A Practical Guide to Fibre Channel over Ethernet (FCoE)

Dennis Martin
President, Demartek

Demartek Company Overview

- Industry analysis with on-site test lab
- Lab includes servers, networking and storage infrastructure
  - Fibre Channel: 4 & 8 Gbps
  - Ethernet: 1 & 10 Gbps (with FCoE)
  - Servers: 8 cores, up to 96GB RAM
  - Virtualization: ESX, Hyper-V, Xen
- We prefer to run real-world applications to test servers and storage solutions
  - Currently testing various SSD and FCoE implementations
- Web: www.demartek.com

Agenda

- The Problem
- Buzzwords and Acronyms
- Key protocols and standards
- Technologies available now
- Advantages and disadvantages
- Effect on currently installed storage networks
- Demartek lab test results
- Futures, Commentary, Roadmaps

The Problem: Too Many Parts

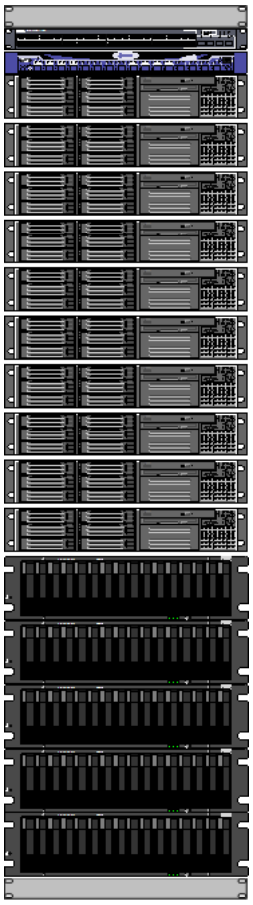

- Rack Servers (each)
  - 4 NIC ports (1Gb)
  - 2 FC ports (4Gb)
- Disk Array
  - 2 NIC ports (1Gb)
  - 4 FC ports (4Gb)
- Totals?
  - Cables & switch ports
  - Adapter cards
  - Maximum bandwidth

The Solution: Converged Network

Combine lossless features of Fibre Channel with ubiquity of Ethernet

- Within a rack (short-term)
- Entire infrastructure (long-term)

New Buzzwords & Acronyms

- Converged Network: combined LAN and SAN network
- Data Center Bridging (DCB)
  - CEE: Converged Enhanced Ethernet
  - DCE: Data Center Ethernet (Cisco trademark)
  - EEDC: Enhanced Ethernet for Data Center
- FCoE: Fibre Channel over Ethernet
  - FCoCEE: FC over CEE
- CNA: Converged Network Adapter

How Can This Work?

- Enhance Ethernet so that it properly handles storage traffic
- Data Center Bridging (DCB)
  - A collection of architectural Ethernet extensions designed to improve Ethernet networking and management in the data center
- FCoE is the first major application for DCB

Data Center Bridging (DCB)

- Traffic Differentiation
  - Can distinguish LAN, SAN and IPC traffic
- Lossless Fabric
  - Required for SAN traffic
- Optimal Bridging
  - Allows shortest path bridging within data center
- Configuration Management
  - Works with existing systems

Key Protocols and Standards

- Ethernet
  - IEEE 802.1
    - 802.1Qau Congestion Notification
    - 802.1Qaz Enhanced Transmission Selection
    - 802.1Qbb Priority-based Flow Control
  - DCB standards ratification expected in 2010
  - TRILL – Multipathing alternative
- Fibre Channel (FC)
  - INCITS T11: FC-BB-5 (“FCoE”)
    - Approved June 2009

Fibre Channel over Ethernet

- FCoE places the FC protocol on a new physical link
  - Uses Lossless Ethernet (DCB) physical links
  - Protocol and behavior is the same as traditional FC
- FCoE fabrics must be built with FCoE/DCB switches
  - Interoperate with traditional FC fabrics
  - Support all FC advanced features
  - Operate identically on FCoE and FC fabrics

