

What are some of the essential settings to check to improve performance on Linux Myri10GE/DBL/Sniffer10G?

Model:

N/A

Software:

Myri10gE/DBL/Sniffer10G

Operating System:

Supports the Linux Operating System.

Information:

Basic Checks

Check that the adapter is installed into a PCI-Express slot that negotiates the correct PCI-Express link width (x8).

Run the **Myri_info** command to verify that the negotiated PCI-Express link is the correct x8 width.

For single port adapters, the speed can be inferred by running ‘**ethtool -S**’ and checking the **read_dma_bw_MB**s and **write_dma_bw_MB**s and verify the dma bandwidth is consistent with x8 pci express link.

Check if the firewall is enabled, and disable it if necessary.

System-config-firewall can be used to disable it.

“**/etc/init.d/iptables status**” and “**/etc/init.d/ip6tables status**” can be used to check if it is disabled.

Check if SELinux is enabled, and disable it if necessary. Auditd may also need to be disabled with: “**/etc/init.d/auditd stop**”.

Check how IRQs are being balanced. In some situations, performance can be impacted negatively by interrupt balancing. Disabling irq balancing (on RHEL it is the daemon **irqbalance**), and manually assigning the irq to a specific CPU and then binding the tasks to a specific CPU (via **taskset** or **numactl**) can improve performance.

Disable CPU Frequency Scaling on the processor.

Myri10GE specific tuning

Make sure LRO is turned on (some virtualization support in RHEL6 can turn it off). You can enable it through **ethtool** if it is disabled.

To remove virtualization: Run “**rpm -q -a | grep virt**” then remove the virtualization related packages. A reboot will be necessary.

If virtualization is required, and LRO can’t be used, turning on GRO will provide some performance improvement and is compatible with virtualization.

Network buffer sizes:

For best performance with Myri10GE or Ethernet Emulation with MX-10G, we recommend that several network buffer sizes be increased from their defaults. Please add the following lines to **/etc/sysctl.conf**:

```
net.core.rmem_max = 16777216  
net.core.wmem_max = 16777216
```

```
net.ipv4.tcp_rmem = 4096 87380 16777216
net.ipv4.tcp_wmem = 4096 65536 16777216
net.core.netdev_max_backlog = 250000
```

For **RHEL6** you may also want to add the following:

```
net.ipv4.tcp_congestion_control = bic' to sysctl.conf
```

Finally, run:

```
sysctl -p /etc/sysctl.conf
```

DBL and Sniffer10G specific tuning

Dbl and Sniffer10G are kernel-bypass drivers. As such, DBL and Sniffer10G do not require much tuning, but it helps to be aware of the following items that affect latency/packet rate.

CPU Binding:

CPU binding can be very useful (e.g., using taskset or numactl on Linux) for optimizing performance for DBL and Sniffer10G, as well as Myri10GE.

On hosts with multiple CPU sockets, some CPUs are physically closer to the adapter and/or memory and perform better than others.

Some systems have multiple PCI-Express root ports, for example AMD machines, and some PCI-Express slots can connect to one root port and are more easily accessible from one CPU socket.

PCI Bridging

Some PCI-Express slots on a machine may have deeper PCI bridging than others. Extra bridge chips between the CPU and the adapter will result in higher latency. If there seem to be lots of bridge chips, trying a different PCI-Express slot may improve latency. (The tool **Myri_info** can be used to see the bridges between the CPU and adapter).

Hyper threading:

We do not recommend using more than one SNF receive thread (or process) on each physical core. Hyper threading can be enabled as long as the processes/threads are mapped correctly.

Reducing Jitter:

Consider reducing jitter by only allowing your SNF/DBL application on a specific core and excluding any other processes from that core.

<u>Revision</u>	<u>Date</u>	<u>Change</u>
1	8/3/2016	Initial Draft