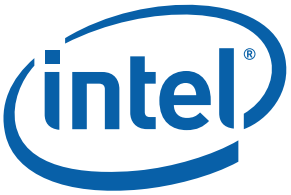

How to Configure Intel® X520 Ethernet Server Adapter Based Virtual Functions on Citrix* XenServer 6.0*

Technical Brief v1.0

December 2011

Contents

Revision History

Date	Revision	Description
December 2011	1.0	Initial release

1.0 Introduction

Processor and platform advances, along with progress in operating systems and applications, have spurred data centers to quickly adopt 10 Gigabit Ethernet (10 GbE) as the standard for interconnects. The Intel® Ethernet Server Adapter X520 family of adapters introduces numerous industry-leading features that are helping data center administrators implement innovative solutions for difficult and challenging connectivity problems. I/O Virtualization is one of the fastest growing usage models within the data center. The X520 family of adapters provides the ability to create Virtual Functions (VFs) that are identical instantiations of the Physical Functions (PFs). VFs are capable of providing 10 GbE connectivity to Virtual Machines (VMs) within a virtualized operating system framework. This document shows how to make use of VFs using Citrix XenServer* 6.0.

1.1 Hardware Requirements

- An Intel® Ethernet X520 Server Adapter
- A server platform that supports Intel® Virtualization Technology for Directed I/O (VT-d) and the PCI-SIG* Single Root I/O Virtualization and Sharing (SR-IOV) specification1
- A server platform with an available PCI Express*:X8 5.0 Gb/s (Gen2) slot

1.2 Software Requirements

- Citrix XenServer* 6.0

2.0 Installation and Configuration

1. Install the Intel® Ethernet X520 server adapter in an available PCI-Express x8 slot. (Ensure that the x8 slot is electrically connected as a x8; some slots are physically x8 but electrically support only x4. Verify this with your server manufacturer or system documentation.)
2. Power up the server.
3. Enter the server's BIOS setup and make sure the virtualization technology and Intel® VT-d features are enabled.
4. Install Citrix XenServer 6.0.
5. Log in to the newly-installed Citrix XenServer operating system using the Citrix XenCenter application.

2.1 Enabling I/O Memory Management Unit (IOMMU)

I/O Memory Management Unit (IOMMU) support is not enabled by default in Citrix XenServer 6.0. IOMMU support is required for a VF to function properly when assigned to a VM. The following kernel boot parameters are required to enable IOMMU support for the Citrix XenServer 6.0 kernel:

```
iommu=1
```

This parameter can be appended to the kernel boot entry by executing the following command at the XenServer console. See Figure 1.

```
/opt/xensource/libexec/xen-cmdline --set-xen iommu=1
```

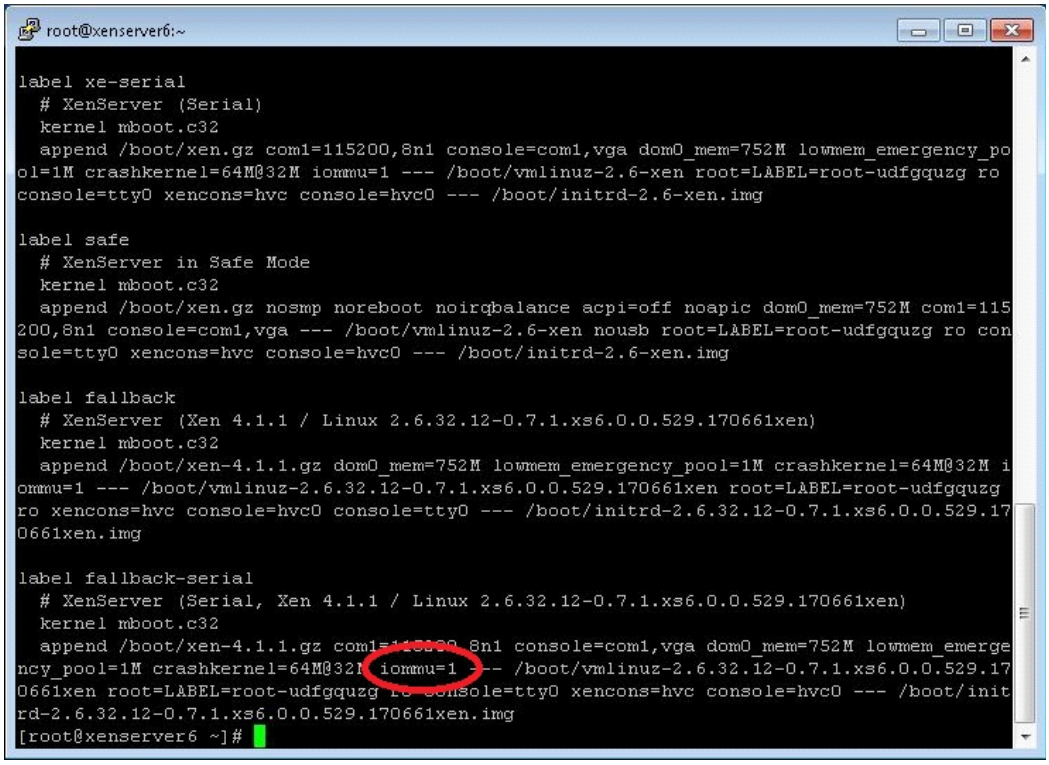

Figure 1. Enabling "iommu" by Appending It to the Boot Kernel Entry

Figure 1 shows the contents of the `/boot/extlinux.conf` file. Reboot the server for the iommu change to take effect (required).

