

Xen VGA Passthrough with Xen 4.2-unstable Changeset 25070 in Ubuntu 11.10 amd64 Final Release Dom0

Version 1.0

Author: Teo En Ming (Zhang Enming)
Website #1: <http://www.teo-en-ming.com>
Website #2: <http://www.zhang-enming.com>
Email #1: teo.en.ming@gmail.com
Email #2: teo-en-ming@teo-en-ming.com
Email #3: teo-en-ming@zhang-enming.com
Mobile Phone(s): +65-8369-2618 / +65-9323-5112 / +65-9465-2119
Country: Singapore
Date: 21 March 2012 Wed 2:59 P.M. Singapore Time

1 Preparing the USB Flash Drive to Extract VGA Card EEPROM

Reference Documentation URL: <http://www.davidgis.fr/blog/index.php?2011/12/07/860-xen-42unstable-patches-for-vga-pass-through>

```
wget http://www.davidgis.fr/download/nvflash\_5.100.1\_usb.iso.tar.bz2  
tar xfvj nvflash_5.100.1_usb.iso.tar.bz2
```

Plug in your USB flash drive.

```
dmesg
```

In my case, the USB flash drive is detected as /dev/sdb.

```
mount | grep sdb
```

```
/dev/sdb1 on /media/C06F-905B type vfat  
(rw,nosuid,nodev,uid=1000,gid=1000,shortname=mixed,dmask=0077,utf8=1,showexec,flush,uhel  
per=udisks)
```

```
sudo umount /media/C06F-905B/
```

```
sudo dd if=nvflash_5.100.1_usb.iso of=/dev/sdb
```

Reboot your computer with the USB flash drive plugged in.

```
nvflash.exe -list (OPTIONAL)  
nvflash.exe -save vgabios.rom
```

Unplug your USB flash drive. Reboot your computer back into Linux Xen Dom0. Plug in your USB flash drive again.

```
cp /media/LEXAR/VGABIOS.ROM /home/teo-en-ming/2nd-palit-nvidia-geforce-8400gs-vgabios.rom
```

2 Patching Xen 4.2-unstable Changeset 25070 for Xen VGA Passthrough

```
cd
hg clone http://xenbits.xen.org/xen-unstable.hg xen-unstable.hg-cs25070-vga-passthrough
cd xen-unstable.hg-cs25070-vga-passthrough
./configure
make world
make clean
```

Download Xen VGA Passthrough patches from David Techer's (Frenchman) website.

```
wget http://www.davidgis.fr/download/xen-4.2\_rev24798\_gfx-passthrough-patches.tar.bz2
tar xfvj xen-4.2_rev24798_gfx-passthrough-patches.tar.bz2
```

Patching Xen 4.2-unstable changeset 25070 source tree.

```
patch -p1 < xen-4.2_rev24798_gfx-passthrough-patches/patch_Makefile
patch -p1 < xen-4.2_rev24798_gfx-passthrough-patches/patch_dsdt.asl
patch -p1 < xen-4.2_rev24798_gfx-passthrough-patches/patch_hvmloder.c
patch -p1 < xen-4.2_rev24798_gfx-passthrough-patches/patch_rombios.c
patch -p1 < xen-4.2_rev24798_gfx-passthrough-patches/patch_pci.c
patch -p1 < xen-4.2_rev24798_gfx-passthrough-patches/patch_pass-through.c
```

3 Configuring MMIO BARS

```
lspci | grep VGA
```

```
01:00.0 VGA compatible controller: nVidia Corporation GT218 [GeForce 8400 GS] (rev a2)
```

```
dmesg | grep 01:00.0 | grep "pci.*mem"
```

```
[ 0.120488] pci 0000:01:00.0: reg 10: [mem 0xd2000000-0xd2ffffff]
[ 0.120508] pci 0000:01:00.0: reg 14: [mem 0xc0000000-0xcfffffff 64bit pref]
[ 0.120528] pci 0000:01:00.0: reg 1c: [mem 0xd0000000-0xd1ffffff 64bit pref]
[ 0.120556] pci 0000:01:00.0: reg 30: [mem 0xd3000000-0xd307ffff pref]
```

