

[patch 12/14] remap_file_pages protection support: also set VM_NONLINEAR on nonuniform VMAs

[patch 12/14] remap_file_pages protection support: also set VM_NONLINEAR on nonuniform VMAs

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-04/msg06742.html>

- *From:* blaisorblade@xxxxxxxx
 - *Date:* Sun, 30 Apr 2006 19:30:05 +0200
-

From: Paolo 'Blaisorblade' Giarrusso <blaisorblade@xxxxxxxx>

To simplify the VM code, and to reflect expected application usage, we decide to also set VM_NONLINEAR when setting VM_MANYPROTS. Otherwise, we'd have to possibly save nonlinear PTEs even on paths which cope with linear VMAs. It's possible, but intrusive (it's done in one of the next patches).

Obviously, this has a performance cost, since we potentially have to handle a linear VMA with nonlinear handling code. But I didn't know of any application which might have this usage.

XXX: update: glibc wants to replace mprotect() with linear VM_MANYPROTS areas, to handle guard pages and data mappings of shared objects.

Signed-off-by: Paolo 'Blaisorblade' Giarrusso <blaisorblade@xxxxxxxx>

Index: linux-2.6.git/mm/fremap.c

--- linux-2.6.git.orig/mm/fremap.c

+++ linux-2.6.git/mm/fremap.c

@@ -206,8 +206,9 @@ retry:

```
if (!vma->vm_private_data || (vma->vm_flags & VM_NONLINEAR)) {
```

```
/* Must set VM_NONLINEAR before any pages are populated. */
```

```
- if (pgoff != linear_page_index(vma, start) &&
```

```
- !(vma->vm_flags & VM_NONLINEAR)) {
```

```
+ if (!(vma->vm_flags & VM_NONLINEAR) &&
```

```
+ (pgoff != linear_page_index(vma, start) ||
```

```
+ pgprot_val(pgprot) != pgprot_val(vma->vm_page_prot))) {
```

```
if (!(vma->vm_flags & VM_SHARED))
```

```
goto out_unlock;
```

```
if (!has_write_lock) {
```

```
@@ -224,19 +225,19 @@ retry:
```

```
vma_nonlinear_insert(vma, &mapping->i_mmap_nonlinear);
```

```
flush_dcache_mmap_unlock(mapping);
```

```
spin_unlock(&mapping->i_mmap_lock);
```

```
- }
```

```
- if (pgprot_val(pgprot) != pgprot_val(vma->vm_page_prot) &&
```

[patch 12/14] remap_file_pages protection support: also set VM_NONLINEAR on nonuniform VMAs

[patch 12/14] remap_file_pages protection support: also set VM_NONLINEAR on nonuniform VMAs

```
-!(vma->vm_flags & VM_MANYPROTS)) {
- if (!(vma->vm_flags & VM_SHARED))
- goto out_unlock;
- if (!has_write_lock) {
- up_read(&mm->mmap_sem);
- down_write(&mm->mmap_sem);
- has_write_lock = 1;
- goto retry;
+ if (!(vma->vm_flags & VM_MANYPROTS) &&
+ pgprot_val(pgprot) != pgprot_val(vma->vm_page_prot)) {
+ if (!(vma->vm_flags & VM_SHARED))
+ goto out_unlock;
+ if (!has_write_lock) {
+ up_read(&mm->mmap_sem);
+ down_write(&mm->mmap_sem);
+ has_write_lock = 1;
+ goto retry;
+ }
+ vma->vm_flags |= VM_MANYPROTS;
+ }
- vma->vm_flags |= VM_MANYPROTS;
}
```

```
err = vma->vm_ops->populate(vma, start, size, pgprot, pgoff,
```

--

Inform me of my mistakes, so I can keep imitating Homer Simpson's "Doh!".

Paolo Giarrusso, aka Blaisorblade (Skype ID "PaoloGiarrusso", ICQ 215621894)

<http://www.user-mode-linux.org/~blaisorblade>

-

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/>

[patch 12/14] remap_file_pages protection support: also set VM_NONLINEAR on nonuniform VMAs