

Outline of Tutorial M13

1. Compute center I/O and IPC needs (20 min)
 - a. Typical I/O scenarios
 - i. Compute nodes, storage, gateways/grids
 - b. CPU bypass/offload – RDMA
 - i. Maximize compute time/communication time
2. InfiniBand Architecture overview (40 min)
 - a. Open standard
 - i. Industry-wide initiative
 - b. Switched fabric
 - i. Top-level functionality
 - c. Hardware components
 - i. Channel adapters, switches, gateways
 - d. Software design/components
 - i. User- and kernel-level protocols and interfaces
 - ii. Transport architecture: queue pairs, virtual lanes, etc.
 - e. Interoperability with legacy networks
 - i. IPoIB, Fibre Channel, ...
3. HPC over IBA in practice (1 hour 30 min)
 - a. In the box – Host channel adapters
 - i. Where do they go? – PCI-X, PCI-Express
 - ii. Features and limitations
 - b. Switches
 - i. Switch requirements and available solutions
 - ii. Modular designs – chassis and gateways
 - c. Management software
 - i. Subnet managers – right routing
 - d. SourceForge
 - i. Open source ULPs and Verbs interface
 - ii. Decouples user interface and vendor specific modules
 - e. Point-to-point performance
 - i. Bandwidth and latency
4. Applications over Message Passing Interface (30 min)
 - a. MPI implementations
 - i. Straight to verbs – MVAPICH (Ohio State)
 - ii. Through a middle layer – VMI (NCSA), Scali MPI, ...
 - b. Application performance
 - i. Xeon cluster – LS-DYNA, Fluent, Cactus, ...
 - ii. Itanium cluster – LS-DYNA, ...