4 Calculating Differences

4.1 First Range

Maximum = 0xd2ffffff = 3539992575

Minimum = 0xd2000000 = 3523215360

Difference = Max – Min + 1 = 3539992575 – 3523215360 + 1 = 16777216 = 0x01000000

Hence,

Max = 0xD2FFFFFF

Min = 0xD2000000

Diff = 0x01000000

4.2 Second Range

Maximum = 0xcfffffff = 3489660927

Minimum = 0xc0000000 = 3221225472

Difference = Max – Min + 1 = 3489660927 – 3221225472 + 1 = 268435456 = 0x10000000

Hence,

Max = 0xCFFFFFFF

Min = 0xC0000000

Diff = 0x10000000

4.3 Third Range

Maximum = 0xd1ffffff = 3523215359

Minimum = 0xd0000000 = 3489660928

Difference = Max – Min + 1 = 3523215359 – 3489660928 + 1 = 33554432 = 0x02000000

Hence,

Max = 0xD1FFFFFF

Min = 0xD0000000

Diff = 0x02000000

5 Important Mathematical Tool (Online)

Link: <http://easycalculation.com/hex-converter.php>

6 Modifying tools/firmware/hvmloder/acpi/dsdt.asl

vi tools/firmware/hvmloder/acpi/dsdt.asl

```
/* reserve MMIO BARs of gfx for 1:1 mapping */
DWordMemory(
    ResourceProducer, PosDecode, MinFixed, MaxFixed,
    Cacheable, ReadWrite,
    0x00000000,
    0xD2000000,
    0xD2FFFFFF,
    0x00000000,
    0x01000000)

DWordMemory(
    ResourceProducer, PosDecode, MinFixed, MaxFixed,
    NonCacheable, ReadWrite,
    0x00000000,
    0xC0000000,
    0xCFFFFFFF,
    0x00000000,
    0x10000000)

DWordMemory(
    ResourceProducer, PosDecode, MinFixed, MaxFixed,
    Cacheable, ReadWrite,
    0x00000000,
    0xD0000000,
    0xD1FFFFFF,
    0x00000000,
    0x02000000)
```

7 Copying the VGA BIOS of Palit NVIDIA Geforce 8400 GS PCI-e x16

```
cp /home/teo-en-ming/2nd-palit-nvidia-geforce-8400gs-vgabios.rom
tools/firmware/vgabios/vgabios-pt.bin
```

8 Building and Installing Xen 4.2-unstable Changeset 25070

```
make world
sudo make install
```

9 pciback (Not Using)

```
sudo nano /etc/grub.d/40_custom
```

```
menuentry 'Ubuntu 11.10 Release with Xen 4.1.3-rc1-pre and Kernel 3.3.0-xen-teo.en.ming-sgp'
--class gnu-linux --class gnu --class os {
    recordfail
    insmod part_msdos
    insmod ext2
    search --no-floppy --fs-uuid --set=root fd1ee157-7822-4a08-8549-56f4ae96f0dc
    set root='(/dev/sda,msdos1)'
    search --no-floppy --fs-uuid --set=root fd1ee157-7822-4a08-8549-56f4ae96f0dc
    multiboot /boot/xen.gz
    module /boot/vmlinuz-3.3.0-xen-teo.en.ming-sgp placeholder root=UUID=fd1ee157-7822-
4a08-8549-56f4ae96f0dc dom0_mem=1024 console=tty quiet splash vt.handoff=7 xen-
pciback.hide=(01:00.0)
    module /boot/initrd.img-3.3.0-xen-teo.en.ming-sgp
}
```

```
sudo update-grub
```

