

Fibre Channel over Ethernet

From Hype to Reality

Frederick Knight



- ❑ **FCoE is real**
- ❑ **Demoing for 2 years**
 - ❑ **First public demo @ SNW – Oct 16, 2007**
 - ❑ **POC done in customer labs**
- ❑ **Pre-standard Configurations**
 - ❑ **Available for about 1 year**
- ❑ **Standard products now available**

- **FCoE Hardware**
 - **Then and Now**
 - **Hosts, Switches, Targets**
- **Ethernet Enhancements (DCB)**
- **FCoE Topologies**
 - **Then and Now**
- **FCoE Tomorrow**

FCoE Hardware

FCoE Converged Network Adapter

- ❑ Then (1st generation - discrete components)
 - ❑ used lots of power
 - ❑ required full size slots
 - ❑ early adopters /
proof of concept
- ❑ Now (integrated solution)
 - ❑ low power
 - ❑ smaller size
 - ❑ general adoption



FCoE NIC chips

- ❑ Then (1st generation)
 - ❑ All Software
 - ❑ Some parts of DCB hard to do in S/W
- ❑ Now (2nd generation)
 - ❑ Some H/W offload added
 - ❑ Lower CPU overhead



- ❑ Then (1st generation)
 - ❑ Stable Hardware
 - ❑ Firmware based on early draft of standard
 - ❑ Interoperated w/others at same draft level
- ❑ Now (2nd generation)
 - ❑ The Same Hardware
 - ❑ Firmware update to standard compliant
 - ❑ Broader Interoperability

FCoE Native Targets

- ❑ Then (1st generation)
 - ❑ Based on 1st generation CNA
 - ❑ Limited feature set
 - ❑ Limited Availability
- ❑ Now (2nd generation)
 - ❑ Based on 2nd generation CNA
 - ❑ Broad feature set
 - ❑ General Availability

Ethernet Enhancements (DCB)


DCB Feature set at IEEE

- ❑ Priority Flow Control (802.1Qbb)
In task group ballot
Current implementations compatible with the standard
- ❑ Congestion notification (802.1Qau)
In working group ballot
- ❑ Enhanced Transmission Selection(802.1Qaz)
Still in the working group
- ❑ Expected by November 2009

FCoE Topologies


Then and Now

- ❑ FCoE Interoperates with existing FC
 - ❑ Existing FC devices
 - ❑ Existing FC switches
- ❑ FCoE multi-vendor switch interoperation is the same as FC multi-vendor switch interoperation




Then

- ❑ FCoE End points wired directly to FCoE switch
 - ❑ From FCoE initiators, or from FC initiators
 - ❑ To FCoE targets, or to FC targets.
- ❑ Some FCoE link errors were not detected
- ❑ No intermediate Ethernet switches



