Neterion Xframe ESX Drivers

Table of Contents

General Information

Driver Name	s2io.o Device Drivers Download
Supported Adapters	All Neterion Adapters listed here.
Supported ESX versions	ESX and ESXi 3.5 through Update 5; ESX and ESXi 4.0 through Update 1
Hardware Architectures	x86, x86-64

Release Notes and Readme files

For information such as driver requirements, change logs, available downloads, and known issues, see the Release Notes.

For additional information regarding installation, performance suggestions, known issues, etc., please see the Readme.

Measuring Baseline Performance

```
Please e-mail your performance results (along with server platform description) to
Neterion support team at neterionsupport@exar.com. Our support team will then be able to
confirm if the results are what should be expected for a given platform.
```

Download Netperf version 2.4.1

Back to TOC

VMware Driver Utilities

```
To obtain comprehensive MAC-level and driver statistics for the Xframe
adapter run the following command from the console:
ESX# ethtool -S
```

Back to TOC

Frequently Asked Questions (FAQ):

How to install, load and unload the driver?

```
#Intall driver update
   ESX# rpm -Uvh --force VMware-esx-drivers-net-s2io-11904-00000.i386.rpm

#Load the driver
   ESX# vmkload_mod s2io

#Unload the driver
   ESX# vmkload_mod -u s2io

Note: If you have installed one of Neterion's web released driver updates and
you update from ESX 3.5 to ESX 3.5 Update 1, you will need to reinstall the
updated Neterion driver once your ESX update is complete.  The ESX Update over-
writes the driver update, reverting back to the original native version.
```

Useful commands

```
      ESX# vmware -v                          -> ESX Server version
      ESX# vi /proc/vmware/log                -> Log messages
      ESX# vi /var/log/vmkernel               -> kernel log
      ESX# cat /proc/vmware/interrupts        -> Interrups
```

How to add the default loadable parameters?

```
Method 1.
      ESX# esxcfg-module -s "intr_type=2 rx_ring_num=8" s2io
This adds below line to /etc/vmware/esx.conf file
      /vmkernel/module/s2io.o/options = "intr_type=2 rx_ring_num=8"

Method 2.
Edit the /etc/vmware/esx.conf file and append these lines at the end
      /vmkmodule[000X]/enabled = "true"
      /vmkmodule[000X]/module = "s2io.o"
      /vmkmodule[000X]/type = "nic"
      /vmkmodule[000X]/options = "intr_type=2"

where X is the index for each module starting with 0
```

How to establish network connections for virtual machine?

```
A network "port group" needs to be configured to establish network connections
virtual machine(s).

      #Create a new virtual switch
      ESX# esxcfg-vswitch -a vSwitch1

      #Specify a portgroup for operation
      ESX# esxcfg-vswitch --add-pg "VM Network1" vSwitch1

      #Set pnic as an uplink for the vswitch
      ESX# esxcfg-vswitch -L vmnic1 vSwitch1
      ESX# esxcfg-vswitch -L vmnic2 vSwitch1   (To create a team)
      ...
      ESX# esxcfg-vswitch -L vmnicN vSwitch1   (To create a team)

      #Restart the service
      ESX# service mgmt-vmware restart

Select the virtual machine. Move the mouse to the virtual machine name in
the Virtual Machines Tree panel.
Click the right-hand mouse button and select "Edit Settings..."
Select the "Hardware" tab page on the "Virtual Machine Properties" dialog,
click Add button.
Select "Ethernet Adapter" for type of device, click next. Select this vSwitch1
from the Network connection dropdown list.
Make sure the "Connect at power on" device status is checked.

      ESX# esxcfg-vswitch -l

      ##Delete the virtual switch
      ESX# esxcfg-vswitch -d vSwitch1
```

