



## X3100 Series Linux Drivers

### Table of Contents

1. General Information
2. Release Notes and Readme files
3. Measuring Baseline Performance
4. Linux Driver Utilities
5. Frequently Asked Questions (FAQ):
  1. Q1: Which driver version am I using?
  2. Q2: How can I list detailed PCI information about my X3100 adapter(s)?
  3. Q3: How can I tell which physical adapter corresponds to a given ethX name?
  4. Q4: What does it mean when I see the kernel error "swapper: page allocation failure. order:2, mode:0x20"?
  5. Q5: How do I use Receive Side Scaling (aka Receive Traffic Handling)

### General Information

Driver Name	vxge.ko, vxge.o	<a href="#">Device Drivers Download</a>
Supported Adapters	All X3100 Series adapters listed	<a href="#">here</a> .
Supported Linux Kernels	2.6 (including Xen v3.2 and above)	
Shipping in Linux Distros	RHEL 5.5 and above; SLES 11.1 and above	
Hardware Architectures	x86, x86-64, Itanium, PPC 64-bit, SPARC	

### 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](#).

[Back to TOC](#)

### Measuring Baseline Performance

While GbE NICs typically run at GbE line rate on modern systems, 10GbE NICs cause "fast network, slow host" scenario and may or may not run at full 10GbE on a particular system. Typical bottlenecks on a particular system include pci bus speed, system memory bandwidth, and/or cpu utilization. Also, network performance is very application-specific. For suggestions to improve performance, please see the [Linux Readme](#)

To measure baseline TCP throughput for your Neterion 10GbE Adapter, please connect two similar servers with Neterion NICs back-to-back, apply performance tuning tips from the [Linux Readme](#) and run the following commands:

- **iPerf**

System 1

```
#./iperf -s
```

System 2

```
#./iperf -c <remotehost> -l89600 -t60 -i5 -P4
```

Download iPerf version 2.0.2

**Latency can be measured with tools like Netperf or Netpipe.**

- **Netperf**

```
#netperf -H <remotehost> -t TCP_RR
```

Download Netperf version 2.4.1

Please e-mail your performance results (along with server platform description) to [neterionsupport@exar.com](mailto:neterionsupport@exar.com). Our support team will be able to confirm whether your results are expected for a given platform.

[Back to TOC](#)

---

## Linux Driver Utilities

To obtain adapter statistics for X3100 adapter assigned eth2:

```
# ethtool -S eth2
```

To obtain adapter register dump for X3100 adapter assigned eth2:

```
# ethtool -d eth2
```

[Back to TOC](#)

---

## Frequently Asked Questions (FAQ):

### Q1: Which driver version am I using?

```
A: Use the command 'ethtool -i ethX'  
# ethtool -i eth2  
driver: vxge  
version: 1.0.0  
firmware-version: 0.35.0  
bus-info: 0000:19:00.0
```

If you don't want to load the kernel module but want to know the version, `grep` for version in modinfo with the following command:

```
# modinfo [module.ko]|grep ^version:  
(Result will be the string reported to MODULE_VERSION()when driver was compiled)
```

### Q2: How can I list detailed PCI information about my X3100 adapter(s)?

---

A: Use the command 'lspci -d 17d5:\* -v'

```
# lspci -d 17d5:* -v
19:00.0 Ethernet controller: S2io Inc. Unknown device 5833 (rev 02)
  Subsystem: S2io Inc. Unknown device 6030
  Flags: bus master, fast devsel, latency 0, IRQ 209
  Memory at f7000000 (64-bit, prefetchable) [size=8M]
  Memory at f6ff0000 (64-bit, prefetchable) [size=4K]
  Memory at f6fe0000 (64-bit, prefetchable) [size=256]
  Capabilities: [40] Power Management version 3
  Capabilities: [70] Express Endpoint IRQ 8
  Capabilities: [a0] MSI-X: Enable+ Mask- TabSize=68
  Capabilities: [c0] Vital Product Data
  Capabilities: [50] Message Signalled Interrupts: Mask+ 64bit+ Queue=0/3 Enable-
  Capabilities: [100] Unknown (14)
  Capabilities: [110] Advanced Error Reporting
  Capabilities: [150] Power Budgeting
  Capabilities: [160] Device Serial Number 00-0c-fc-00-00-00-b0-33
```

