

[PATCH 1/3] KVM: Fix GFP_KERNEL alloc in atomic section bug

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-12/msg08131.html>

- *From:* Avi Kivity <avi@xxxxxxxxxxxx>
 - *Date:* Sun, 31 Dec 2006 13:19:51 -0000
-

From: Ingo Molnar <mingo@xxxxxxx>

KVM does kmalloc() in an atomic section while having preemption disabled via vcpu_load(). Fix this by moving the ->*_msr setup from the vcpu_setup method to the vcpu_create method.

(This is also a small speedup for setting up a vcpu, which can in theory be more frequent than the vcpu_create method).

Signed-off-by: Ingo Molnar <mingo@xxxxxxx>

Signed-off-by: Avi Kivity <avi@xxxxxxxxxxxx>

Index: linux-2.6/drivers/kvm/vmx.c

```
=====
--- linux-2.6.orig/drivers/kvm/vmx.c
+++ linux-2.6/drivers/kvm/vmx.c
@@ -1094,14 +1094,6 @@ static int vmx_vcpu_setup(struct kvm_vcp
rdmsrl(MSR_IA32_SYSENTER_EIP, a);
vmcs_writel(HOST_IA32_SYSENTER_EIP, a); /* 22.2.3 */

- ret = -ENOMEM;
- vcpu->guest_msrs = kmalloc(PAGE_SIZE, GFP_KERNEL);
- if (!vcpu->guest_msrs)
- goto out;
- vcpu->host_msrs = kmalloc(PAGE_SIZE, GFP_KERNEL);
- if (!vcpu->host_msrs)
- goto out_free_guest_msrs;
-
for (i = 0; i < NR_VMX_MSR; ++i) {
u32 index = vmx_msr_index[i];
u32 data_low, data_high;
@@ -1155,8 +1147,6 @@ static int vmx_vcpu_setup(struct kvm_vcp

return 0;

-out_free_guest_msrs:
- kfree(vcpu->guest_msrs);
out:
```

[PATCH 1/3] KVM: Fix GFP_KERNEL alloc in atomic section bug

```
return ret;
}
@@ -1906,13 +1896,33 @@ static int vmx_create_vcpu(struct kvm_vc
{
struct vmcs *vmcs;

+ vcpu->guest_msrs = kmalloc(PAGE_SIZE, GFP_KERNEL);
+ if (!vcpu->guest_msrs)
+ return -ENOMEM;
+
+ vcpu->host_msrs = kmalloc(PAGE_SIZE, GFP_KERNEL);
+ if (!vcpu->host_msrs)
+ goto out_free_guest_msrs;
+
vmcs = alloc_vmcs();
if (!vmcs)
- return -ENOMEM;
+ goto out_free_msrs;
+
vmcs_clear(vmcs);
vcpu->vmcs = vmcs;
vcpu->launched = 0;
+
return 0;
+
+out_free_msrs:
+ kfree(vcpu->host_msrs);
+ vcpu->host_msrs = NULL;
+
+out_free_guest_msrs:
+ kfree(vcpu->guest_msrs);
+ vcpu->guest_msrs = NULL;
+
+ return -ENOMEM;
}
```

```
static struct kvm_arch_ops vmx_arch_ops = {
```

```
-
```

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/>