Upon a successful server reboot, use the `dmesg` command at the XenServer console to confirm iommu is successfully initialized. See Figure 2.

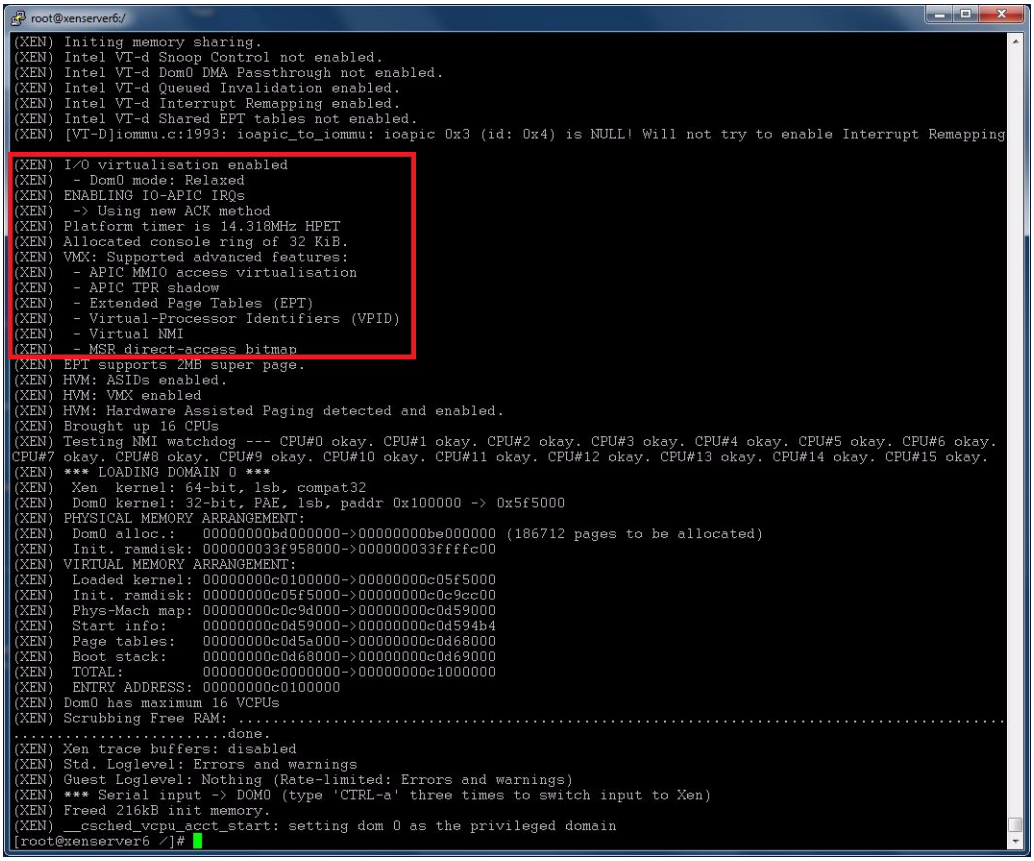

Figure 2. XenServer Console Display Showing Successful "iommu" Initialization

2.2 Intel® X520 Ethernet Server Adapter Driver Configuration

1. Citrix XenServer 6.0 includes the X520 Server Adapter driver. XenServer loads the driver at boot time and creates 40 virtual functions (VFs) per port by default. The X520 Server Adapter supports up to 63 VFs per port.
2. The number of VFs created can be adjusted by using the `ixgbe` driver load-time parameter called "`max_vfs`." This parameter is located in the driver configuration file called "`ixgbe`." The "`ixgbe`" file is located in the /etc/modprobe.d/ folder in Citrix XenServer 6.0. Figure 3 shows the default `ixgbe` configuration file contents:

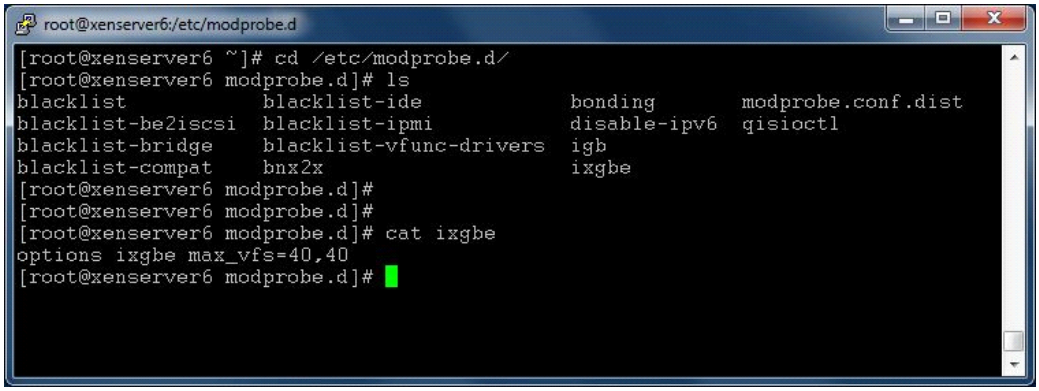

Figure 3. Default Contents of the "ixgbe" Configuration File

Any warnings, errors and information messages during the `ixgbe` driver load are logged in the /var/log/messages file. It is a good practice to review this file to confirm that the driver loaded successfully without warnings or errors.