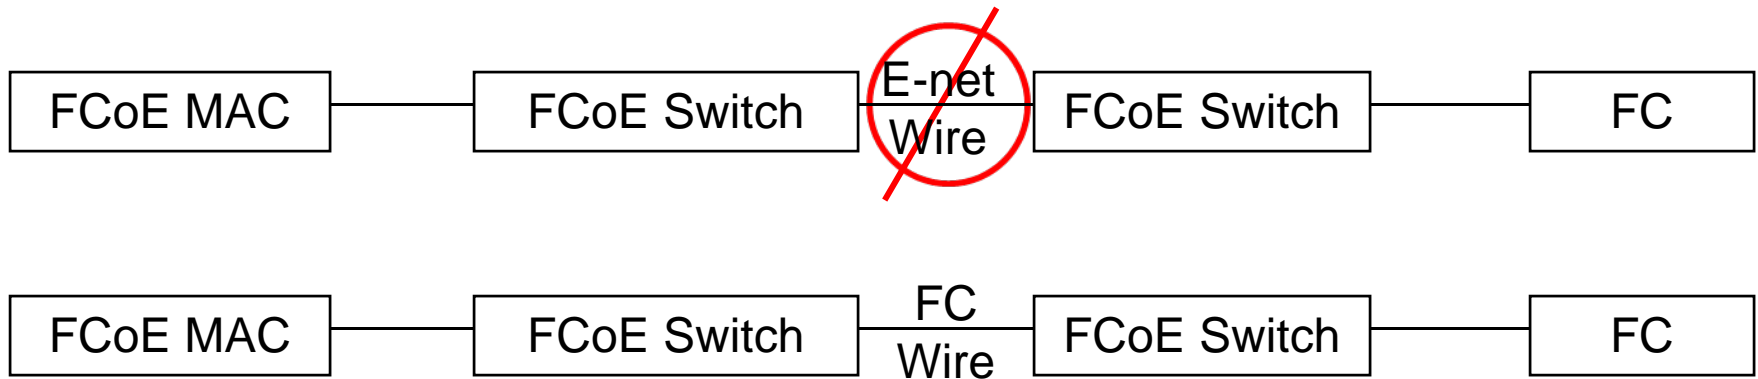
Now

- ❑ FCoE End points wired directly to FCoE switch
 - ❑ From FCoE initiators, or from FC initiators
 - ❑ To FCoE targets, or to FC targets.
- ❑ Improved link error detection (via FIP)
- ❑ FCoE through 10Gb Ethernet DCB switches
- ❑ All traffic must still transit the FCoE switch




Tomorrow


- ❑ Fabrics of FCoE switches
 - ❑ FCoE E-Node connectivity (Ethernet connecting FCoE switches) still coming
- ❑ Can connect FC ports today



FCoE in Top of RACK

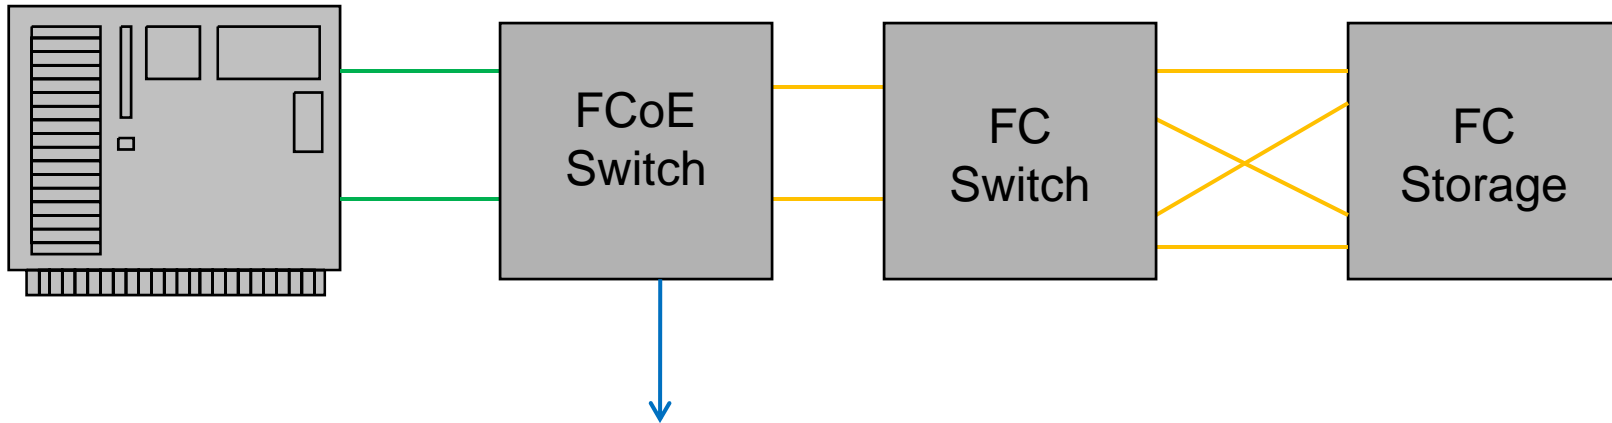


FCoE in Top of RACK




Sample FCoE Deployments

Customer (many)




- ❑ Many host platforms
- ❑ Many host O/S (ESX, Linux, Windows)
- ❑ Regular FC storage

Customer Before FCoE



26 Servers	Ethernet	FC	Total
Adapters	26	26	52
Switches	2	2	4
Cables	56	56	112
Uplinks	4	4	8

Customer After FCoE



26 Servers	CNA	Total	Savings
Adapters	26	26	50%
Switches	2	2	50%
Cables	60	60	46%
Up links	8	8	0%


Customer Cost Analysis

Cost per server

Component	FC and Ethernet	FCoE
Cable	\$300	\$200
FC HBA (2-port)	\$1200	-
1GbE NIC (2-port)	\$800	-
10GbE FCoE CNA (2-port)	-	\$1800
FC switch (2-ports)	\$2400	-
1GbE switch (8-ports)	\$2800	-
FCoE switch (2-ports)	-	\$4000
Total	\$7500	\$6000

