

Ethernet for Performance Computing

Moray McLaren, Technical Director (CTO)

moray@quadrics.com

www.quadrics.com

Why data centre Ethernet is of interest to Quadrics

- Opportunity to influence standard to support HPTC (high performance technical computing) requirements
- Take ideas developed in HPTC and apply to wider markets.

Delivering low network latency

- Packet sizes are increasing
- Packet length has significant effect on latency in busy networks
- Latency increase with number of switch stages.
 - $\sim 7\mu\text{s}$ for 8Kbytes on 10GE.
 - Worst case = $7\mu\text{s} * \text{number of hops traversed}$.
- Packet priorities with “pre-emption” allow mix of low latency high priority control packets with, large, efficient bulk data transfer packets.

Delivery low latency to the application

- In parallel computing worst case latency can be as important as average latency
- Packet retry schemes requiring software intervention have huge penalty on retried packets over delivered 1st time packets
- Can effect scaling of parallel applications
 - Very unlikely that a packet will have to be retried twice.
 - BUT on a large application very likely that in each time step at least one packet will have to be retried.

Quadrics contribution to data centre Ethernet development

- Participate in standards process
- Contribute experience in scalable, ultra low-latency networks
- Contribute experience of SAN type protocols for reliable delivery of packets
- Experience of real large scale, low latency deployments.