Use the `lspci` command to confirm that the VF was successfully created. An example entry:

```
#lspci | grep 82599
```

Figure 4 shows the results of this command.

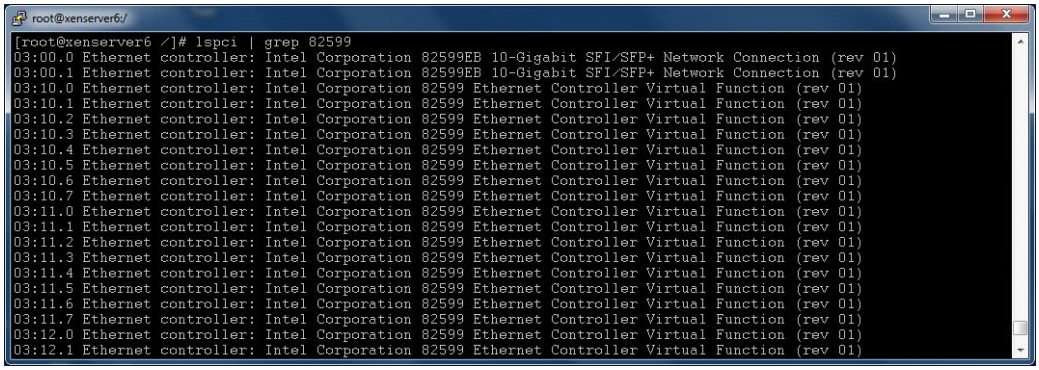

Figure 4. VFs Per Port: Bus #, Device #, and Function #

Figure 4 shows the VFs per port for the X520 Server Adapter. Each VF is identified by an unique bus, device, and function number. In the example, the first VF is assigned Bus #3, Device #10, Function #0. VFs ending with an even function number belong to PF 0, that is, port 0. VFs ending with an odd function number belong to PF 1, that is, port 1.

2.3 Preparing Virtual Functions for Assignment

Citrix XenServer 6.0 requires SR-IOV-compliant, X520 Server Adapter-based Virtual Devices (VFs) not to be claimed by the Xen Domain-0 if the VFs are intended to be assigned to Virtual Machines. VF bus, device and function numbers are required for

preparing VFs for Virtual Machine assignment. Execute the following two commands at the XenServer console for each VF that is intended for Virtual Machine assignment. See Figure 5.

```
echo 0000:03:10.2 > /sys/bus/pci/drivers/pciback/new_slot

echo 0000:03:10.2 > /sys/bus/pci/drivers/pciback/bind
```

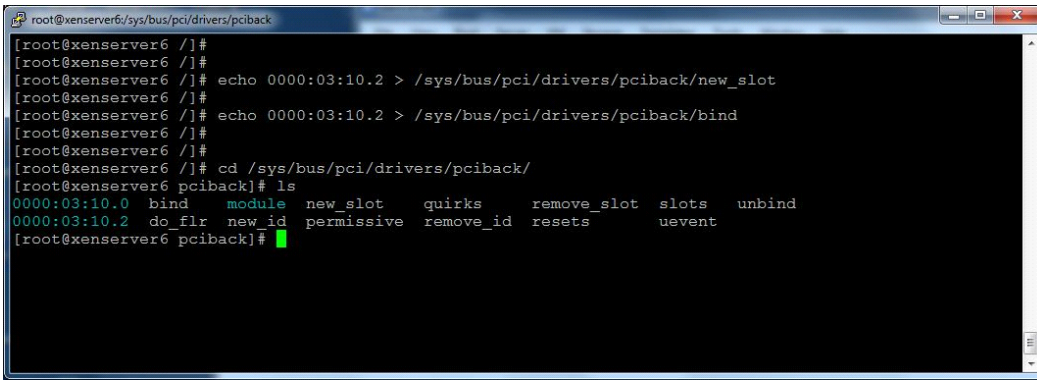

Figure 5. Results of `new_slot` and `bind` Commands With Link Node Information for the VF Identified by Bus #3, Device #10, and Function #2

Figure 5 also shows the link node for the VF identified by Bus # 3, Device # 10 and Function #2 that was created as a result of the commands. Citrix XenServer 6.0 will create these nodes for all the VFs that are intended for VM assignment

Note: Intel strongly recommends creating a script with the two commands above for all the VFs that are intended for VM assignment. This script should run every time XenServer is rebooted.