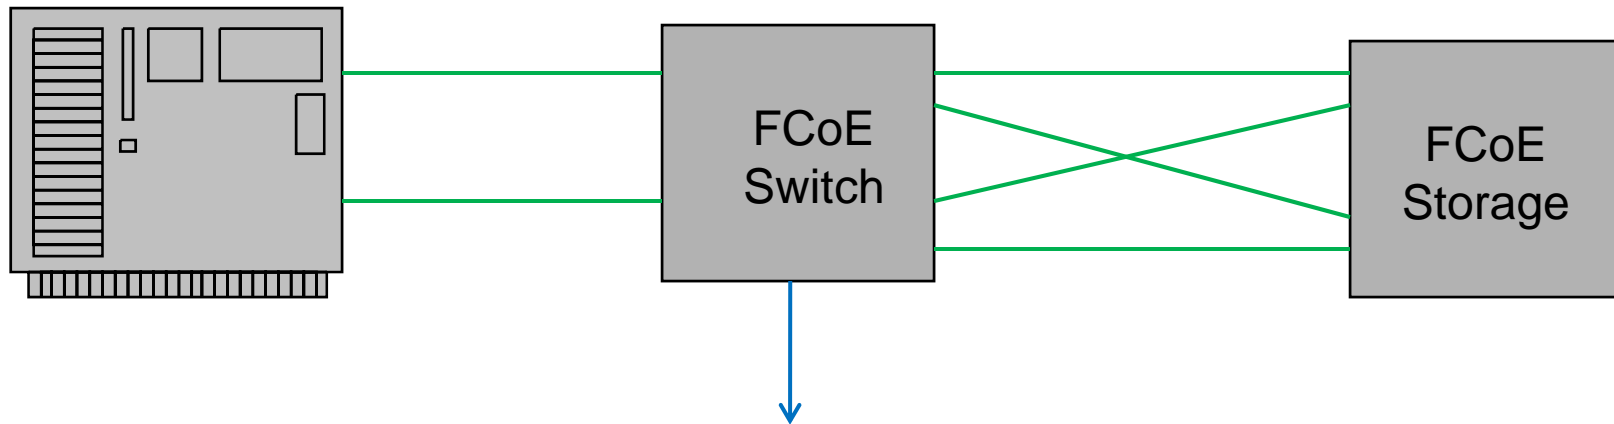
20% 

Customer 2




- ❑ HP DL580 w/ESX 3.5 U3 (w/MPIO) – MS 2k3 & 2k8 guests – QLE8042
- ❑ Nexus 5010/5020
- ❑ FAS3140 native FCoE

Customer 3



- ❑ Dell PowerEdge R900 w/ESX 3.5 U3 (w/MPIO) – MS 2k3 & 2k8 guests – QLE8042
- ❑ Nexus 5010/5020
- ❑ FAS6080 native FCoE


Customer 4



- ❑ Dell PowerEdge R900 w/ESX 3.5 U3 (w/MPIO) – MS 2k3 & 2k8 guests – QLE8042
- ❑ Nexus 5010/5020
- ❑ FAS3140 native FCoE