10 XL Domain Configuration File for Windows XP Home Edition SP3 HVM domU

```
# XL domain configuration file for Windows XP Home Edition SP3 HVM domU
# Please refer to "man xl.cfg" for further explanations.
# See also docs/misc/xl-network-configuration.markdown and
# docs/misc/xl-disk-configuration.txt

# Written by Teo En Ming (Zhang Enming)
# Email: teo.en.ming@gmail.com
# Mobile Phone: +65-8369-2618
# Country: Singapore
# Date: 18 Mar 2012 Sun

name="WindowsXPHomeEditionSP3"

builder="hvm"

vcpus=2

memory=512

on_poweroff="destroy"
on_reboot="restart"
on_crash="destroy"

disk=[ 'format=raw, vdev=hda, access=rw, target=/var/lib/libvirt/images/Windows-XP-Home-
Edition.img' ]

vif=[ 'bridge=virbr0,type=ioemu,model=rtl8139' ]

#boot=[c|d|n]
```

```
#      Selects the emulated virtual device to boot from. Options are hard disk (c), cd-rom (d) or
network/PXE (n).
#      Multiple options can be given and will be attempted in the order they are given. e.g. to
boot from cd-rom
#      but fallback to the hard disk you can give dc. The default is cd.

boot="dc"

acpi=1

xen_platform_pci=1

viridian=1

stdvga=1

vnc=1
vnclisten="192.168.1.2"
vncdisplay=1
vncunused=1
vncpasswd=""
sdl=0

usb=1
usbdevice="tablet"

gfx_passthru=1

pci = [ '01:00.0' ]
```

11 pci-stub

Prevents nouveau kernel module/vga driver from loading.

```
sudo nano /etc/modprobe.d/blacklist.conf
```

```
blacklist nouveau
```

Uninstall the lightdm display manager. Previous versions of Ubuntu uses gdm.

```
sudo apt-get remove lightdm
```

Reboot your computer.

```
sudo reboot
```

Load the pci_stub module.

```
sudo modprobe pci-stub
```

```
lsmod | grep pci_stub
```

Palit NVIDIA Geforce 8400 GS PCI Express x16 VGA card

```
lspci | grep VGA
```

```
lspci -n | grep "01:00.0"
```

```
01:00.0 0300: 10de:10c3 (rev a2)
```

Create a shell script to start Windows HVM domU.

```
cd  
nano start-windows
```

```
#!/bin/sh  
set -x  
sudo chmod o+w /sys/bus/pci/drivers/pci-stub/new_id  
sudo chmod o+w /sys/bus/pci/devices/0000:01:00.0/driver/unbind  
sudo chmod o+w /sys/bus/pci/drivers/pci-stub/bind  
echo "10de 10c3" > /sys/bus/pci/drivers/pci-stub/new_id  
echo "0000:01:00.0" > /sys/bus/pci/devices/0000:01:00.0/driver/unbind  
echo "0000:01:00.0" > /sys/bus/pci/drivers/pci-stub/bind  
sudo xl create /etc/xen/WindowsXPHomeEditionSP3
```

```
sudo chmod +x start-windows
```

You may have to execute the shell script start-windows twice.

```
./start-windows  
./start-windows
```

12 Checking Whether Intel VT-d is Enabled

```
sudo xl dmesg | grep 'I/O virtualisation'
```

```
(XEN) I/O virtualisation enabled
```

13 Error Logs in /var/log/xen

13.1 xl-WindowsXPHomeEditionSP3.log

```
Waiting for domain WindowsXPHomeEditionSP3 (domid 1) to die [pid 3084]
```

