

How to Configure Intel[®] X520 Ethernet Server Adapter Based Virtual Functions on SuSE* Enterprise Linux Server* using Xen*

Technical Brief v1.0

September 2011



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

IMPORTANT - PLEASE READ BEFORE INSTALLING OR USING INTEL® PRE-RELEASE PRODUCTS.

Please review the terms at http://www.intel.com/netcomm/prerelease_terms.htm carefully before using any Intel® pre-release product, including any evaluation, development or reference hardware and/or software product (collectively, "Pre-Release Product"). By using the Pre-Release Product, you indicate your acceptance of these terms, which constitute the agreement (the "Agreement") between you and Intel Corporation ("Intel"). In the event that you do not agree with any of these terms and conditions, do not use or install the Pre-Release Product and promptly return it unused to Intel.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information.

Products may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Hyper-Threading Technology requires a computer system with an Intel® Pentium® 4 processor supporting HT Technology and a HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See http://www.intel.com/products/ht/Hyperthreading_more.htm for additional information.


Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel Leap ahead., Intel Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel ViiV, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2011, Intel Corporation. All Rights Reserved.



Contents

1.0	Introduction	5
1.1	Hardware Requirements	5
1.2	Software Requirements	5
2.0	Installation and Configuration	5
2.1	Enabling IOMMU.....	6
2.2	Intel X520 Ethernet Server Adapter Driver Installation	6
2.3	Hiding Virtual Functions	9
2.4	Virtual Machine Installation and Configuration	11
3.0	Summary	15
4.0	Customer Support	15
5.0	For Product Information	16



Revision History

Date	Revision	Description
September 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 these VFs using Linux* Xen*, which is an integral part of SuSE* Enterprise Linux* Server version 11.

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) specification.
- A server platform with an available PCI Express*: X8, 5.0Gb/s (Gen2) slot

1.2 Software Requirements

- SuSE Enterprise Linux Server Version 11 or newer
- Intel® Ethernet X520 server adapter driver packages. Download the latest available Physical Function and Virtual Function drivers called "ixgbe" and "ixgbev" from <http://sourceforge.net/projects/e1000/files>.

2.0 Installation and Configuration

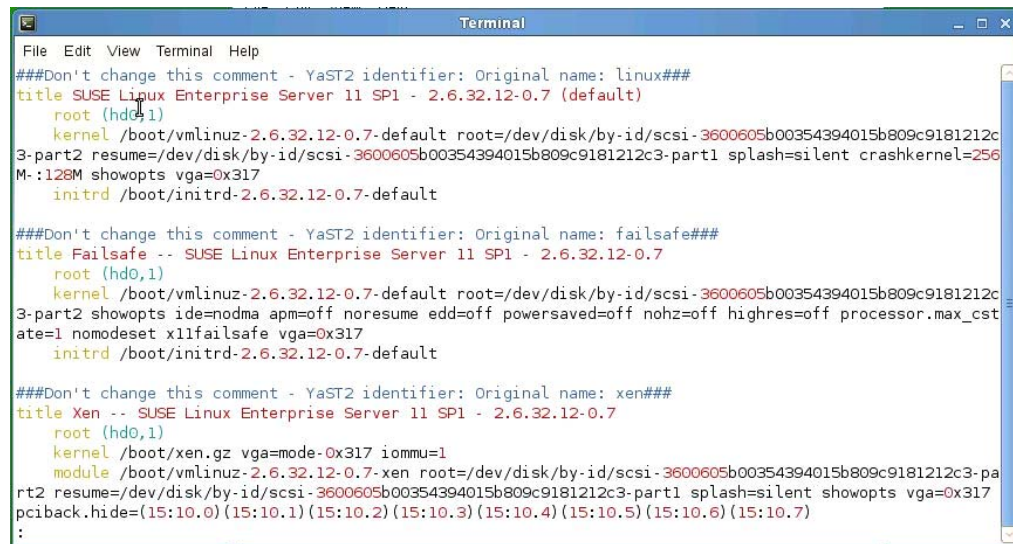
1. Install the Intel® Ethernet X520 server adapter in an available PCI-Express x8 slot. (Ensure that the x8 slot is electrically connected as a x8, some slots are physically x8 but electrically support only x4. Verify this with your server manufacturer or system documentation.)
2. Power up the server.
3. Enter the server's BIOS setup and make sure the virtualization technology and Intel® VT-d features are enabled.
4. Install SuSE Enterprise Linux Server 11 (or newer) on the server.
5. Make sure all Xen modules, libraries, user tools, and utilities have been installed during the operation system installation.
6. The SuSE Enterprise Linux installation process may require a server reboot upon successful operating system install.
7. Log in to the newly-installed SuSE Enterprise Linux operating system using the "root" user account and password.

