

How do I isolate the source of packet loss when using the Sniffer10G driver?

Model:

10G-PCIE2-8CS-2S

Software:

Sniffer10G

Operating System:

Supports both Linux and Windows operating systems.

Information:

The following suggestions will help to isolate the source of packet loss when using the Sniffer10G driver.

If you need assistance, please send the answers to the following **Checklist** questions and the following output:

Linux:

Please send us the output of **/opt/snf/sbin/myri_bug_report**.

This script must be run as root and it is very important that it is executed from within the **/sbin directory**, otherwise important diagnostic information will not be collected.

Windows:

Please send us the output of **[INSTALL_DIR] \bin\myri_counters.exe**.

If there are any error messages in the event viewer (Start-Run, enter “eventvwr”), please also send us those messages.

If you are running Windows 2009 or later and powershell is installed, please send us the output of **[INSTALL_DIR] \sbin\myri_dmesg.ps1**.

This script must be run as administrator and the execution of the script requires the following setting:

Set-executionpolicy remotesigned

It is very important that this script is executed from within the **\sbin directory**, otherwise important diagnostic information will not be collected.

If you cannot use powershell, inspect the **Event Viewer-> Windows Logs** for error messages related to the Sniffer10G

software.

Checklist:

1. Which OS are you using?
2. Are the Myri-10G network adapters connected to a 10GbE switch? Is flow control disabled on the switch and on the Myri-10G network adapters? By default, link level flow control is disabled on the adapters (via the Sniffer10G driver). See the Sniffer10G README for details.
By default, Ethernet flow control is disabled in Sniffer10G. We assume that all users will want to passively capture packets. There is never a good reason to enable flow control with Sniffer10G since Sniffer10G will try to sink the traffic as best effort.
3. With which tool are you measuring packet rate (tcpdump, wireshark, dumpcap), and how is the tool generating output (to stdout, to the screen, writing to a disk file)?
Please be aware that while Sniffer10G can support up to 14.8Mpps, tcpdump will print output to standard out, and wireshark scrolls output to the screen, which will severely limit achievable packet rates.
4. Do your switch counters report packet loss?
5. How are you measuring the packet loss? With which benchmark? What type of traffic? Unicast? Multicast? TCP? UDP? Does the application use multiple connections?
6. Collect the Sniffer10G **myri_counters**, **myri_dmabench**, and **myri_info** output on the receiver at the time when packet loss is being experienced.
(On **Linux**, these tools can be found in `/opt/snf/sbin/` and `/opt/snf/bin/`.)
(On **Windows**, these tools can be found in `[INSTALL_DIR]\SBIN\` and `[INSTALL_DIR]\bin\`.)

Examine the Sniffer10G myri_info output for the following information:

```
# /opt/snf/sbin/myri_info
```

Verify that our x8 PCIe network adapter is installed into a PCIe slot running at x8 speed (link_width:8).

If you are using a two-port Myri-10G network adapter (10G-PCIE2-8C2-2* or 10G-PCIE2-8B2-2*) ensure that the adapter is installed into a Gen2 (and not Gen1) x8 PCIe slot.

If this is not the case, carefully read the technical specifications for the motherboard to determine if there is a PCIe slot on the motherboard that will run at x8 PCIe speed if a x8 PCIe adapter is installed.

Examine the Sniffer10G myri_counters output for the following information:

```
$ /opt/snf/bin/myri_counters
```

or for a two-port adapter:

```
$ /opt/snf/bin/myri_counters -p 0  
$ /opt/snf/bin/myri_counters -p 1
```

Is the value of **Ethernet recv overrun** non-zero? A non-zero value for this counter indicates that the OS is not keeping up

(due to buffer sizes) and drop packets.

Are the values of **SNF drop ring full**, **SNF drop desc credits**, and **SNF drop data credits** non-zero? If these counters are non-zero, then the Sniffer10G application (or the application using libpcap) does not consume the packets fast enough, so packets are dropped because the host ring is full. Examine the output when running the application with **SNF_DEBUG_MASK=0x3**. What are the values of the environment variables **SNF_NUM_RINGS**, **SNF_DATARING_SIZE**, and **SNF_DESCRING_SIZE**?

Is the value of **Net bad PHY** non-zero? If this counter is non-zero, packets are dropped due to bad PHY check or bad Ethernet CRC. If the number is high, the link is bad. Check for a damaged cable, transceiver module, or adapter connector.

Is the value of **Net overflow drop** non-zero?

For Sniffer10G, the value of **Net overflow drop** should always be zero.

If this counter has a non-zero value, did you verify that **myri_dmabench** reports sufficient bandwidth for 10Gb/s traffic?

If this counter has a non-zero value, you should also verify that CPU Frequency Scaling is disabled for the processor.

If this counter has a non-zero value, are you using a PCIe Gen3 motherboard? If yes, do you also see non-zeropcie_* counters in the output of **/sbin/myri_info -vvv**? If yes, check if the BIOS on this PCIe Gen3 motherboard has a Gen2 setting, and switch to that setting.

Draft	Date	Change
1	7/13/2016	Initial Draft
2	7/25/2016	Feedback