### Q3: How can I tell which physical adapter corresponds to a given ethX name?

A: Use the command 'ethtool -p ethX'

```
# ethtool -p eth2
(the LED of the physical adapter assigned interface eth2 will blink)
```

### Q4: What does it mean when I see the kernel error "swapper: page allocation failure. order:2, mode:0x20"?

A. This message is due to a memory allocation failure. If you are using Jumbo Frames larger than 4000 bytes, it is likely a memory fragmentation problem and there probably aren't enough contiguous pages available in memory to allocate.

The workaround is to use a MTU size of 4000 bytes or less. Since handling of a frame this size requires allocation of no more than a single page, memory fragmentation will no longer be an issue. Using LRO in conjunction with this reduced MTU size will help to keep performance at par with 9000 byte MTU performance.

```
# insmod vxge.ko lro=1
# ifconfig eth# <ip_address> mtu 4000
```

### Q5: How do I use Receive Side Scaling (aka Receive Traffic Handling)

A: To disable/enable RSS, the system must have a newer version of ethtool (version 2.6.34 or newer)

RSS is enabled by default (see receive-hashing below)

```
# modprobe vxge
# dmesg
vxge: Copyright(c) 2002-2010 Exar Corp.
vxge: Driver version: 2.0.10.20927-k
vxge 0000:09:00.0: PCI INT A -> GSI 18 (level, low) -> IRQ 18
vxge 0000:09:00.0: setting latency timer to 64
eth4: SERIAL NUMBER: SXJ0843014
eth4: PART NUMBER: X3110SR0003
eth4: Neterion X3110 Single-Port SR 10GbE Server Adapter
eth4: MAC ADDR: 00:0c:fc:00:b8:4d
```

```
eth4: Link Width x8
eth4: Firmware version : 1.7.1 Date : 07/07/2010
eth4: Single Function Mode Enabled
eth4: 2 Vpath(s) opened
eth4: Interrupt type MSI-X
eth4: RTH steering enabled for TCP_IPV4
eth4: Tx multiqueue steering enabled
eth4: Generic receive offload enabled
eth4: Rx doorbell mode enabled
eth4: VLAN tag stripping Enabled
eth4: Ring blocks : 2
eth4: Fifo blocks : 14
# ethtool -k eth4
Offload parameters for eth4:
rx-checksumming: on
tx-checksumming: on
scatter-gather: on
tcp-segmentation-offload: on
udp-fragmentation-offload: off
generic-segmentation-offload: on
generic-receive-offload: on
large-receive-offload: off
ntuple-filters: off
receive-hashing: on

To disable RSS:
# ethtool -K eth4 rxhash off
# ethtool -k eth4
Offload parameters for eth4:
rx-checksumming: on
tx-checksumming: on
scatter-gather: on
tcp-segmentation-offload: on
udp-fragmentation-offload: off
generic-segmentation-offload: on
generic-receive-offload: on
large-receive-offload: off
ntuple-filters: off
receive-hashing: off

To enable RSS:
# ethtool -K eth4 rxhash on
# ethtool -k eth4
Offload parameters for eth4:
rx-checksumming: on
tx-checksumming: on
scatter-gather: on
tcp-segmentation-offload: on
udp-fragmentation-offload: off
generic-segmentation-offload: on
generic-receive-offload: on
large-receive-offload: off
ntuple-filters: off
receive-hashing: on
```

[Back to TOC](#)

---

**Download in other formats:**

Plain Text



Powered by **Trac 0.8**  
By Edgewall Software.

