



10 Gigabit Ethernet: current state


Internet2 Performance Workshop

6-Dec-2005 v0.1

Caveat: we personally have limited experience with 10GE servers (just introduced in HOPI)

Reporting current “best of breed”

- Some general principles
- Land speed record candidates
- What High-Energy Physics is doing



10GE: setting expectations

~833,400 pps at 1500 MTU

~138,900 pps at 9000 MTU

Bus limitations

- PCI-X 1.0 is ~8Gbps
- PCI-X 2.0 (PCI Express?) should help, but chipsets and interface cards are just now becoming available



Three Card Vendors

Intel

Neterion (nee S2IO)

Chelsio

You want dual AMD 64 bit processors
PCI-X 1.0 at 133 Mhz has achieved
7.5Gbps TCP on Linux (Neterion)

PCI-X 1.0 at 100 Mhz has achieved 5.0
Gbps

UDP: Esnet had some problems;
Harvey's group at CALTECH has
achieved 6.5 Gbps



DREN [October 2004]

1st gen Intel 10GE NICs

Dell 2650

2.4 GHz Xeon, 400Mhz FSB

3.8Gbps TCP



Caltech [October 2004]

Use 64 bit architecture

S2IO XFrame (now Neterion)

- 6.5Gbps

1st Gen Intel: 5.8 Gbps

Opteron slightly more performant than
Itanium

Opteron

- Tyan Thunder K8S Pro (S2882), 2x Opteron 2.2GHz, 2GB mem

Itanium

- HP rx 4640, 4x1.5GHz Itanium-2, zx1 chipset, 8GB memory



ESnet [April 2005]

Tyan S2895A2NRF IDE Version

“thunderk8we”

Dual AMD 252 Opteron CPUs

2GB DDR 400 ECC/Reg.

Neterion X-Frame PCI-X NIC

Linux (2.6?)

7.5Gbps TCP



NCAR [September 2005]

“Dual Opteron, PCI-X”

Chelsio N210 cards

6.5Gbps

MTU 8160, a bunch of other tuning
parms

“it just worked”



MAX [September 2005]

Supermicro X6DHE-XG2, Dual Xeon
800MHz FSB motherboard

(2) Intel Xeon64 3.2GHz, 1MB Cache,
800MHz FSB Processors

Chelsio S210 card



Linux tuning

Talk to your NIC vendor

MTU 9000 (8160?)

Iptables off (and unloaded)

Actual tweaks vary by version (and have changed a lot recently)



A running list

<http://e2epi.internet2.edu/net-perf-wkshp/binder-docs/10GEexamples.html>



INTERNET®

www.internet2.edu