

# [PATCH 17/40] KVM: align valid EFER bits with the features of the host system

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-03/msg12020.html>

---

- *From:* Avi Kivity <[avi@xxxxxxxxxxxxx](mailto:avi@xxxxxxxxxxxxx)>
  - *Date:* Mon, 31 Mar 2008 17:37:01 +0300
- 

From: Joerg Roedel <[joerg.roedel@xxxxxxx](mailto:joerg.roedel@xxxxxxx)>

This patch aligns the bits the guest can set in the EFER register with the features in the host processor. Currently it lets EFER.NX disabled if the processor does not support it and enables EFER.LME and EFER.LMA only for KVM on 64 bit hosts.

Signed-off-by: Joerg Roedel <[joerg.roedel@xxxxxxx](mailto:joerg.roedel@xxxxxxx)>

Signed-off-by: Avi Kivity <[avi@xxxxxxxxxxxxx](mailto:avi@xxxxxxxxxxxxx)>

---

arch/x86/kvm/svm.c | 3 +++  
arch/x86/kvm/vmx.c | 4 ++++  
arch/x86/kvm/x86.c | 10 ++++++++--  
3 files changed, 16 insertions(+), 1 deletions(-)

```
diff --git a/arch/x86/kvm/svm.c b/arch/x86/kvm/svm.c
index 1a582f1..ff3bc74 100644
--- a/arch/x86/kvm/svm.c
+++ b/arch/x86/kvm/svm.c
@@ -403,6 +403,9 @@ static __init int svm_hardware_setup(void)
 set_msr_interception(msrpm_va, MSR_IA32_SYSENTER_ESP, 1, 1);
 set_msr_interception(msrpm_va, MSR_IA32_SYSENTER_EIP, 1, 1);
```

```
+ if (boot_cpu_has(X86_FEATURE_NX))
+ kvm_enable_efer_bits(EFER_NX);
+
```

```
for_each_online_cpu(cpu) {
r = svm_cpu_init(cpu);
if (r)
```

```
diff --git a/arch/x86/kvm/vmx.c b/arch/x86/kvm/vmx.c
index 1157e8a..a509910 100644
--- a/arch/x86/kvm/vmx.c
+++ b/arch/x86/kvm/vmx.c
@@ -1117,6 +1117,10 @@ static __init int hardware_setup(void)
{
if (setup_vmcs_config(&vmcs_config) < 0)
return -EIO;
+
```

[PATCH 17/40] KVM: align valid EFER bits with the features of the host system

```
+ if (boot_cpu_has(X86_FEATURE_NX))
+ kvm_enable_efer_bits(EFER_NX);
+
return alloc_kvm_area();
}
```

```
diff --git a/arch/x86/kvm/x86.c b/arch/x86/kvm/x86.c
index ec9265b..db16f23 100644
--- a/arch/x86/kvm/x86.c
+++ b/arch/x86/kvm/x86.c
@@ -41,7 +41,15 @@
| X86_CR4_OSXMMEXCPT | X86_CR4_VMXE))
```

```
#define CR8_RESERVED_BITS (~(unsigned long)X86_CR8_TPR)
-static u64 __read_mostly efer_reserved_bits = 0xfffffffffff2fe;
+/* EFER defaults:
+ * - enable syscall per default because its emulated by KVM
+ * - enable LME and LMA per default on 64 bit KVM
+ */
+#ifdef CONFIG_X86_64
+static u64 __read_mostly efer_reserved_bits = 0xfffffffffffafeULL;
+#else
+static u64 __read_mostly efer_reserved_bits = 0xfffffffffffefeULL;
+#endif
```

```
#define VM_STAT(x) offsetof(struct kvm, stat.x), KVM_STAT_VM
#define VCPU_STAT(x) offsetof(struct kvm_vcpu, stat.x), KVM_STAT_VCPU
--
1.5.4.5
```

```
--
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/
```