

[PATCH 31/40] KVM: SVM: enable LBR virtualization

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-03/msg12002.html>

- *From:* Avi Kivity <avi@xxxxxxxxxxxxx>
 - *Date:* Mon, 31 Mar 2008 17:37:15 +0300
-

From: Joerg Roedel <joerg.roedel@xxxxxxx>

This patch implements the Last Branch Record Virtualization (LBRV) feature of the AMD Barcelona and Phenom processors into the kvm-amd module. It will only be enabled if the guest enables last branch recording in the DEBUG_CTL MSR. So there is no increased world switch overhead when the guest doesn't use these MSRs.

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

Signed-off-by: Markus Rechberger <markus.rechberger@xxxxxxx>

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

arch/x86/kvm/svm.c | 39 ++++++++++++++++++++++++++++++++++++++-----
1 files changed, 37 insertions(+), 2 deletions(-)

diff --git a/arch/x86/kvm/svm.c b/arch/x86/kvm/svm.c

index 281a2ff..7d73e93 100644

--- a/arch/x86/kvm/svm.c

+++ b/arch/x86/kvm/svm.c

@@ -47,6 +47,8 @@ MODULE_LICENSE("GPL");

#define SVM_FEATURE_LBRV (1 << 1)

#define SVM_DEATURE_SVML (1 << 2)

+#define DEBUGCTL_RESERVED_BITS (~(0x3fULL))

+

/* enable NPT for AMD64 and X86 with PAE */

#if defined(CONFIG_X86_64) || defined(CONFIG_X86_PAE)

static bool npt_enabled = true;

@@ -387,6 +389,28 @@ static void svm_vcpu_init_msrpm(u32 *msrpm)

set_msr_interception(msrpm, MSR_IA32_SYSENTER_EIP, 1, 1);

}

+static void svm_enable_lbrv(struct vcpu_svm *svm)

+

+ u32 *msrpm = svm->msrpm;

+

+ svm->vmcb->control.lbr_ctl = 1;

+ set_msr_interception(msrpm, MSR_IA32_LASTBRANCHFROMIP, 1, 1);

[PATCH 31/40] KVM: SVM: enable LBR virtualization

```
+ set_msr_interception(msrpm, MSR_IA32_LASTBRANCHTOIP, 1, 1);
+ set_msr_interception(msrpm, MSR_IA32_LASTINTFROMIP, 1, 1);
+ set_msr_interception(msrpm, MSR_IA32_LASTINTTOIP, 1, 1);
+}
+
+static void svm_disable_lbrv(struct vcpu_svm *svm)
+{
+ u32 *msrpm = svm->msrpm;
+
+ svm->vmcb->control.lbr_ctl = 0;
+ set_msr_interception(msrpm, MSR_IA32_LASTBRANCHFROMIP, 0, 0);
+ set_msr_interception(msrpm, MSR_IA32_LASTBRANCHTOIP, 0, 0);
+ set_msr_interception(msrpm, MSR_IA32_LASTINTFROMIP, 0, 0);
+ set_msr_interception(msrpm, MSR_IA32_LASTINTTOIP, 0, 0);
+}
+
+static __init int svm_hardware_setup(void)
+{
+ int cpu;
+ @@ -1231,8 +1255,19 @@ static int svm_set_msr(struct kvm_vcpu *vcpu, unsigned ecx, u64 data)
+ svm->vmcb->save.sysenter_esp = data;
+ break;
+ case MSR_IA32_DEBUGCTLMR:
+ - pr_unimpl(vcpu, "%s: MSR_IA32_DEBUGCTLMR 0x%llx, nop\n",
+ - __FUNCTION__, data);
+ + if (!svm_has(SVM_FEATURE_LBRV)) {
+ + pr_unimpl(vcpu, "%s: MSR_IA32_DEBUGCTL 0x%llx, nop\n",
+ + __FUNCTION__, data);
+ + break;
+ + }
+ + if (data & DEBUGCTL_RESERVED_BITS)
+ + return 1;
+ +
+ + svm->vmcb->save.dbgctl = data;
+ + if (data & (1ULL<<0))
+ + svm_enable_lbrv(svm);
+ + else
+ + svm_disable_lbrv(svm);
+ break;
+ case MSR_K7_EVNTSEL0:
+ case MSR_K7_EVNTSEL1:
```

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/>