Storage and the OSI Model

<table>
<tr><td colspan="6">Operating System and Applications</td></tr>
<tr><td colspan="6">SCSI Layer</td></tr>
<tr><td>FCP</td><td>FCP</td><td>FCP</td><td>FCP</td><td>iSCSI</td><td>SRP/iSER</td></tr>
<tr><td rowspan="3"></td><td rowspan="2"></td><td>FCIP</td><td>iFCP</td><td></td><td rowspan="3"></td></tr>
<tr><td>TCP</td><td>TCP</td><td>TCP</td></tr>
<tr><td>FCoE</td><td>IP</td><td>IP</td><td>IP</td></tr>
<tr><td>FC</td><td colspan="4">Ethernet</td><td>IB</td></tr>
</table>

Switch Technology Available

- DCB/FCoE switches
  - Blade Networks Virtual Fabric 10G Switch & RackSwitch G8124
  - Brocade 8000 FCoE switch, FCOE10-24 blade for DCX Backbone
  - Cisco Nexus 5000 series switches, 4000 blades
- Contain technology for:
  - 10Gb Enhanced Ethernet (lossless)
  - Optional: 4 or 8Gb Fibre Channel
  - Support for FCoE traffic
  - Support for iSCSI traffic

Adapter Technology Available

- Converged Network Adapters (CNA)
  - Brocade 1010/1020
  - Emulex OCE10102
  - QLogic QLE8x00 series (8000, 8100, etc.)
  - Mellanox 10GBase-T adapter (w/FCoE firmware)
- 10GbE NICs
  - Intel X520, Broadcom BCM57712
  - Uses native FCoE features in operating systems
- These adapters require PCI-Express slots
  - Newer adapters work best with PCIe 2.0

Adapter Technology Available

- Offload characteristics
  - CNAs: FC & FCoE supported in hardware
  - 10GbE NICs: FC & FCoE supported by software
  - Ethernet: Similar to good server-class NIC
- Connectors
  - Copper: CX4 and SFP+ (10GBASE-CR)
  - Optical: SFP+ (10GBASE-SR)
- Future: 10GBASE-T (RJ45, Cat 6 and 7)
  - May be available by end of 2010?

CNA Generations

- Generation 1
  - Separate Ethernet and FC ASICs on card
  - 10GbE and 4GbFC
  - Emulex and QLogic
- Generation 2
  - Single, combined Ethernet & FC ASIC
  - 10GbE and 10GbFC
  - Brocade, Emulex and QLogic
- Generation 3 soon to be released

Advantages & Disadvantages

- Advantages
  - Reduced number of cables
  - Reduced number of adapters
  - Reduced number of switches (over the long term)
  - Retain existing management software
- Disadvantages
  - Possible single points-of-failure
    - Single adapter
    - Single switch
  - Organizational issues

Organizational Issues

- In typical large shops today, networking and storage are separate departments
  - Networking: Dynamic (more changes)
  - Storage: Stable (fewer changes)
- Other areas of convergence
  - Consider voicemail & email
- Those that learn networking and storage will be in the best position

Effect on Current Storage Networks

Can use and coexist with existing storage networks

- Converged switches can pass FC traffic to existing FC SAN switches or to FC targets
- Existing storage management software should work with FCoE technology

Expected Deployment Phases

Expecting a slow, deliberate process

- 2008 – 2009
  - Early adopters, top-of-rack-switch, connects to existing storage networks
- 2010 – 2011
  - Core networking support and wider adoption of FCoE adapters, some FCoE storage targets
- 2011+
  - More native FCoE storage targets

Demartek Lab Test Results

- Past Testing:
  - Conducted FCoE “First Look” in May 2008
  - Participated in FCoE “Test Drive” in June 2008
  - NetApp native FCoE storage in January 2010
  - IBM/QLogic FCoE-to-FC storage in May 2010
  - Evaluation of Emulex OneCommand in Aug. 2010
  - Evaluation of Intel X520 FCoE/iSCSI in Sep. 2010
- Current Testing:
  - Testing with various adapters, switches and storage in the Demartek lab in 2010