Use the following command to see available VFs and their identification attributes for assignment to a Virtual Machine.

```
#xl pci-list-assignable
```

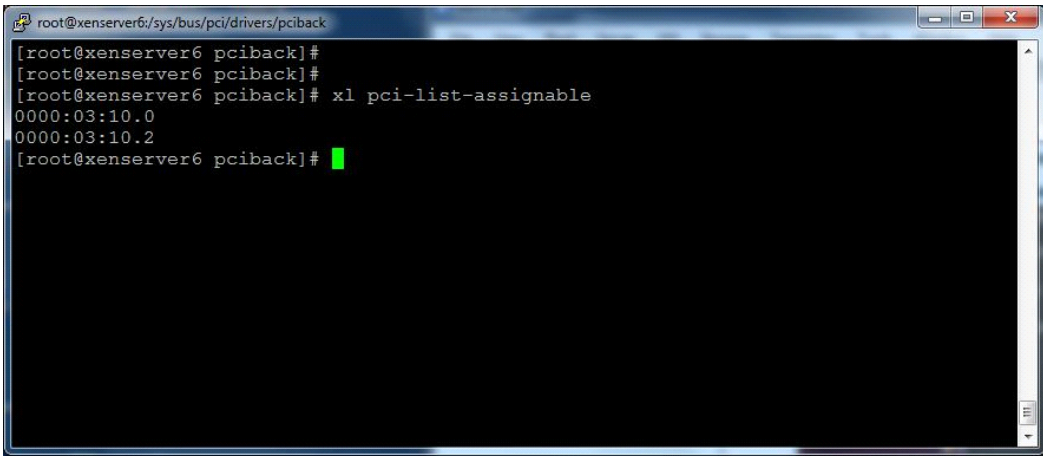

Figure 6. Display Showing Available VFs and Their ID Attributes

Figure 6 above shows the list of VFs that are available for assignment to Virtual Machines and their respective bus, device and function numbers.

2.4 Virtual Machine Installation and Configuration

Citrix XenServer 6.0 includes tools for creating and managing VMs. These tools offer both Command Line Interface (CLI) and Graphical User Interface (GUI). Citrix XenCenter is a GUI tool for creating and managing VMs.

1. Use Citrix XenCenter to create a VM.
2. Install the operating system of choice on the newly created VM. For the purposes of this document, Red Hat Enterprise Linux* was installed in the VM. See Figure 7 for an example:

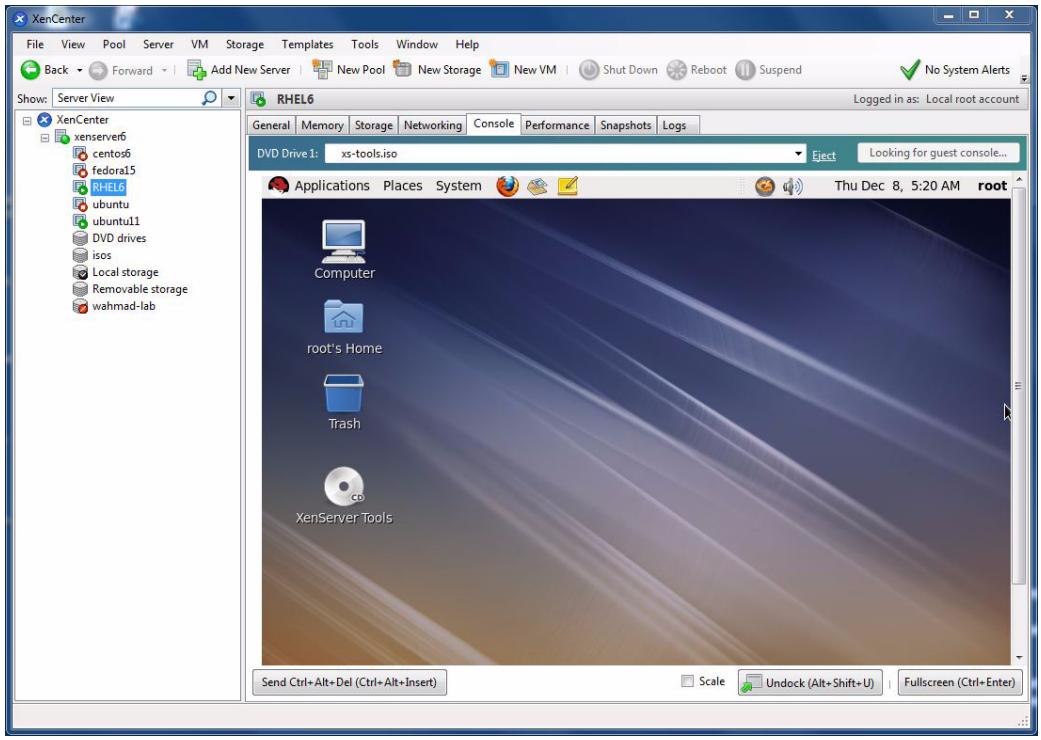

Figure 7. Red Hat Enterprise Linux* Installation Example

3. Citrix XenCenter doesn't provide a method via the GUI to assign PCI devices to Virtual Machines. Use the XenServer console commands below to list all available Virtual Machines and to assign VFs to a VM.

```
#xl list-vm

#xe vm-param-add other-config:pci=0/0000:03:10.0 uuid=bdf5...
```