2.1 Enabling IOMMU

I/O Memory Management Unit (IOMMU) support is not enabled by default in SuSE Enterprise Linux Server 11 distributions. IOMMU support is required for a VF to function properly when assigned to a VM. The following kernel boot parameter are required to enable IOMMU support for Linux kernels:

```
iommu=1
```

This parameter can be appended to the kernel boot entry in `/boot/grub/menu.lst` configuration file. See Figure 1.



```
File Edit View Terminal Help
##Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 SP1 - 2.6.32.12-0.7 (default)
  root (hd0,1)
  kernel /boot/vmlinuz-2.6.32.12-0.7-default root=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part2 resume=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part1 splash=silent crashkernel=256M:128M showopts vga=0x317
  initrd /boot/initrd-2.6.32.12-0.7-default

##Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 SP1 - 2.6.32.12-0.7
  root (hd0,1)
  kernel /boot/vmlinuz-2.6.32.12-0.7-default root=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part2 showopts ide=nodma apm=off noresume eed=off powersaved=off nohz=off highres=off processor.max_cstate=1 nomodeset x11failsafe vga=0x317
  initrd /boot/initrd-2.6.32.12-0.7-default

##Don't change this comment - YaST2 identifier: Original name: xen###
title Xen -- SUSE Linux Enterprise Server 11 SP1 - 2.6.32.12-0.7
  root (hd0,1)
  kernel /boot/xen.gz vga=mode-0x317 iommu=1
  module /boot/vmlinuz-2.6.32.12-0.7-xen root=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part2 resume=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part1 splash=silent showopts vga=0x317
  pciback.hide=(15:10.0)(15:10.1)(15:10.2)(15:10.3)(15:10.4)(15:10.5)(15:10.6)(15:10.7)
  :
```

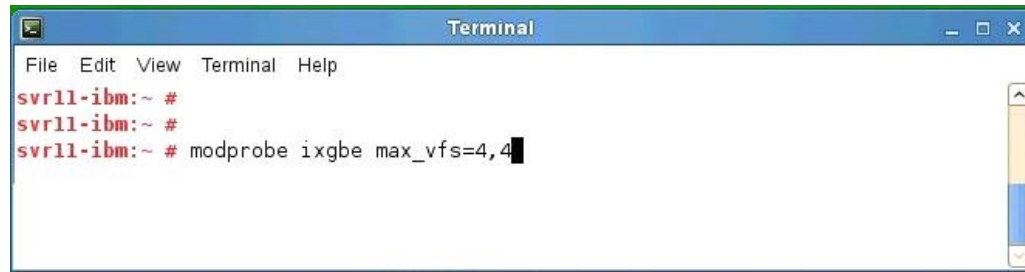
Figure 1. Appending a Parameter

Reboot the server for the `iommu` change to take effect.

2.2 Intel X520 Ethernet Server Adapter Driver Installation

1. PF driver for the X520 server adapter included in SuSE Enterprise Linux Server 11 distributions doesn't support Virtual Functions.
2. Download, uncompress, compile and install `ixgbe` and `ixgbev` drivers.
3. The SuSE Enterprise Linux Server 11 installation does not create VF by default. The X520 server adapter supports up to 63 VFs per port. VFs can be created by using the `ixgbe` driver load time parameter called `max_vfs`. The example below (Figure 2) shows the creation of four VFs per port.

```
#modprobe ixgbe max_vfs=4, 4
```



```
Terminal
File Edit View Terminal Help
svr11-ibm:~ #
svr11-ibm:~ #
svr11-ibm:~ # modprobe ixgbe max_vfs=4,4
```