FCoE Test Drive

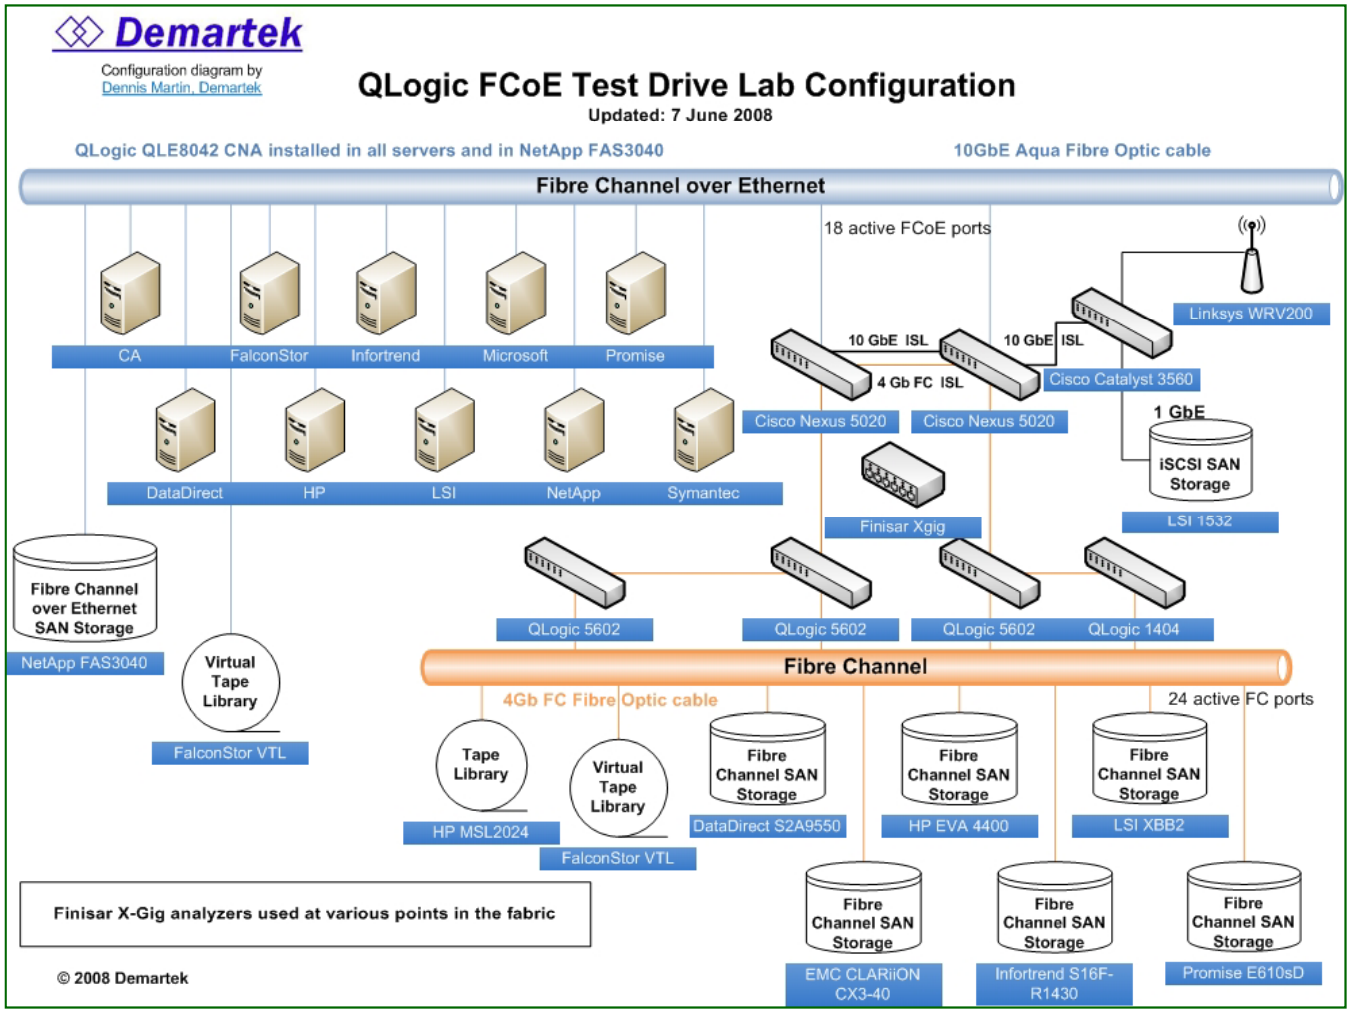

Servers connected via CNAs to FCoE switch connected to native FCoE storage and FC storage