How to enable TSO?

```
TCP Segmentation Offload (TSO) support is added to the TCP/IP stack in ESX
Server 3 version 3.5.
TSO is enabled on the VMkernel interface by default, but must be enabled at
the virtual machine level.

TSO support through the Enhanced vmxnet network adapter is available for virtual
machines running the following guest operating systems:
    - Microsoft Windows 2003 Enterprise Edition with Service Pack 2 (32 bit and 64 bit)
    - Red Hat Enterprise Linux 4 (64 bit)
    - Red Hat Enterprise Linux 5 (32 bit and 64 bit)
    - SuSE Linux Enterprise Server 10 (32 bit and 64 bit)

To enable TSO at the virtual machine level, you must replace the existing vmxnet or
Flexible virtual network adapters with Enhanced vmxnet virtual network adapters.
This may result in a change in the MAC address of the virtual network adapter.

To enable TSO support for a virtual machine
-------------------------------------------
```

```
1 Log in to the VI Client and select the virtual machine from the inventory panel.
The hardware configuration page for this server appears.
2 Click the Summary tab, and click Edit Settings.
3 Select the network adapter from the Hardware list.
4 Record the network settings and MAC address that the network adapter is using.
5 Click Remove to remove the network adapter from the virtual machine.
6 Click Add.
7 Select Ethernet Adapter and click Next.
8 In the Adapter Type group, select Enhanced vmxnet.
9 Select the network setting and MAC address that the old network adapter was
  using and click Next.
10 Click Finish.
11 Click OK.
12 If the virtual machine is not set to upgrade VMware Tools at each power on, you
   must upgrade VMware Tools manually. See the Basic System Administration Guide.

To check that TSO is enabled on a VMkernel interface
----------------------------------------------------
1 Log in directly to your ESX Server 3 hosts console.
2 Use the esxcfg-vmknic -l command to display a list of VMkernel interfaces.
  Each TSO enabled VMkernel interface should appear on the list with TSO MSS set
  to 40960.

TSO is enabled by default on a VMkernel interface. If TSO gets disabled for a particular
VMkernel interface, the only way to enable TSO is to delete that VMkernel interface and
recreate it with TSO enabled.

Caveat: Use TCP workload, since UDP does not use TSO.
```

How to enable MSI-X in the driver?

```
MSI-X can be enabled on platforms which support it, resulting in noticeable
performance improvement. To enable MSI-X in the driver, use the
esxcfg-module -s "intr_type=2" s2io command to enable MSI-X on the s2io module.
```

How to enable Netqueue?

```
1 Log in to the VI Client and select the server from the inventory panel.
  The hardware configuration page for this server appears.
2 Click the Configuration tab, and click Advanced Settings.
3 Select VMkernel.
4 Select VMkernel.Boot.netNetQueueEnable and click OK.
5 Log in directly to your ESX Server 3 hosts console.
6 Use the esxcfg-module -e s2io command to enable the s2io module.
7 Use the esxcfg-module -s "intr_type=2 rx_ring_num=8" s2io command to
  enable NetQueue on the s2io module.
8 Reboot the ESX Server 3 host.
```

How to change the default virtual network adapter emulated inside the guests?

```
The default virtual network adapter emulated inside 32-bit guests is the AMD
PCnet32 device configured with VMware's vlance driver (e1000 for 64-bit guests).
However, vmxnet provides much better performance than vlance and should be
used for optimal performance.
```

```
To use the vmxnet network adapter, install the vmxnet driver (available
in VMware Tools) on your virtual machines.

VMware Tools is a suite of utilities which enhances the performance of a
virtual machine's guest operating system and improves management of the
virtual machine

Step 1.
  1. Log-in to the guest operating-system as root.
  2. On the VI Client control pull-down menu, select:
     [Inventory] -> [Virtual machine] -> [Install VMware tools]
     or on the VM Console pull-down menu, select:
     [VM] -> [Install VMware Tools]
  3. Open a shell terminal window on the VM Console display, then run the
     commands:
     GOS# cd /mnt/cdrom
     GOS# rpm -ivh VMwareTools-e.x.p-xxxxx.i386.rpm
  4. Run the command:
     GOS# /usr/bin/vmware-config-tools.pl

Step 2.
  1. Stop the guest, gracefully.
  2. unregister the vmx file for this VM
     COS:> vmware-cmd -s unregister <full path to vmx file>
  3. Edit the VM VMX configuration file to use vmxnet driver instead e1000 or
     vlance (pcnet32) driver in your Linux guest VMs as follows:

eg. Edit /vmfs/volumes/461d2d02-459f5c38-dfcf-000cfc00140c/RHELAS3U8/RHELAS3U8.vmx
------------------
......
ethernet[n].virtualDev = "vmxnet"
......
------------------
  4. Reregister the vmx file with hostd
     vmware-cmd -s register <full path to vmx file>
  5. Restart the guest
```