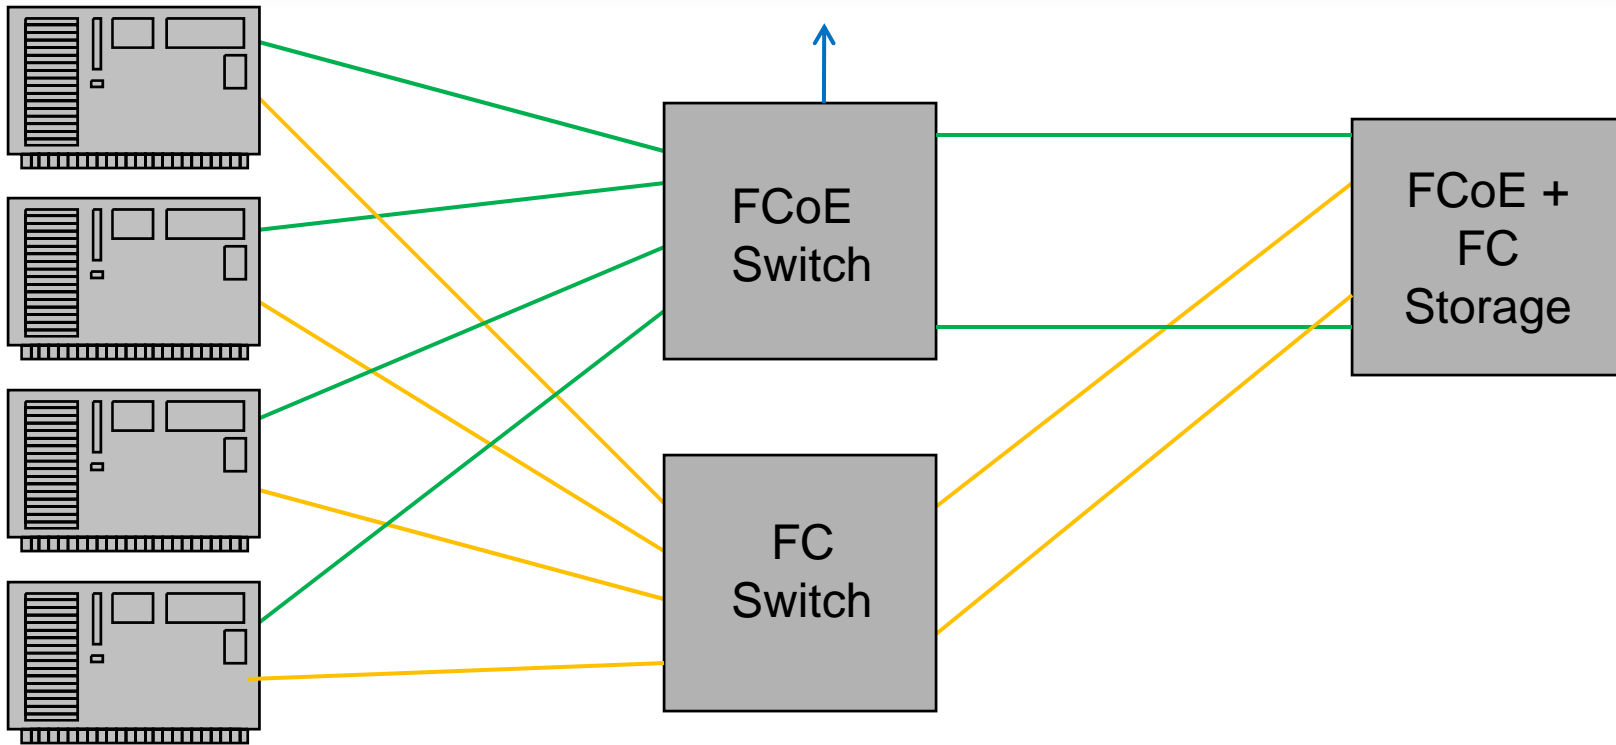
"We have been testing our new FCoE infrastructure over the last few months and have been impressed most by the performance," said Howard Eddy, vice president of Information Technology at Subaru New England. "For instance, we are seeing over 75% greater performance on SQL versus our existing DAS solution. We've also been able to dramatically reduce our overall complexity for our storage and networking by streamlining down to a single network technology: Ethernet."

Customer 5




- ❑ Dell PowerEdge R6850 and Sun X4100 w/ESX 3.5 U3 (w/MPIO) – MS 2k3 & 2k8 guests – LPe21002
- ❑ Nexus 5010/5020
- ❑ FAS3140 native FCoE

Customer 6




- ❑ HP DL380 w/ESX 3.5 U3 – MS 2k3 and 2k8 – LPe21002 and QLE8042
- ❑ Nexus 5010/5020
- ❑ FAS3170 native FC and native FCoE