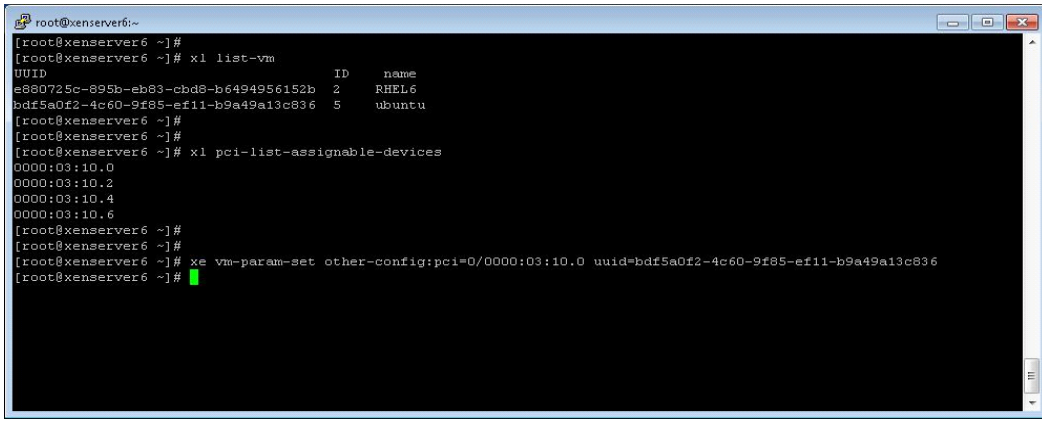

Figure 8. List of VMs, VFs Available, and Assignment to VM "RHEL6"

Figure 8 above shows the list of Virtual Machines configured on the server, the list of VFs available for assignment and subsequent VF assignment to Virtual Machine named "RHEL6."

4. Power up the RHEL6 VM. Log into the VM using the credentials created during the VM installation process. The VM must be powered down and powered back up in case the VF assignment took place while the VM was running.
5. Open the Linux Terminal. Use the Linux `lspci` utility to confirm that the assigned VF is visible within the VM's PCI-Express hierarchy. See Figure 9:

