



Linux Clusters Institute: Introduction to Networking

Ben Matthews, Software Engineer, NCAR

Types of Cluster Networks

- High-Speed Interconnect
- Internet/External Access
- Management/Boot/Lights-Out

Requirements

- High-Speed Interconnect
 - Low Latency
 - High Bandwidth
 - Reliable
 - User software support (MPI, etc)

Requirements

- Internet/External Access
 - Bandwidth
 - Access to/from external sites
 - Security
 - Reliable

Requirements

- Management
 - Cheap
 - Reliable Hardware/Software Support
 - PXE
 - IPMI
 - Secure (Isolation from users)

External

- Ethernet
- Many Vendors
 - Arista
 - Cisco
 - Juniper
 - Mellanox

Interconnect

- InfiniBand
 - Mellanox
- OmniPath (Intel)
- Cray
- Ethernet

Management

- Cheap GigE usually good enough
- NIC may or may not be shared with Lights-out management
- Should meet the requirements of your cluster manager (xCAT, SMC, HP CMU, Bright, Cobbler/Puppet)
- VLAN support desirable for isolating insecure hardware (managed switches, IPMI, etc)

Software

- Network gear exists to support your software – consider what you want to run before buying network hardware (or any hardware)
- MPI?
- Other middleware?
- RDMA? TCP? ...?
- LibFabrics? Verbs? PSM? Portals? ...?

Questions?



Extra Time: Physical Layer

