Demystifying Network and I/O Convergence: What You Need to Know and Why?

COMPUTERWORLD

Simon Gordon Senior Product Line Manager Juniper Networks



Abstract

 Many customers are currently evaluating options for upgrading their data center network infrastructures and have a natural desire to deploy the ideal products that provide the lowest cost of ownership and highest return on investment as they start to implement some level of I/O or network convergence.

COMPUTERWORL

Abstract

 During this session, Simon will explain the standards involved in network and I/O convergence, the purpose each of these standards serve within the network, what standards are in use with products in the market today and the benefits they provide, and the implications of the still evolving standards on investment protection with these products.



Abstract

- Attendees will:
 - 1. Learn how to apply the various networking standards when considering data center network upgrades
 - 2. Learn how to take advantage of the benefits these standards and incorporated products provide to optimize their data center networks
 - 3. Hear best practices and benefits for successful data center network infrastructure upgrades

COMPUTERWORLD

Agenda

- How is the Data Center changing
 - What are the protocols that apply
 - What operational changes apply
- What can be done to prepare
 - Layout the data center for the future
- How best can I benefit today
 - How to take the first steps
- What comes next
 - Unified fabrics and full network convergence

COMPUTERWORLD

Agenda

- How is the Data Center changing
 - What are the protocols that apply
 - What operational changes apply
- What can be done to prepare
 - Layout the data center for the future
- How best can I benefit today
 - How to take the first steps
- What comes next
 - Unified fabrics and full network convergence



What Storage Protocol

- Most of what I will say applies no matter the actual storage protocol in question...
 - iSCSI
 - FCoE
 - NAS
 - Other
- Some is specific to FCoE
 - Specific FCoE Deployment models
 - FCoE Operation

How is the Data Center changing

• Old world

- Servers had fixed roles
- Servers had multiple NICs, HBAs
- Datacenter had multiple networks
- New World
 - Fungible resources
 - I/O & Network Consolidation
 - Pooling of resources at much larger scale



- **Data Center LAN**
 - North South Client Server traffic
 - Limited East West Server Virtualization traffic
 - High Oversubscription possible with limited server mobility

Storage Area NW

- Heavily Pooled
- Low Oversubscription
- Low Latency
- Dual Rail

DATACENTER DESIGN THE SERVER EYE VIEW







Data Center Networks In reality...



Data Center LANs

- North South Client Server traffic
- Limited East West Server Virtualization traffic
- High Oversubscription possible with limited server mobility

Storage Area NW

- Heavily Pooled
- Low Oversubscription
- Low Latency
- Dual Rail

Server View Load Balancing & High Availability

COMPUTERWORLD



Campus network Cluster/Application Network Backup network Storage Network Management Network

. . .

COMPUTERWORLD **Server View** Load Balancing & High Availability SCSI SCSI AP Nic AA Nic IP Load Load AP Nic AA Nic IP Load Load Team Team Balancing Balance Team Team Balancing Balance **Different networks** setup with different topologies



Moving to the new world

- the convergence protocols
 - Mostly DCB implications
 - Some FCoE implications
- Load balancing & ha deployments
 - Across multiple Ethernet services
 - Ethernet vs. FC/FCoE model





Data Center LAN

COMPUTERWORLD

 Collapsing the multiple Ethernet network segments in to a common access 10G layer

Storage Area NW

- Heavily Pooled
- Low Oversubscription
- Low Latency
- Dual Rail

1Wire 8 Channels – PFC & ETS



COMPUTERWORLD

DCBX is just DCB auto negotiation

The standard gives flexibility on implementation



Server View Load Balancing & High Availability



Having collapsed many NICs and Networks each of which may have been configured and operated differently it is necessary to consider the load balancing and high availability model in a single cohesive way for the new converged network...



Ethernet Load Balancing & High Availability



COMPUTERWORLD



The Physical Switch deployment model as well as the L2/L3 Topology will determine your options for NIC teaming, HA, Load Balancing



Server View Load Balancing & High Availability



FCoE Clarifications – Multi Hop

Classic FC Definition

- Cascaded FCoE based FCFs
- Everything is exactly the same as classic FC except I now have point to point connections on ethernet wires instead of on fibre channel wires
- VE-VE port ISLs instead of E-E ISLs

FCoE Clarifications – Multi Hop

- Another possible definition
 - I have pure L2 hops (L2 Ethernet switches) in between my FC aware devices
 - L2 Switches in between the servers and the FCoE enabled FCF
 - L2 Switches in between a group of FCoE enabled FCFs





FCoE Implications

- In both cases
 - L2 Visibility whether physical or through VLANS determines what possible virtual FC connections can be established
 - FC layer has no understanding of the underlying Ethernet topology which is perhaps unfortunate and places heavy requirements on L2 partitioning

FCoE Load Balancing Virtual Connection Status





FCoE & Multiple Fabrics



FCoE Load Balancing & High Availability

COMPUTERWORLD



NOTE: we can largely consider the SAN (from the perspective of the Ethernet network) as just another end point on that network...