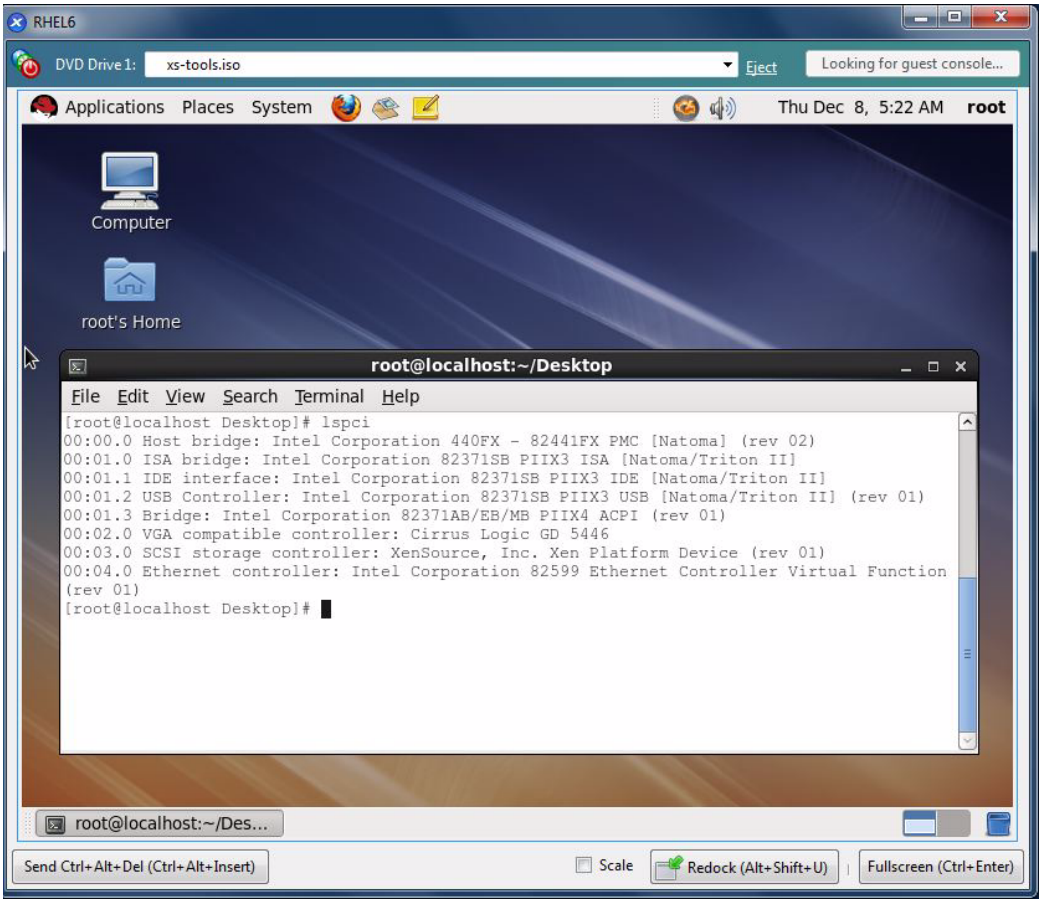

Figure 9. Linux Terminal Running the "`lspci`" Utility

6. Use the Linux `lsmod` utility to confirm that the driver for the VF has loaded successfully, as shown in Figure 10:

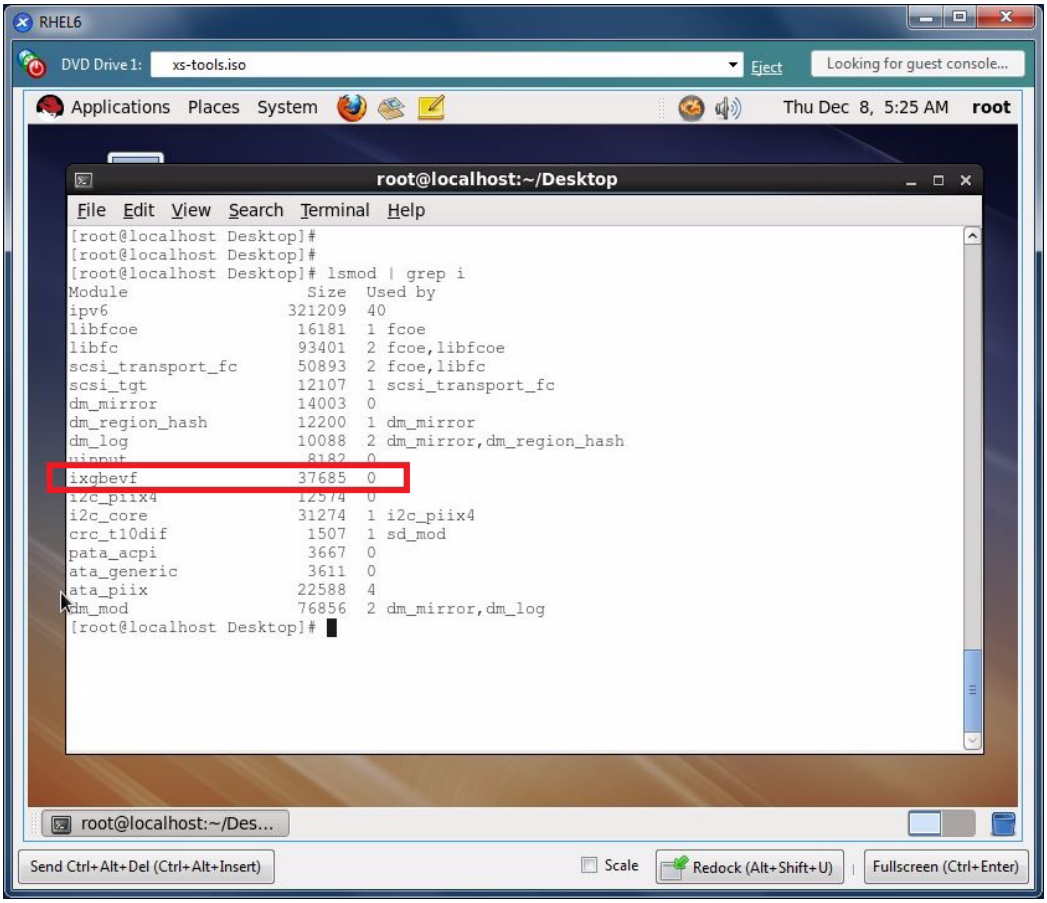

Figure 10. Linux Terminal Using "`lsmod`" Utility to Show Successful Driver Load

7. Use the Linux ifconfig utility to confirm that the newly assigned VF is ready for use (Figure 11).

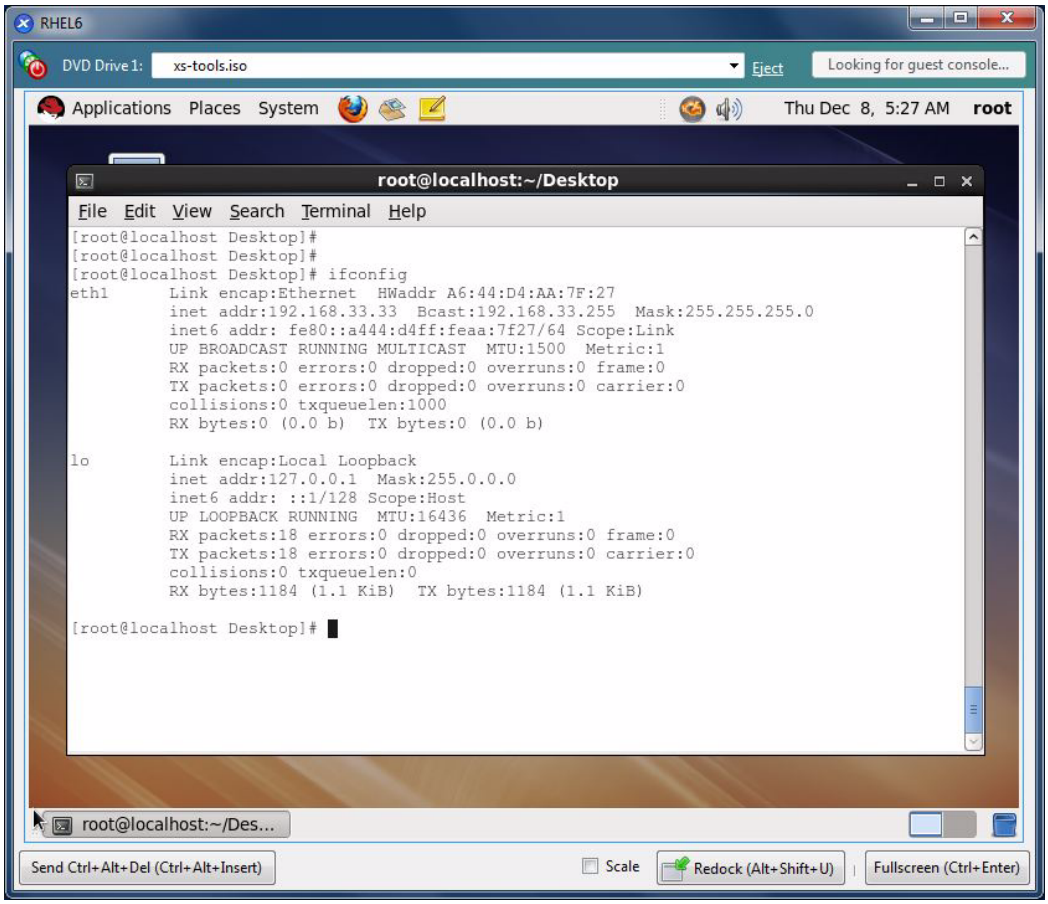

Figure 11. Linux Terminal Using "`ifconfig`" Utility Confirming Assigned VF Is Ready

The VF can be configured for DHCP or static IP address assignment. The VF is ready to communicate once it has an IP address assigned.

2.5 Enabling Hot-Plug Device Support