Visit the Trac open source project at  
<http://trac.edgewall.com/>



## X3100 Series Linux Drivers

### Table of Contents

1. General Information
2. Release Notes and Readme files
3. Measuring Baseline Performance
4. Linux Driver Utilities
5. Frequently Asked Questions (FAQ):
  1. Q1: Which driver version am I using?
  2. Q2: How can I list detailed PCI information about my X3100 adapter(s)?
  3. Q3: How can I tell which physical adapter corresponds to a given ethX name?
  4. Q4: What does it mean when I see the kernel error "swapper: page allocation failure. order:2, mode:0x20"?
  5. Q5: How do I use Receive Side Scaling (aka Receive Traffic Handling)

### General Information

Driver Name	vxge.ko, vxge.o	<a href="#">Device Drivers Download</a>
Supported Adapters	All X3100 Series adapters listed	<a href="#">here</a> .
Supported Linux Kernels	2.6 (including Xen v3.2 and above)	
Shipping in Linux Distros	RHEL 5.5 and above; SLES 11.1 and above	
Hardware Architectures	x86, x86-64, Itanium, PPC 64-bit, SPARC	

### 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](#).

[Back to TOC](#)

### Measuring Baseline Performance

While GbE NICs typically run at GbE line rate on modern systems, 10GbE NICs cause "fast network, slow host" scenario and may or may not run at full 10GbE on a particular system. Typical bottlenecks on a particular system include pci bus speed, system memory bandwidth, and/or cpu utilization. Also, network performance is very application-specific. For suggestions to improve performance, please see the [Linux Readme](#)

To measure baseline TCP throughput for your Neterion 10GbE Adapter, please connect two similar servers with Neterion NICs back-to-back, apply performance tuning tips from the [Linux Readme](#) and run the following commands:

- **iPerf**

System 1

```
#./iperf -s
```

System 2

```
#./iperf -c <remotehost> -l89600 -t60 -i5 -P4
```

Download iPerf version 2.0.2

**Latency can be measured with tools like Netperf or Netpipe.**

- **Netperf**

```
#netperf -H <remotehost> -t TCP_RR
```

Download Netperf version 2.4.1

Please e-mail your performance results (along with server platform description) to [neterionsupport@exar.com](mailto:neterionsupport@exar.com). Our support team will be able to confirm whether your results are expected for a given platform.

[Back to TOC](#)

---

## Linux Driver Utilities

To obtain adapter statistics for X3100 adapter assigned eth2:

```
# ethtool -S eth2
```

To obtain adapter register dump for X3100 adapter assigned eth2:

```
# ethtool -d eth2
```

[Back to TOC](#)

---

## Frequently Asked Questions (FAQ):

### Q1: Which driver version am I using?

```
A: Use the command 'ethtool -i ethX'  
# ethtool -i eth2  
driver: vxge  
version: 1.0.0  
firmware-version: 0.35.0  
bus-info: 0000:19:00.0
```

If you don't want to load the kernel module but want to know the version, `grep` for version in modinfo with the following command:

```
# modinfo [module.ko]|grep ^version:  
(Result will be the string reported to MODULE_VERSION()when driver was compiled)
```

### Q2: How can I list detailed PCI information about my X3100 adapter(s)?

---

A: Use the command 'lspci -d 17d5:\* -v'

```
# lspci -d 17d5:* -v
19:00.0 Ethernet controller: S2io Inc. Unknown device 5833 (rev 02)
  Subsystem: S2io Inc. Unknown device 6030
  Flags: bus master, fast devsel, latency 0, IRQ 209
  Memory at f7000000 (64-bit, prefetchable) [size=8M]
  Memory at f6ff0000 (64-bit, prefetchable) [size=4K]
  Memory at f6fe0000 (64-bit, prefetchable) [size=256]
  Capabilities: [40] Power Management version 3
  Capabilities: [70] Express Endpoint IRQ 8
  Capabilities: [a0] MSI-X: Enable+ Mask- TabSize=68
  Capabilities: [c0] Vital Product Data
  Capabilities: [50] Message Signalled Interrupts: Mask+ 64bit+ Queue=0/3 Enable-
  Capabilities: [100] Unknown (14)
  Capabilities: [110] Advanced Error Reporting
  Capabilities: [150] Power Budgeting
  Capabilities: [160] Device Serial Number 00-0c-fc-00-00-00-b0-33
```