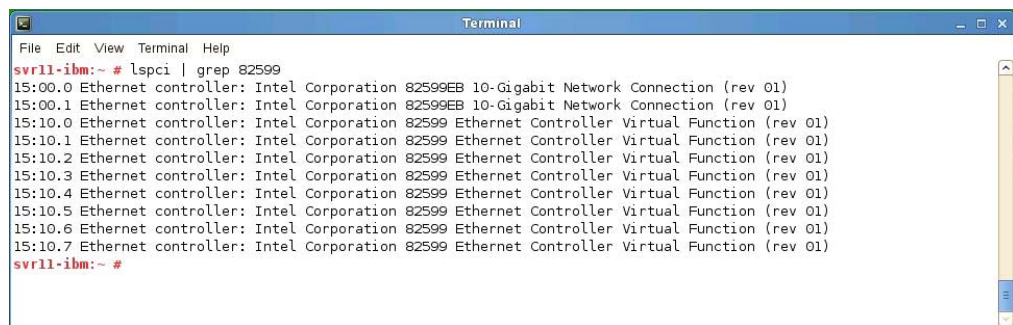
Figure 2. Creating VFs

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

4. Use the `lspci` command to confirm that the VF was successfully created. Example entry:

```
#lspci | grep 82599
```

See the following, Figure 3, for the result of this command.




```
Terminal
File Edit View Terminal Help
svr11-ibm:~ # lspci | grep 82599
15:00.0 Ethernet controller: Intel Corporation 82599EB 10-Gigabit Network Connection (rev 01)
15:00.1 Ethernet controller: Intel Corporation 82599EB 10-Gigabit Network Connection (rev 01)
15:10.0 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
15:10.1 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
15:10.2 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
15:10.3 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
15:10.4 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
15:10.5 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
15:10.6 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
15:10.7 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
svr11-ibm:~ #
```

Figure 3. Result of `#lspci | grep 82599`

Figure 3 shows four VFs per port for the X520 server adapter. Each VF is identified by a unique bus, device, and function number. In the example, the first VF is assigned Bus #15, 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.

Module options are not persistent from one boot to next. SuSE Enterprise Linux Server 11 distribution provides /etc/modprobe.conf.local file for passing parameters to various drivers during the boot process. Add the desired number of VFs that are needed so the VF can be created during server boot time. See Figure 4.



```

Terminal
File Edit View Terminal Help
# Increase the number of loopback devices available for vm creation
options loop max_loop=128
~
~
~
~
~
:wq

```

Figure 6. Setting the Number of Devices to 128

2.3 Hiding Virtual Functions

SuSE Enterprise Linux Server 11 includes the Xen 4.0 open source hypervisor. Xen 4.0 requires SR-IOV-compliant Intel X520 Ethernet Adapter-based Virtual Functions (VFs) not to be claimed by the Xen Domain-0 if VFs are intended to be assigned to Virtual Machines. VF bus, device, and function numbers are required for hiding VFs from Xen Domain-0. See Figure 7.

```
pciback.hide=(15:10.0)(15:10.1) . . .
```

```

Terminal
File Edit View Terminal Help
###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 SP1 - 2.6.32.12-0.7 (default)
root (hd0,1)
kernel /boot/vmlinuz-2.6.32.12-0.7-default root=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part2 resume=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part1 splash=silent crashkernel=256M:128M showopts vga=0x317
initrd /boot/initrd-2.6.32.12-0.7-default

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 SP1 - 2.6.32.12-0.7
root (hd0,1)
kernel /boot/vmlinuz-2.6.32.12-0.7-default root=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part2 showopts ide=nodma apm=off noresume edd=off powersaved=off nohz=off highres=off processor.max_cstate=1 nomodeset x11failsafe vga=0x317
initrd /boot/initrd-2.6.32.12-0.7-default

###Don't change this comment - YaST2 identifier: Original name: xen###
title Xen -- SUSE Linux Enterprise Server 11 SP1 - 2.6.32.12-0.7
root (hd0,1)
kernel /boot/xen.gz vga=mode-0x317 iommu=1
module /boot/vmlinuz-2.6.32.12-0.7-xen root=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part2 resume=/dev/disk/by-id/scsi-3600605b00354394015b809c9181212c3-part1 splash=silent showopts vga=0x317
pciback.hide=(15:10.0)(15:10.1)(15:10.2)(15:10.3)(15:10.4)(15:10.5)(15:10.6)(15:10.7)
:


```

Figure 7. Hiding VFs

Figure 7 shows all eight VFs that were created earlier hidden from Xen Domain-0. These VFs are intended to be assigned to various Virtual Machines later. SuSE Enterprise Linux server will need to be rebooted for these changes to take in effect.

After the server reboots successfully load the pciback driver to claim the VFs for assigning the VFs to Virtual Machines. Use the following command below to load the driver.

```
#modprobe pciback
```



Warnings, errors and informational messages during pciback driver load are logged in the /var/log/messages file. It is good practice to review this file to confirm that the driver loaded successfully without warnings or errors. These messages can also be examined using Linux dmesg utility. Figure 8 below shows dmesg output.

