

[PATCH 02/40] KVM: MMU: Simplify hash table indexing

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-03/msg11991.html>

- *From:* Avi Kivity <avi@xxxxxxxxxxxx>
- *Date:* Mon, 31 Mar 2008 17:36:46 +0300

From: Dong, Eddie <eddie.dong@xxxxxxxx>

Signed-off-by: Yaozu (Eddie) Dong <eddie.dong@xxxxxxxx>

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

arch/x86/kvm/mmu.c | 10 +++++-----
include/asm-x86/kvm_host.h | 3 ++-
2 files changed, 7 insertions(+), 6 deletions(-)

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

index 28f9a44..6f8392d 100644

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

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

@@ -559,7 +559,7 @@ static void kvm_mmu_free_page(struct kvm *kvm, struct kvm_mmu_page *sp)

static unsigned kvm_page_table_hashfn(gfn_t gfn)

```
{  
- return gfn;  
+ return gfn & ((1 << KVM_MMU_HASH_SHIFT) - 1);  
}
```

static struct kvm_mmu_page *kvm_mmu_alloc_page(struct kvm_vcpu *vcpu,

@@ -663,7 +663,7 @@ static struct kvm_mmu_page *kvm_mmu_lookup_page(struct kvm *kvm, gfn_t gfn)

struct hlist_node *node;

pgprintk("%s: looking for gfn %lx\n", __FUNCTION__, gfn);

- index = kvm_page_table_hashfn(gfn) % KVM_NUM_MMU_PAGES;

+ index = kvm_page_table_hashfn(gfn);

bucket = &kvm->arch.mmu_page_hash[index];

hlist_for_each_entry(sp, node, bucket, hash_link)

if (sp->gfn == gfn && !sp->role.metaphysical) {

@@ -701,7 +701,7 @@ static struct kvm_mmu_page *kvm_mmu_get_page(struct kvm_vcpu *vcpu,

}

pgprintk("%s: looking gfn %lx role %x\n", __FUNCTION__,

gfn, role.word);

- index = kvm_page_table_hashfn(gfn) % KVM_NUM_MMU_PAGES;

+ index = kvm_page_table_hashfn(gfn);

bucket = &vcpu->kvm->arch.mmu_page_hash[index];

[PATCH 02/40] KVM: MMU: Simplify hash table indexing

```
hlist_for_each_entry(sp, node, bucket, hash_link)
if (sp->gfn == gfn && sp->role.word == role.word) {
@@ -840,7 +840,7 @@ static int kvm_mmu_unprotect_page(struct kvm *kvm, gfn_t gfn)

pgprintk("%s: looking for gfn %lx\n", __FUNCTION__, gfn);
r = 0;
- index = kvm_page_table_hashfn(gfn) % KVM_NUM_MMU_PAGES;
+ index = kvm_page_table_hashfn(gfn);
bucket = &kvm->arch.mmu_page_hash[index];
hlist_for_each_entry_safe(sp, node, n, bucket, hash_link)
if (sp->gfn == gfn && !sp->role.metaphysical) {
@@ -1450,7 +1450,7 @@ void kvm_mmu_pte_write(struct kvm_vcpu *vcpu, gpa_t gpa,
vcpu->arch.last_pt_write_count = 1;
vcpu->arch.last_pte_updated = NULL;
}
- index = kvm_page_table_hashfn(gfn) % KVM_NUM_MMU_PAGES;
+ index = kvm_page_table_hashfn(gfn);
bucket = &vcpu->kvm->arch.mmu_page_hash[index];
hlist_for_each_entry_safe(sp, node, n, bucket, hash_link) {
if (sp->gfn != gfn || sp->role.metaphysical)
diff --git a/include/asm-x86/kvm_host.h b/include/asm-x86/kvm_host.h
index 4702b04..d6db0de 100644
--- a/include/asm-x86/kvm_host.h
+++ b/include/asm-x86/kvm_host.h
@@ -57,7 +57,8 @@

#define KVM_PERMILLE_MMU_PAGES 20
#define KVM_MIN_ALLOC_MMU_PAGES 64
-#define KVM_NUM_MMU_PAGES 1024
+#define KVM_MMU_HASH_SHIFT 10
+#define KVM_NUM_MMU_PAGES (1 << KVM_MMU_HASH_SHIFT)
#define KVM_MIN_FREE_MMU_PAGES 5
#define KVM_REFILL_PAGES 25
#define KVM_MAX_CPUID_ENTRIES 40
--
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/>