Hot plug device assignment to Virtual Machines is not enabled by default in Citrix XenServer 6.0. It is possible to enable hot pluggable device assignment support to Virtual Machines in XenServer 6.0. Citrix doesn't support this feature at this time; use this feature at your own risk.

1. Hot plug device assignment can be enabled by editing the "`qemu-dm-wrapper`" file that is located in the /opt/xensource.libexec/ folder. Add the following line in "`def main(argv):`" function, as shown in Figure 12:

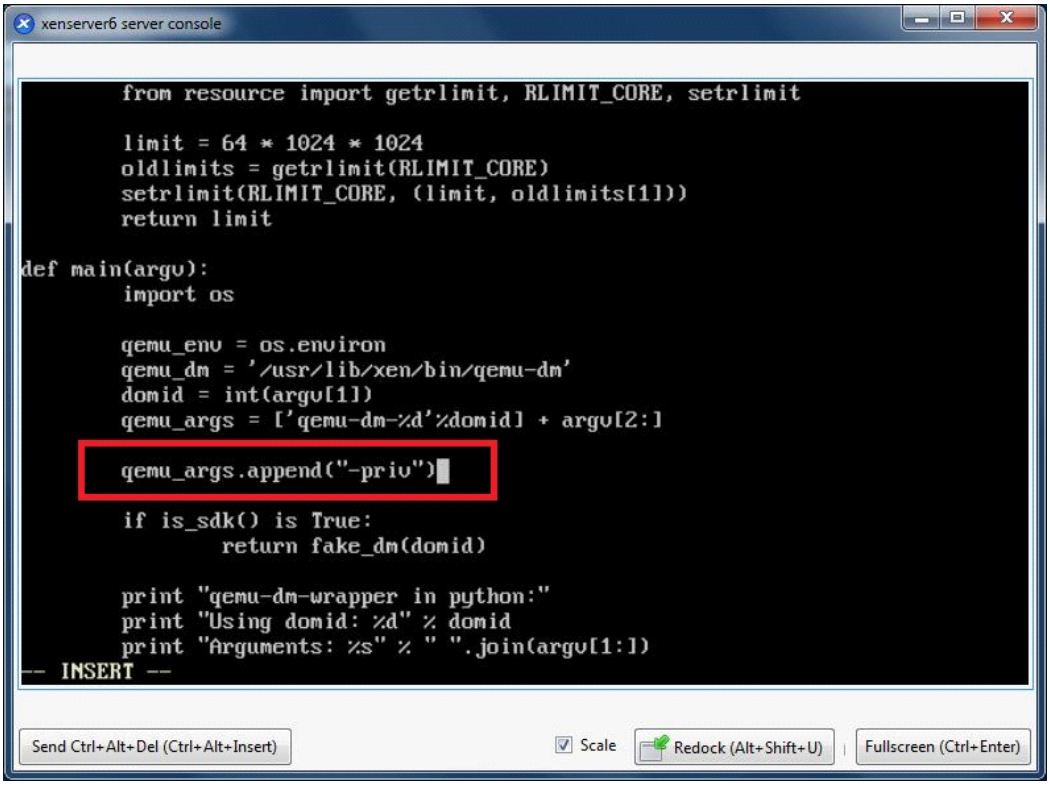

Figure 12. Amending the "def main(argv):" Function

2. Use the XenServer console commands below to list all available Virtual Machines and to assign desired VF to a VM of choice.

```
xl list-vm

xl pci-attach 1 0000:03:10.6
```