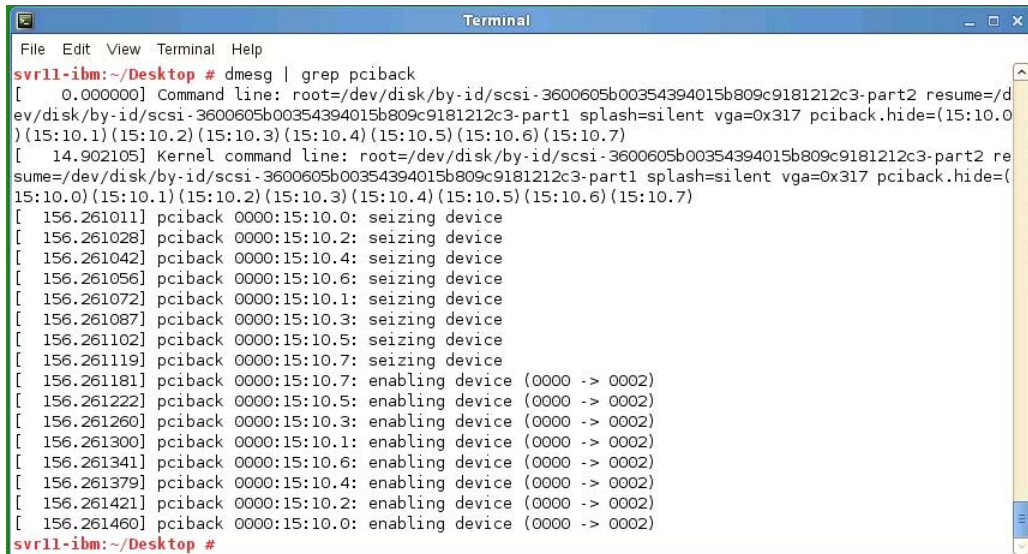


Figure 8. Viewing Messages Using the dmesg Command

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

```
#xm pci-list-assignable-devices
```

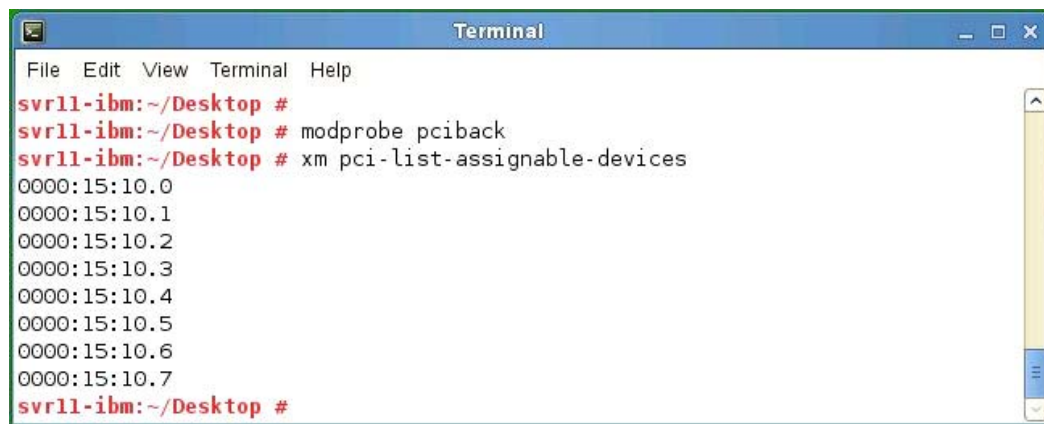


Figure 9. Viewing Available VFs

Figure 9 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

SuSE Enterprise Linux Server 11 includes tools for creating and managing VMs. These tools offer both Command Line (CLI) and Graphical User (GUI) interfaces. Virt-Manager is a GUI tool for creating and managing VMs.

