

[PATCH 16/40] KVM: make EFER_RESERVED_BITS configurable for architecture code

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-03/msg12008.html>

- *From:* Avi Kivity <avi@xxxxxxxxxxxxx>
 - *Date:* Mon, 31 Mar 2008 17:37:00 +0300
-

From: Joerg Roedel <joerg.roedel@xxxxxxx>

This patch give the SVM and VMX implementations the ability to add some bits the guest can set in its EFER register.

Signed-off-by: Joerg Roedel <joerg.roedel@xxxxxxx>

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

arch/x86/kvm/x86.c | 11 ++++++++--
include/asm-x86/kvm_host.h | 1 +
2 files changed, 10 insertions(+), 2 deletions(-)

diff --git a/arch/x86/kvm/x86.c b/arch/x86/kvm/x86.c
index 6b01552..ec9265b 100644

--- a/arch/x86/kvm/x86.c
+++ b/arch/x86/kvm/x86.c
@@ -41,7 +41,7 @@
| X86_CR4_OSXMMEXCPT | X86_CR4_VMXE))

```
#define CR8_RESERVED_BITS (~(unsigned long)X86_CR8_TPR)
-#define EFER_RESERVED_BITS 0xffffffff2fe
+static u64 __read_mostly efer_reserved_bits = 0xffffffff2fe;
```

```
#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
@@ -428,7 +428,7 @@ static u32 emulated_msrs[] = {
```

```
static void set_efer(struct kvm_vcpu *vcpu, u64 efer)
{
- if (efer & EFER_RESERVED_BITS) {
+ if (efer & efer_reserved_bits) {
printk(KERN_DEBUG "set_efer: 0x%llx #GP, reserved bits\n",
efer);
kvm_inject_gp(vcpu, 0);
@@ -452,6 +452,13 @@ static void set_efer(struct kvm_vcpu *vcpu, u64 efer)
```

[PATCH 16/40] KVM: make EFER_RESERVED_BITS configurable for architecture code

#endif

```
+void kvm_enable_efer_bits(u64 mask)
+{
+ efer_reserved_bits &= ~mask;
+}
+EXPORT_SYMBOL_GPL(kvm_enable_efer_bits);
+
+
+/*
+ * Writes msr value into into the appropriate "register".
+ * Returns 0 on success, non-0 otherwise.
diff --git a/include/asm-x86/kvm_host.h b/include/asm-x86/kvm_host.h
index 67ae307..92668fa 100644
--- a/include/asm-x86/kvm_host.h
+++ b/include/asm-x86/kvm_host.h
@@ -429,6 +429,7 @@ void realmode_lmsw(struct kvm_vcpu *vcpu, unsigned long msw,
unsigned long realmode_get_cr(struct kvm_vcpu *vcpu, int cr);
void realmode_set_cr(struct kvm_vcpu *vcpu, int cr, unsigned long value,
unsigned long *rflags);
+void kvm_enable_efer_bits(u64);
int kvm_get_msr(struct kvm_vcpu *vcpu, u32 msr_index, u64 *data);
int kvm_set_msr(struct kvm_vcpu *vcpu, u32 msr_index, u64 data);
```

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/>