Particularly with HW based CNAs we can also largely consider the FCoE HA model separately to the general Ethernet HA model – though some may be concerned about the risk to classic dual rail SAN

FCoE Logical Connectivity Single Ethernet Dual SAN



Server View Software Based FCoE Stack



Server View Hardware based FCoE Stack



COMPUTERWORLD

Other Stuff

- Solving the spanning tree problem
 TRILL, 802.1aq, Stacking, Distributed Lag
- Options for congestion management
 ECN, TCP, QCN, ICMP Source Quench
- Handling Virtual Servers
 - VEPA etc
 - Potential overlap between virtual switches and storage over Ethernet

COMPUTERWORLD

Agenda

- How is the Data Center changing

 What are the protocols that apply
 What operational changes apply
- What can be done to prepare
 - Layout the data center for the future
- How best can I benefit today
 - How to take the first steps
- What comes next
 - Unified fabrics and full network convergence



Physical Preparations

- Understand the cabling requirements
 10GbE / 40GbE / 100GbE
 - Copper connectivity to the server
- Implications for Data Center Layout
 - Servers / Racks / Rows / Pods
 - Location of networking equipment

COMPUTERWORLD

The short story

- 10G First hop most likely copper
 - Server to tor
 - limit in m's
- Subsequent hop optical 40G+
 - ToR to Aggregation / Fabric Core
 - 150m on OM4 for 40G & 100G
- Data Center Layout
 - PODS of up 32 to 128 Racks
 - Multiple PODS making up the data center





COMPUTERWORLD

Aggregation/core layer QSFP/CXP QSFP/CXP MTP connector MMF OM3 cable minimum(100m) OM4 preferred for slightly greater reach. Structured cabling (12 pair) **Access layer**

SFP+ optical transceiver 300 meter MMF with SR

Twinax copper cable

1, 3, 5, 7 meter



SFP transceiver Copper (1000Base-T)



40GbE/100GbE ready Access cabling

- Use MTP terminated trunk cables between access to aggregation switches
- MTP to LC cassettes and LC to LC breakout cables for access layer connection
- Patch panel housing in main distribution area
- MTP to LC harnesses from patch panel to switches



COMPUTERWORLD

Agenda

- How is the Data Center changing – What are the protocols that apply
 - What operational changes apply
- What can be done to prepare
 - Layout the data center for the future
- How best can I benefit today
 - How to take the first steps
- What comes next
 - Unified fabrics and full network convergence



Can I benefit today ?

- Yes plan for deploy in phases in 2011
 - Ethernet I/O Consolidation
 - Understand and deploy DCB
 - Flatten/Simplify the Ethernet network
 - Readiness for convergence and virtualization
 - Convergence access layer using FCoE
 - FCoE Transit Switch
 - FCoE-FC Gateway using NPIV
 - Or converge with iSCSI, NAS, etc

Data Center Networks FCoE Based Convergence



Data Center LAN

- Common access layer FCoE Transit Switch or FCoE-FC Gateway
- Separate Ethernet Aggregation / FC Backbone
- Uplinks to SAN Backbone could be either FC or FCoE
- What about blade servers ?



FCOE TRANSIT SWITCH FCOE-FC GATEWAY



SERVER ACCESS LAYER CONVERGENCE OPTIONS

FCoE Transit Switch	FCoE-FC Gateway
Comparatively Low Cost Just a DCB Switch with FIP Snooping	Higher cost for FC Personality Software License and/or Physical Ports
Interesting at ToR	Interesting at ToR
May occur as embedded blade server switch instead of pass through module	Investment protection if can buy as a DCB switch and upgrade later to a gateway
Just DCB + FIP Snooping	Converts from FCoE to FC Also works as a DCB switch. Not a FCF – NPIV based proxy
For FC-BB-5 cannot provide end to end FCoE solution – still required a full FC Fabric somewhere	Requires external FC Fabric, can also support upstream FCoE Transit Switches
Clean management separation as nothing for the SAN team to actively manage	Some overlap of management as gateway is active in the data-path and load balancing





FCoE-FC Gateway at ToR No Multi-hop FCoE





FCoE-FC Gateway at ToR Limited Multi-hop FCoE







FCoE Transit Switch at ToR Limited Multi-hop FCoE



FCoE Transit Switch at ToR Multi-hop FCoE

COMPUTERWORLD

SNI



COMPUTERWORLD

Agenda

- How is the Data Center changing

 What are the protocols that apply
 What operational changes apply
- What can be done to prepare
 - Layout the data center for the future
- How best can I benefit today
 - How to take the first steps
- What comes next
 - Unified fabrics and full network convergence



What comes next

- The FABRIC is coming
 - Again understand the full implications
 - Its not just about eliminating spanning tree
 - Implications of fully overlaying the networks
- FCoE is still evolving FC-BB-6
 - Scale UP
 - Scale DOWN
 - Implications of overlay



FC-BB-6 (EXPECTED 2012)

FCOE FC-BB-6 Large Scale

FCoE FC-BB-6 Direct Attach

Traditional FC SAN FCoE: FC-BB-5 Converged Access

Thank You – Questions ?

- 2011 is the year of...
 - Deploying 10G at the edge
 - Getting ready for 40G/100G at the core
 - Considering Fabric implications
 - Deploying full convergence at the edge
- 2012

COMPUTERWORLD

- Will it be the year for full convergence ?



Driving Innovation Through the Information Infrastructure

SPRING 2011