1. Use virt-manager to create a VM.
2. Install the operating system of choice on the newly created VM. For the purpose of this document, Ubuntu 11.04 32-Bit desktop Linux was installed in the VM. See Figure 10 example below:



Figure 10. Creating a VM

3. SuSE Enterprise Linux Server 11 provided Virtual Manager graphical user interface doesn't properly assign PCI devices to Virtual Machine. Use the commands below to list all available Virtual Machines and to assign desired VF to a VM of choice.

```
#xm list
```

```
#xm pci-attach Ubuntu11 15:10.0
```

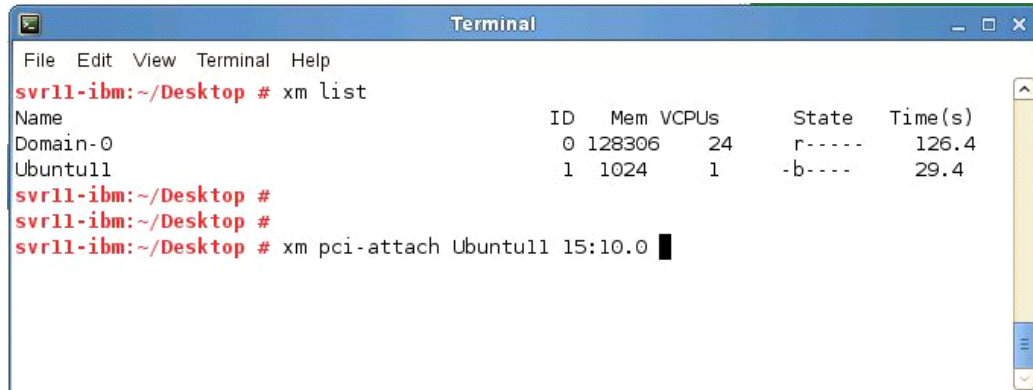


Figure 11. Listing VMs

Figure 11 above shows list of Virtual Machines configured on the server and subsequent VF assignment to Virtual Machine named "Ubuntu11".

4. Click on the icon to edit or examine VM properties.

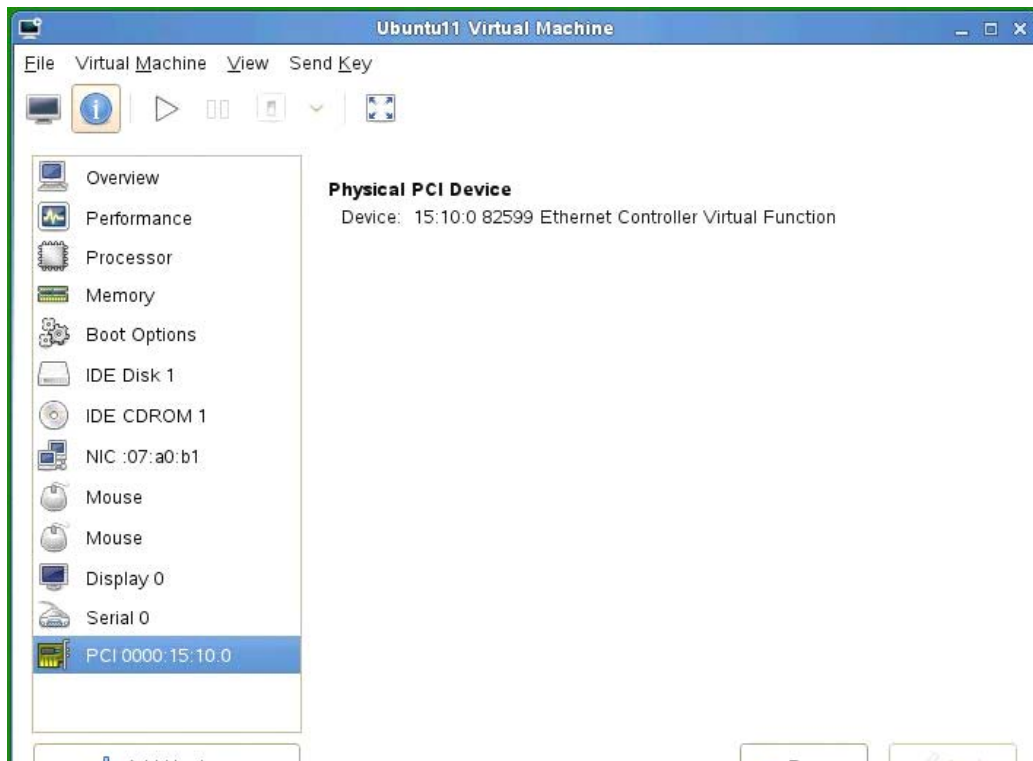



Figure 12. VF Assignment

Figure 12 above shows the Intel X510 Ethernet Server Adapter VF (15:10.0) successfully assigned to the VM.