FCoE General Comments

- Storage infrastructure changes slowly
- Should be considered in long-term planning, new equipment acquisitions and data center build-outs
- Standards
  - FCoE (FC-BB-5) is now a standard within the T11 committee
  - INCITS 462-2010 has been approved as a standard, available for purchase in July 2010

Roadmaps

- FC: 16-Gbps by 2011
  - SAN interface has a future
  - Disk drive interface approaching end-of-life
- Ethernet: 40 & 100 Gbps specifications (IEEE 802.3ba) ratified in June 2010
- FCoE will follow Ethernet roadmap
- Infiniband: 10, 20, 40 Gbps now, expecting 80 & higher

Future Technology Outlook

- Higher-speed adapters will require servers with PCI-Express 2.0 slots
- Vendors are scrambling for LAN-on-Motherboard (LOM) design wins
  - Expect to see 10GbE on server motherboards beginning in late 2010 or early 2011
  - Look for possible 10GbE + FCoE on server motherboards in the future
  - Expect to see both copper and optical connections available directly on server motherboards

Future Technology Outlook

- As interface speeds increase, expect increased usage of fiber-optic cables and connectors for most interfaces
  - At higher Gigabit speeds, copper cables and interconnects become too “noisy” except for short distances
- 10GBASE-T, expected in late 2010, may increase deployments of 10GbE and may drag some additional FCoE with it

Cabling

- Single-mode fiber (SMF)
  - 9 µm (microns), very-long distance, yellow
- Multi-mode fiber (MMF)
  - 62.5 µm or 50 µm, medium distance
    - Orange: OM1 (62.5 µm), OM2 (50 µm)
    - Aqua: OM3 and OM4 (both 50 µm)
  - 10GbE is best with OM3 or OM4
  - OM3 and OM4 also will support 40GbE & 100GbE
- Cable deployments change very slowly, so choose 10GbE cabling wisely

Before and After

- Compare parts list with DCB & FCoE
  - Cables & switch ports
  - Adapter cards
  - Maximum bandwidth
- What can be eliminated?

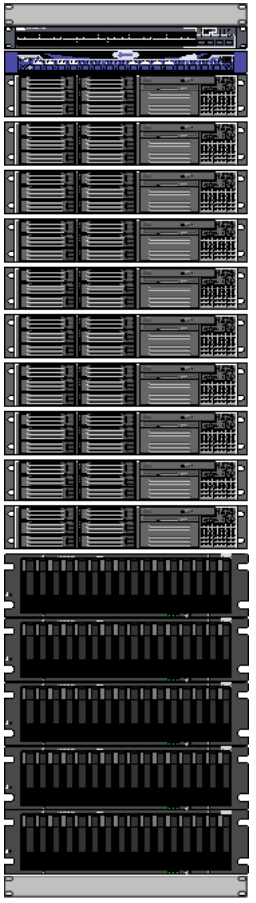

Demartek FCoE Resources

- Demartek has compiled a free comparison reference guide of FCoE and the other storage networking interfaces, which is updated periodically and includes roadmap information. This also includes cable distances and speeds.
  www.demartek.com/Demartek_Interface_Comparison.html

- Demartek FCoE Zone
  - http://www.demartek.com/FCoE.html

Free Monthly Newsletter

- Demartek publishes a free monthly newsletter highlighting recent reports, articles and commentary. Look for the newsletter sign-up at http://www.demartek.com.

Contact Information

Dennis Martin, President
Demartek
(303) 940-7575
dennis@demartek.com

www.linkedin.com/in/dennismartin
http://twitter.com/demartek