

## Hot Interconnects 2003 Panel Discussion

**Greg Pfister**  
**IBM, Austin TX**

### Virtual Garlic

- Any trademarks are the owners'
- Opinions are my own, not IBM's

### General Comment

- Fabrizio's questions all demand long, detailed answers.
- In 15 minutes total...

## Major Question 1

- Which interconnect is the best for high-performance computing and why?
- HPC too broad for any one “best” – from shared memory GUPS-lovers to Grids like United Devices, with various others in the middle (e.g., Enet, IB, various proprietary).

## Major Question 2

- What are the future trends in high-performance networking, and what are the implications of these trends?
- One general trend:
  - If the free ride from shrinking features is ending, then interconnects will increase in importance.
  - But SiTech guys have cried wolf.
- Another:
  - Better understanding of the market. Currently can't answer this question:
    - “If I plan / develop / qualify / launch / market / support a high-speed interconnect, how much money will I make?”

## Specific Question 1

- Are quantitative measures of latency and bandwidth enough to characterize a network interconnect, what other ways should we be evaluating interconnects?
- Assumption: "latency" includes software overhead.
- Management
- Virtualization
- Quality of Service
  - But not mainly, or primarily, multimedia streaming.
    - Response time to end user; or some other measure for HPC

## Specific Question 2

- Will the "status quo" in networking continue? That is, Ethernet as a commodity interconnect that is also used as a cheap commodity solution for clusters with InfiniBand, Quadrics, and Myrinet "relegated" to high-end and more costly clusters.
- **Possible** disruptor: a third, middle tier of IP offload / RDMA-enabled adapters & faster switches.
- Only "possible" because they may not be best of both worlds: Not cheap enough for some areas, not fast enough for others.

### Specific Question 3

- What assumptions must interconnects make about the underlying architecture (or what assumptions would they like to make)? PCI-X? PCI Express, Hyper Transport?
- Implicit assumption here: The interconnect is not an integrated part of the underlying system design.
- Obviously, the best implementations will be those in which system vendors attach the network directly to internal busses.
- (P.S.: PCI-<whatever> is already an interconnect.)

### SQ4

- In five years, how will today's interconnects evolve and/or compete in high-performance computing?
- In 5 years: *commercial* server-volume business in high speed interconnects, unlike the current niche-y situation, as commercial s/w vendors start exploiting their abilities.
- Require standards, high availability, management, virtualization, QoS, robust IOverX implementations, other commercially-important things.
- Volumes in those area will result in bleed-over to HPC as the ecosystem develops around these implementations.

## Q5

- InfiniBand started out as a high-performance I/O technology but has evolved into a general network interconnect for high-performance clusters. Will it replace Myrinet or Quadrics as the costlier high-performance interconnect for high-end clusters?
- See previous answer.
- Also, nothing changes as fast as anybody's PR machine says it will.

## 6

- What features and improvements are needed in communication subsystems to build next generation clusters? (network hardware, communication layers, libraries, programming models, etc.).
- Well, yes, what's listed in parentheses.
- Also, OS modifications to solve send vs. receive problems with general-use nodes.