


Low Latency COTS GigE

62:	8192 bytes	545 times -->	4576.97 Mbps in	229.10 usec
63:	8195 bytes	545 times -->	4410.98 Mbps in	237.81 usec
64:	12285 bytes	525 times -->	4808.78 Mbps in	327.00 usec
65:	12288 bytes	509 times -->	4949.95 Mbps in	317.75 usec
66:	12291 bytes	524 times -->	4927.55 Mbps in	319.28 usec
67:	16381 bytes	261 times -->	4966.26 Mbps in	422.20 usec
68:	16384 bytes	296 times -->	5035.23 Mbps in	416.50 usec
69:	16387 bytes	300 times -->	4993.62 Mbps in	420.04 usec
70:	24573 bytes	297 times -->	5110.70 Mbps in	615.44 usec
71:	24576 bytes	270 times -->	5165.47 Mbps in	608.99 usec
72:	24579 bytes	273 times -->	5143.19 Mbps in	611.70 usec
73:	32765 bytes	136 times -->	5094.44 Mbps in	823.24 usec
74:	32768 bytes	151 times -->	5526.62 Mbps in	758.93 usec
75:	32771 bytes	164 times -->	5042.57 Mbps in	831.85 usec
76:	49149 bytes	150 times -->	4975.27 Mbps in	1264.47 usec
77:	49152 bytes	131 times -->	4832.97 Mbps in	1301.78 usec
78:	49155 bytes	128 times -->	5263.84 Mbps in	1195.29 usec
79:	65533 bytes	69 times -->	4919.98 Mbps in	1704.93 usec
80:	65536 bytes	73 times -->	4730.13 Mbps in	1773.44 usec
81:	65539 bytes	70 times -->	4926.60 Mbps in	1702.79 usec
82:	98301 bytes	73 times -->	4559.99 Mbps in	2759.33 usec
83:	98304 bytes	60 times -->	4800.65 Mbps in	2621.09 usec
84:	98307 bytes	63 times -->	4964.35 Mbps in	2534.73 usec
85:	133260 bytes	33 times -->	4964.35 Mbps in	2534.73 usec

Lessons Learned : IP limitations are difficult to code around. It is easier to build a cluster-specific solution.



Limits; NIC (node) count, NIC Bandwidth, Bus Bandwidth

Low Latency COTS GigE



Scaling an 8 node cluster : You get incremental performance gains with NIC count after 8 NICs and reduced performance after 10. So we kept the node count at 8 and replaced our single-core CPUs with dual-cores. Additionally we went to a diskless system with a 200MB RAM disk on each client node.

Low Latency COTS GigE

GAMMA beats IP for specialized apps in Bandwidth, Latency and CPU usage.

For an 8-node cluster (7 interconnected GAMMA NICs) we observed 5.8 μ sec latencies and 5% CPU overhead.



Future Plans : The Cluster will be reorganized to take advantage of the minimal amount of IP traffic we have. We can devote one NIC to it in production (primarily for NFS file transfers). Additionally, we can now scale the cluster up in traditional fashion by adding switches (as suggested by the FNN approach) to introduce single switch latencies that will no more than double our internode latencies while affording us to scale up to 64 nodes (256 PEs).