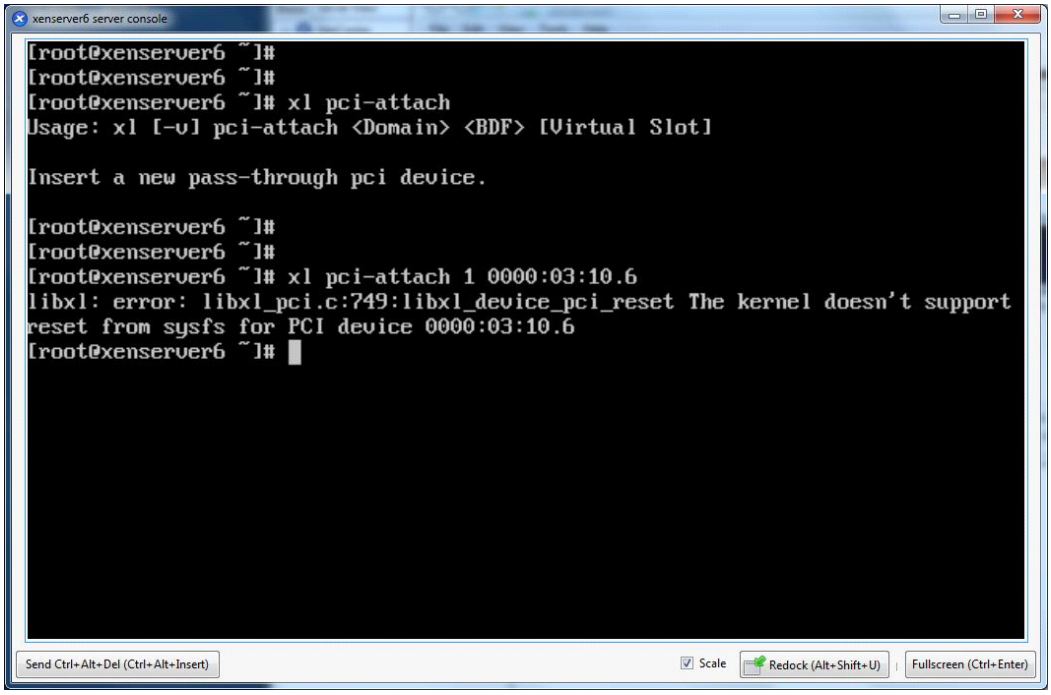

Figure 13. VF (Bus #3, Device #10, and Function #6) Assigned to VM Domain ID #1

Figure 13 shows VF represented by Bus#3, Device#10 and Function#6 assigned to VM domain ID#1. Error message "`libxl: error:`" can be safely ignored. This error doesn't affect VM or VF's functionality.

3. Use the XenServer console command below to verify VF assignment.

```
xl pci-list 1
```

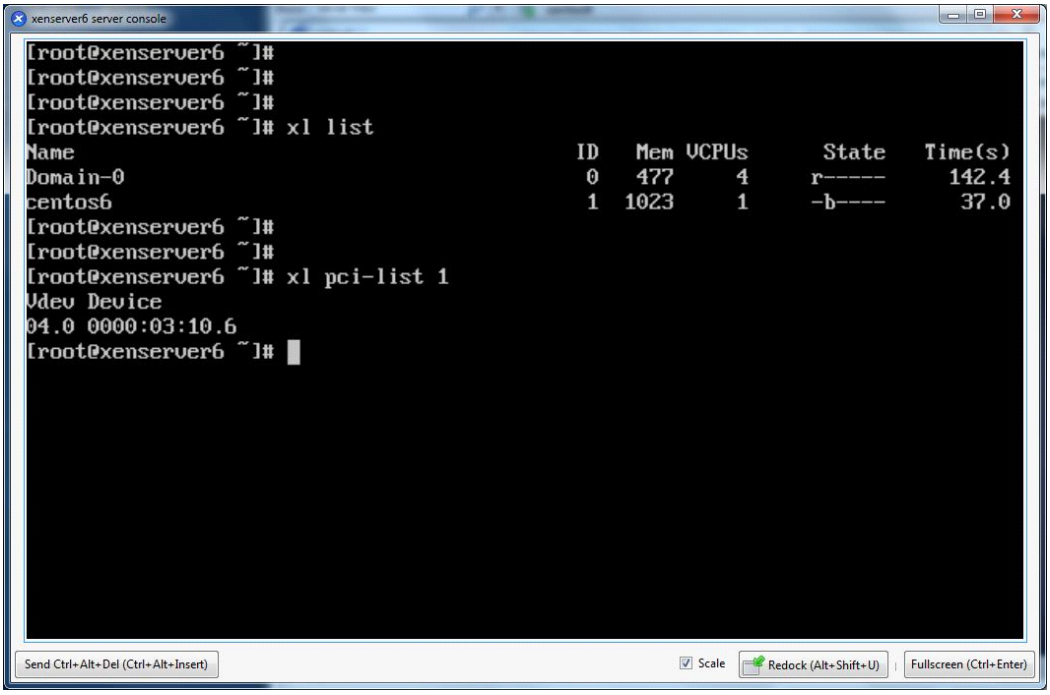

Figure 14. VF Assignment to VM "centos6" with Domain ID #1

3.0 Summary

Intel's best-of-breed 10 GbE solutions are now available with I/O Virtualization capabilities. Customers get world-class Ethernet support along with I/O virtualization support in mainstream Linux distributions--all in a single adapter.

4.0 Customer Support

Intel® Customer Support Services offers a wide selection of programs, including phone support and warranty service. For more information, contact us at:

support.intel.com/support/go/network/ adapter/home.htm

(Service and availability may vary by country.)

5.0 For Product Information

To speak to a customer service representative regarding Intel products, please call 1-800-538-3373 (U.S. and Canada) or visit

support.intel.com/support/go/network/contact.htm

for the telephone number in your area.

§ §