### Q3: How can I tell which physical adapter corresponds to a given ethX name?

A: Use the command 'ethtool -p ethX'

```
# ethtool -p eth2
(the LED of the physical adapter assigned interface eth2 will blink)
```

### Q4: What does it mean when I see the kernel error "swapper: page allocation failure. order:2, mode:0x20"?

A. This message is due to a memory allocation failure. If you are using Jumbo Frames larger than 4000 bytes, it is likely a memory fragmentation problem and there probably aren't enough contiguous pages available in memory to allocate.

The workaround is to use a MTU size of 4000 bytes or less. Since handling of a frame this size requires allocation of no more than a single page, memory fragmentation will no longer be an issue. Using LRO in conjunction with this reduced MTU size will help to keep performance at par with 9000 byte MTU performance.

```
# insmod vxge.ko lro=1
# ifconfig eth# <ip_address> mtu 4000
```

### Q5: How do I use Receive Side Scaling (aka Receive Traffic Handling)

A: To disable/enable RSS, the system must have a newer version of ethtool (version 2.6.34 or newer)

RSS is enabled by default (see receive-hashing below)

```
# modprobe vxge
# dmesg
vxge: Copyright(c) 2002-2010 Exar Corp.
vxge: Driver version: 2.0.10.20927-k
vxge 0000:09:00.0: PCI INT A -> GSI 18 (level, low) -> IRQ 18
vxge 0000:09:00.0: setting latency timer to 64
eth4: SERIAL NUMBER: SXJ0843014
eth4: PART NUMBER: X3110SR0003
eth4: Neterion X3110 Single-Port SR 10GbE Server Adapter
eth4: MAC ADDR: 00:0c:fc:00:b8:4d
```



```
eth4: Link Width x8
eth4: Firmware version : 1.7.1 Date : 07/07/2010
eth4: Single Function Mode Enabled
eth4: 2 Vpath(s) opened
eth4: Interrupt type MSI-X
eth4: RTH steering enabled for TCP_IPV4
eth4: Tx multiqueue steering enabled
eth4: Generic receive offload enabled
eth4: Rx doorbell mode enabled
eth4: VLAN tag stripping Enabled
eth4: Ring blocks : 2
eth4: Fifo blocks : 14
# ethtool -k eth4
Offload parameters for eth4:
rx-checksumming: on
tx-checksumming: on
scatter-gather: on
tcp-segmentation-offload: on
udp-fragmentation-offload: off
generic-segmentation-offload: on
generic-receive-offload: on
large-receive-offload: off
ntuple-filters: off
receive-hashing: on

To disable RSS:
# ethtool -K eth4 rxhash off
# ethtool -k eth4
Offload parameters for eth4:
rx-checksumming: on
tx-checksumming: on
scatter-gather: on
tcp-segmentation-offload: on
udp-fragmentation-offload: off
generic-segmentation-offload: on
generic-receive-offload: on
large-receive-offload: off
ntuple-filters: off
receive-hashing: off

To enable RSS:
# ethtool -K eth4 rxhash on
# ethtool -k eth4
Offload parameters for eth4:
rx-checksumming: on
tx-checksumming: on
scatter-gather: on
tcp-segmentation-offload: on
udp-fragmentation-offload: off
generic-segmentation-offload: on
generic-receive-offload: on
large-receive-offload: off
ntuple-filters: off
receive-hashing: on
```

[Back to TOC](#)

---

**Download in other formats:**

Plain Text



Powered by **Trac 0.8**  
By Edgewall Software.

Visit the Trac open source project at  
<http://trac.edgewall.com/>