & MAP_FIXED) && invalid_64bit_range(addr, len))
+ return -EINVAL;
+
+ return 0;
+}
+
/* Linux version of mmap */
asmlinkage unsigned long sys_mmap(unsigned long addr, unsigned long len,
 unsigned long prot, unsigned long flags, unsigned long fd,
@@ -563,27 +583,11 @@ @@@ asmlinkage unsigned long sys_mmap(unsign
}
flags &= ~(MAP_EXECUTABLE | MAP_DENYWRITE);
len = PAGE_ALIGN(len);
- retval = -EINVAL;
-
- if (test_thread_flag(TIF_32BIT)) {
- if (len >= STACK_TOP32)
- goto out_putf;
-
- if ((flags & MAP_FIXED) && addr > STACK_TOP32 - len)
- goto out_putf;
- } else {
- if (len >= VA_EXCLUDE_START)
- goto out_putf;
-
- if ((flags & MAP_FIXED) && invalid_64bit_range(addr, len))
- goto out_putf;
- }
-
down_write(&current->mm->mmap_sem);
retval = do_mmap(file, addr, len, prot, flags, off);
up_write(&current->mm->mmap_sem);

-out_putf:
if (file)
fput(file);
out:
diff --git a/include/asm-generic/mman.h b/include/asm-generic/mman.h
index 3b41d2b..010ced7 100644
--- a/include/asm-generic/mman.h

```

```

+++ b/include/asm-generic/mman.h
@@ -39,4 +39,10 @@ #define MADV_DOFORK 11 /* do inherit ac
#define MAP_ANON MAP_ANONYMOUS
#define MAP_FILE 0

+ifdef __KERNEL__
+ifndef arch_mmap_check
#define arch_mmap_check(addr, len, flags) (0)
+endif
+endif
+
#endif
diff --git a/include/asm-ia64/mman.h b/include/asm-ia64/mman.h
index 6ba179f..a42a3e6 100644
--- a/include/asm-ia64/mman.h
+++ b/include/asm-ia64/mman.h
@@ -8,6 +8,14 @@ #define _ASM_IA64_MMAN_H
 * David Mosberger-Tang <davidm@hpl.hp.com>, Hewlett-Packard Co
 */

+ifdef __KERNEL__
#define arch_mmap_check ia64_mmap_check
+ifndef __ASSEMBLY__
+int ia64_mmap_check(unsigned long addr, unsigned long len,
+ unsigned long flags);
+endif
+endif
+
#include <asm-generic/mman.h>

#define MAP_GROWSDOWN 0x00100 /* stack-like segment */
diff --git a/include/asm-sparc/mman.h b/include/asm-sparc/mman.h
index 88d1886..cf8e002 100644
--- a/include/asm-sparc/mman.h
+++ b/include/asm-sparc/mman.h
@@ -2,6 +2,14 @@ 
#ifndef __SPARC_MMAN_H__
#define __SPARC_MMAN_H__

+ifdef __KERNEL__
#define arch_mmap_check sparc_mmap_check
+ifndef __ASSEMBLY__
+int sparc_mmap_check(unsigned long addr, unsigned long len,
+ unsigned long flags);
+endif
+endif
+
#include <asm-generic/mman.h>

```

```

/* SunOS'ified... */
diff --git a/include/asm-sparc64/mman.h b/include/asm-sparc64/mman.h
index 6fd878e..900e2ac 100644
--- a/include/asm-sparc64/mman.h
+++ b/include/asm-sparc64/mman.h
@@ -2,6 +2,14 @@
#ifndef __SPARC64_MMAN_H__
#define __SPARC64_MMAN_H__

#ifndef __KERNEL__
#define arch_mmap_check sparc64_mmap_check
#ifndef __ASSEMBLY__
+int sparc64_mmap_check(unsigned long addr, unsigned long len,
+ unsigned long flags);
#endif
#endif
+
#include <asm-generic/mman.h>

/* SunOS'ified... */
diff --git a/mm/mmap.c b/mm/mmap.c
index c1868ec..02ecef8 100644
--- a/mm/mmap.c
+++ b/mm/mmap.c
@@ -913,6 +913,10 @@ unsigned long do_mmap_pgoff(struct file
    if (!len)
        return -EINVAL;

+ error = arch_mmap_check(addr, len, flags);
+ if (error)
+     return error;
+
/* Careful about overflows.. */
len = PAGE_ALIGN(len);
if (!len || len > TASK_SIZE)
@@ -1859,6 +1863,7 @@ unsigned long do_brk(unsigned long addr,
    unsigned long flags;
    struct rb_node **rb_link, *rb_parent;
    pgoff_t pgoff = addr >> PAGE_SHIFT;
+ int error;

len = PAGE_ALIGN(len);
if (!len)
@@ -1867,6 +1872,12 @@ unsigned long do_brk(unsigned long addr,
    if ((addr + len) > TASK_SIZE || (addr + len) < addr)
        return -EINVAL;

```

```
+ flags = VM_DATA_DEFAULT_FLAGS | VM_ACCOUNT | mm->def_flags;
+
+ error = arch_mmap_check(addr, len, flags);
+ if (error)
+ return error;
+
/*
 * mlock MCL_FUTURE?
 */
@@ -1907,8 +1918,6 @@ unsigned long do_brk(unsigned long addr,
if (security_vm_enough_memory(len >> PAGE_SHIFT))
return -ENOMEM;

- flags = VM_DATA_DEFAULT_FLAGS | VM_ACCOUNT | mm->def_flags;
-
/* Can we just expand an old private anonymous mapping? */
if (vma_merge(mm, prev, addr, addr + len, flags,
NULL, NULL, pgoff, NULL))
```
