

Future of Ethernet in Data Centers



OPENFABRICS
ALLIANCE

Gopal Hegde, Cisco
Pramod Srivatsa, Cisco

www.openfabrics.org

Agenda of Topics

- Converged Fabric Overview – Gopal / Pramod
- IEEE Data Center Bridging Update – Manoj Wadekar / Ilango Ganga
- UD Extensions for iWARP verbs – Terry Hulett

Snapshot – March 2009

Server / Access



- 10GbE NIC/CNA Vendors – Broadcom, Chelsio, Emulex, Intel, Mellanox, Myricom, QLogic, Neterion, NetXen, ServerEngines, Sun, Others...
- 10GbE Switch Vendors – Arista, BNT, Cisco, Extreme, Force10, Foundry, Fujitsu, HP, Juniper, Woven, Others...
- Tier-1 server vendors introducing landed on motherboard 10GbE
- 40GbE / 100GbE in IEEE Standardization Process
 - (IEEE 802.3ba – Working Group Ballot; On target for standardization in June 2010)

Behind the 10GbE Traction

At the Server / Access



➤ Performance

- Bandwidth / Latency Requirements

➤ Unified Fabric

- Total Cost of Ownership Savings

➤ Virtualization

- Treat VMs like physical servers from a networking perspective

Performance Drivers:

- Multi-core systems, PCIe Gen2, DDR3 memory etc..
- Increased reliance on numerical methods in enterprise
- Increased number of VMs
- Faster storage access (ex: solid state)

Data Center Bridging

Performance Benefits



Feature	Potential Benefits
<p>Priority-based Flow Control (PFC) ; Lossless Service, High Priority VL</p> <p>IEEE 802.1Qbb</p>	<p>Reduces latency and latency jitter; Provides ability to transport various traffic types (e.g. Storage, RDMA)</p>
<p>CoS Based BW Management (ETS)</p> <p>IEEE 802.1Qaz</p>	<p>Increases throughput, reduces latency and latency jitter of high priority traffic, when non-uniform traffic demands exist</p>
<p>Congestion Management (QCN)</p> <p>IEEE 802.1Qau</p>	<p>Reduces latency, increases throughput by preventing packet discards</p>
<p>L2 Multi-path for Unicast & Multicast</p> <p>IETF – TRILL WG</p>	<p>Build very large L2 fabrics ; Increases throughput by utilizing full Bi-Sectional bandwidth with ECMP</p>
<p>DCB Capability Exchange protocol (DCBX)</p> <p>IEEE 802.1az.org</p>	<p>Auto-negotiation for Data Center Bridging capabilities</p>

Fibre Channel over Ethernet

Delivering Unified I/O



Data Center Bridging Standards

Unified I/O Transport

- Mapping FC frames over Ethernet Transport
- Enables existing HBA stacks to run over a lossless Ethernet medium
- **Open FC** – Open source implementation effort for FCoE
- Single Adapter, less device proliferation, lower power consumption
- No gateways required

T.11 FC-BB-5 WG

Letter Ballot Completed
On track for standardization
on or before June 2009



Enabling Converged Network: DCE/FCoE Support in Linux Kernel



- Support for negotiating DCB links with DCBXP daemon
- Multi-Queue transmit and receive for multiple VLs, including configurable traffic classification for steering packets
- Existing FC Stack works seamlessly over FCoE Converged Network Adapters (CNAs)
- Software FCoE initiator that works with any Ethernet NIC

Thank you!



OPENFABRICS
ALLIANCE

www.openfabrics.org