5. Power up the Ubuntu 11 VM. Log into the VM using the credentials created during the VM installation process.
6. Open the Linux Terminal. Use the Linux `lspci` utility to confirm that the assigned VFis shown within the VM's PCI-Express hierarchy. See Figure 13.




Figure 13. Confirming the VF is Assigned Correctly

7. Use the Linux `lsmod` utility to confirm that the driver for the VF has loaded successfully, as shown in Figure 14.

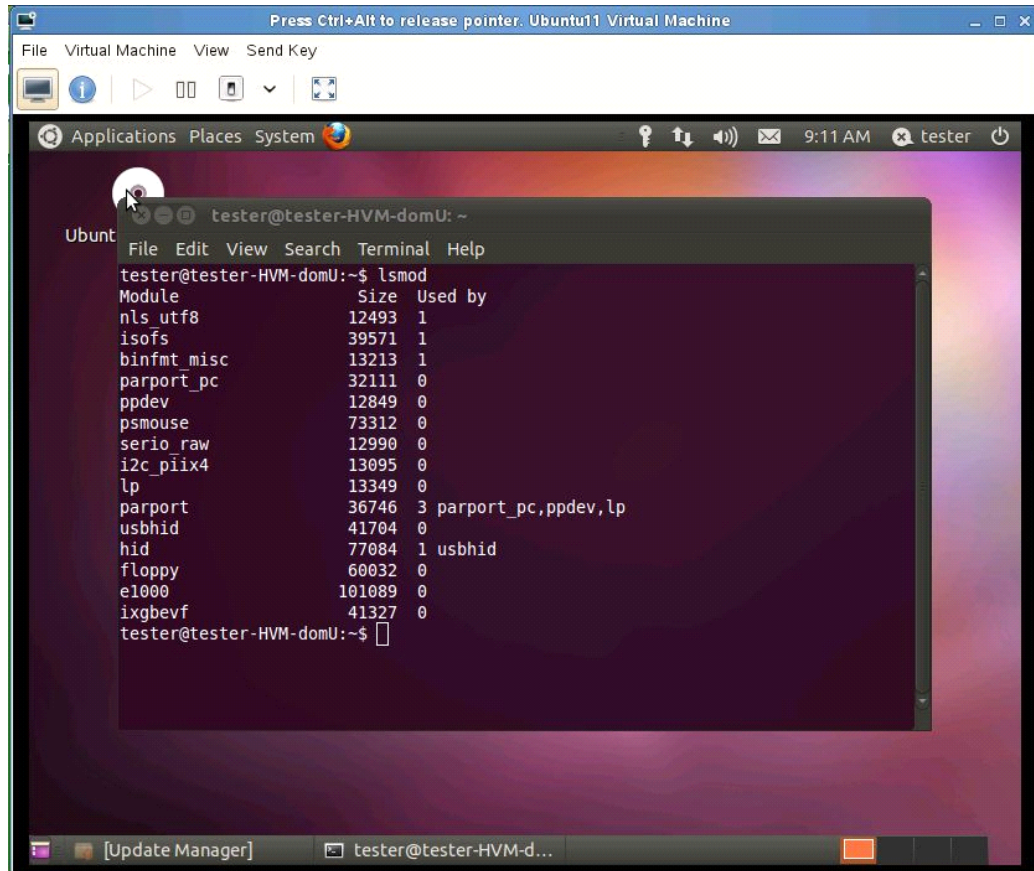



Figure 14. Confirm that the VF has Loaded Correctly

8. Use the Linux ifconfig utility to confirm that the newly assigned VF is ready for use (Figure 15).



```

tester@tester-HVM-domU: ~
File Edit View Search Terminal Help

tester@tester-HVM-domU:~$ ifconfig -a
eth0      Link encap:Ethernet  HWaddr 9a:21:4b:2f:8a:c7
          BROADCAST MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 B)  TX bytes:0 (0.0 B)

eth1      Link encap:Ethernet  HWaddr 00:16:3e:07:a0:b1
          inet addr:10.0.2.57  Bcast:10.0.255.255  Mask:255.255.0.0
          inet6 addr: fe80::216:3eff:fe07:a0b1/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:365577 errors:0 dropped:0 overruns:0 frame:0
          TX packets:65057 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:188514007 (188.5 MB)  TX bytes:5236911 (5.2 MB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
  
```

Figure 15. Confirm that the VF is Ready to Use

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.

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 in a single adapter.

4.0 Customer Support

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

[support.intel.com/support/go/network/ adapter/home.htm](https://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.