How to reserve bandwidth to fifos and rings?

```
  Each fifo and ring can be assigned different guaranteed bandwidth as below.
  1. [root@VMware-ESX]# vmkload_mod s2io rx_ring_num=8 tx_fifo_num=8 bw_percentage=12,24,34
      - In the above example bandwidth are reserved for rings and fifos as below.
        Fifo-0            - 12%
        Fifo-1            - 24%
        Fifo-2            - 34%
        Fifo-3 to Fifo-7 - 6%

        Ring-0            - 12%
        Ring-1            - 24%
        Ring-2            - 34%
        Ring-3 to Ring-7 - 6% each

  2. [root@VMware-ESX]# vmkload_mod s2io rx_ring_num=8 tx_fifo_num=5 bw_percentage=12,24,34
      - In the above example bandwidth are reserved for rings and fifos as below.
```

```
          Fifo-0           - 12%
          Fifo-1           - 24%
          Fifo-2           - 34%
          Fifo-3 to Fifo-4 - 15% each

          Ring-0           - 12%
          Ring-1           - 24%
          Ring-2           - 34%
          Ring-3 to Ring-7 - 6%

    3. [root@VMware-ESX]# vmkload_mod s2io rx_ring_num=5 tx_fifo_num=3 bw_percentage=30,30,40
        - In the above example bandwidth are reserved for rings and fifos as below.
          Fifo-0           - 30%
          Fifo-1           - 30%
          Fifo-2           - 40%

          Ring-0 to Ring-4 - 20% each

    4. [root@VMware-ESX]# vmkload_mod s2io rx_ring_num=5 tx_fifo_num=4
        - In the above example bandwidth are reserved for rings and fifos as below.
          Fifo-0 to Fifo-3 - 25% each

          Ring-0 to Ring-4 - 20% each
```

How to reserve bandwidth on transmit fifos and receive rings with VLAN traffic?

```
       Each fifo and ring can be assigned to a vlan.
      # vmkload_mod s2io bw_percentage=30,40 tx_steering_type=2 rx_steering_type=4 bw_vlan=3,180
tx_fifo_num=8 rx_ring_num=8
       - In the above example
         Fifo-0, Ring-0 to traffic to vlan id 3
         Fifo-1, Ring-1 to traffic to vlan id 180
         30% of the bandwidth is reserved for traffic with vlan id 3
         40% of the bandwidth is reserved for traffic with vlan id 180
         Remaining 6 fifos/rings are used for other traffic (vlan/non vlan)
```

How to Troubleshoot

```
1. "Out of memory in Open" in log message while loading the driver or changing MTU
   Driver load may fail with "Out of memory in Open" failure.  The reason
   for this is the low value of default heap size configured for network
   buffers.  This can be increased with the following steps:
   Edit /etc/vmware/esx.conf and reboot the server:
    /vmkernel/netPktHeapMinSize = "32"
    /vmkernel/netPktHeapMaxSize = "128"

2. Change mtu fails with "Operation not permitted" in VM
   1. Verify that the corresponding virtual switch has the mtu set to equal or more than
      the new mtu we are trying to set for the interface in the VM.
      [root@ESX-3]# esxcfg-vswitch -l ==> view the virtual switches
      [root@ESX-3]# esxcfg-vswitch -m 9000 vSwitch1 ==> chnage the mtu for virtual switch
   2. Verify the Jumbo frames support for the virtual interface
      Look in the configuration file corresponding to the VM
      e.g /vmfs/volumes/461d2d02-459f5c38-dfcf-000cfc00140c/RHELAS3U8/RHELAS3U8.vmx
```

```
---
ethernet[n].features = "15" #This enables TSO as well as Jumbo Frame
---

Please refer to "How to enable TSO?" in the wiki page for more detail.
```

3. Make sure new MTU is set correctly with the ping command from a client

```
  #ping -M do -s 9000 1.1.1.1   ==> Linux
  #ping -f -l 9000 1.1.1.1   ==> Windows
```

Back to TOC

---

*Last updated on **Thu May 26 11:38:36 2011***

Download in other formats:

Plain Text