13.2 qemu-dm-WindowsXPHomeEditionSP3.log

```
domid: 1
Strip off blktap sub-type prefix to /var/lib/libvirt/images/Windows-XP-Home-Edition.img (drv
'aio')
Using file /var/lib/libvirt/images/Windows-XP-Home-Edition.img in read-write mode
Watching /local/domain/0/device-model/1/logdirty/cmd
Watching /local/domain/0/device-model/1/command
Watching /local/domain/1/cpu
qemu_map_cache_init nr_buckets = 10000 size 4194304
shared page at pfn feffd
buffered io page at pfn feffb
Guest uuid = a417a52a-68e4-48a1-89cc-3216184b8ed3
Register xen platform.
Done register platform.
platform_fixed_ioport: changed ro/rw state of ROM memory area. now is rw state.
xs_read(/local/domain/0/device-model/1/xen_extended_power_mgmt): read error
Log-dirty: no command yet.
I/O request not ready: 0, ptr: 0, port: 0, data: 0, count: 0, size: 0
I/O request not ready: 0, ptr: 0, port: 0, data: 0, count: 0, size: 0
vcpu-set: watch node error.
xs_read(/local/domain/1/log-throttling): read error
qemu: ignoring not-understood drive `/local/domain/1/log-throttling'
medium change watch on `/local/domain/1/log-throttling' - unknown device, ignored
dm-command: hot insert pass-through pci dev
register_real_device: Assigning real physical device 01:00.0 ...
pt_iomul_init: Error: pt_iomul_init can't open file /dev/xen/pci_iomul: No such file or directory:
0x1:0x0.0x0
pt_register_regions: IO region registered (size=0x01000000 base_addr=0xd2000000)
pt_register_regions: IO region registered (size=0x10000000 base_addr=0xc000000c)
pt_register_regions: IO region registered (size=0x02000000 base_addr=0xd000000c)
pt_register_regions: IO region registered (size=0x00000080 base_addr=0x0000d001)
pt_register_regions: Expansion ROM registered (size=0x00080000 base_addr=0xd3000002)
setup_vga_pt: vga bios checksum is adjusted!
pt_msi_setup: msi mapped with pirq 37
pci_intx: intx=1
register_real_device: Real physical device 01:00.0 registered successfully!
IRQ type = MSI-INTx
pt_bar_reg_read: first read BARs of gfx
pt_iomem_map: e_phys=d2000000 maddr=d2000000 type=0 len=16777216 index=0 first_map=1
pt_bar_reg_read: first read BARs of gfx
pt_iomem_map: e_phys=c0000000 maddr=c0000000 type=8 len=268435456 index=1
```



```
first_map=1
pt_bar_reg_read: first read BARs of gfx
pt_bar_reg_read: first read BARs of gfx
pt_iomem_map: e_phys=d0000000 maddr=d0000000 type=8 len=33554432 index=3 first_map=1
pt_bar_reg_read: first read BARs of gfx
pt_bar_reg_read: first read BARs of gfx
pt_ioport_map: e_phys=d000 pio_base=d000 len=128 index=5 first_map=1
platform_fixed_ioport: changed ro/rw state of ROM memory area. now is rw state.
platform_fixed_ioport: changed ro/rw state of ROM memory area. now is ro state.
pt_iomem_map: e_phys=ffffffff maddr=d2000000 type=0 len=16777216 index=0 first_map=0
pt_iomem_map: e_phys=ffffffff maddr=c0000000 type=8 len=268435456 index=1 first_map=0
pt_iomem_map: e_phys=ffffffff maddr=d0000000 type=8 len=33554432 index=3 first_map=0
pt_ioport_map: e_phys=ffff pio_base=d000 len=128 index=5 first_map=0
pt_iomem_map: e_phys=d2000000 maddr=d2000000 type=0 len=16777216 index=0 first_map=0
pt_iomem_map: e_phys=c0000000 maddr=c0000000 type=8 len=268435456 index=1
first_map=0
pt_iomem_map: e_phys=d0000000 maddr=d0000000 type=8 len=33554432 index=3 first_map=0
pt_ioport_map: e_phys=d000 pio_base=d000 len=128 index=5 first_map=0
```