Customer 7



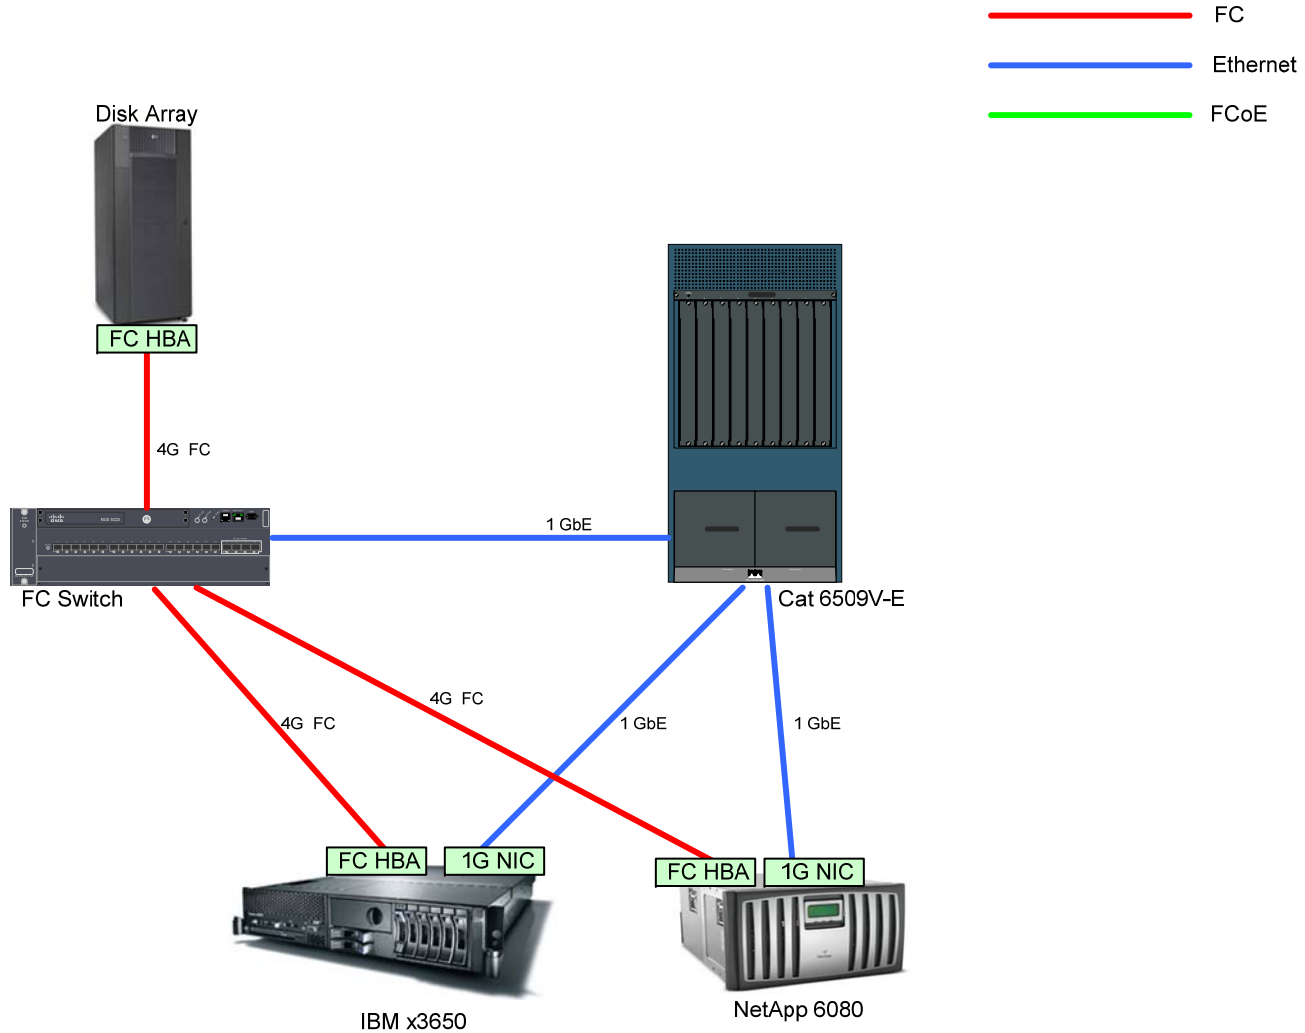
- ❑ HP DL380 w/ESX 3.5 U3 (w/MPIO) – MS 2k3 & 2k8 guests – LPE21002 and QLE8042
- ❑ Nexus 5010/5020
- ❑ FAS3140 native FCoE

Customer 8




- ❑ Intel Xeon EM64T w/MS 2k3 (w/MPIO) – LPE21002
- ❑ Nexus 5010/5020
- ❑ FAS3140 native FCoE

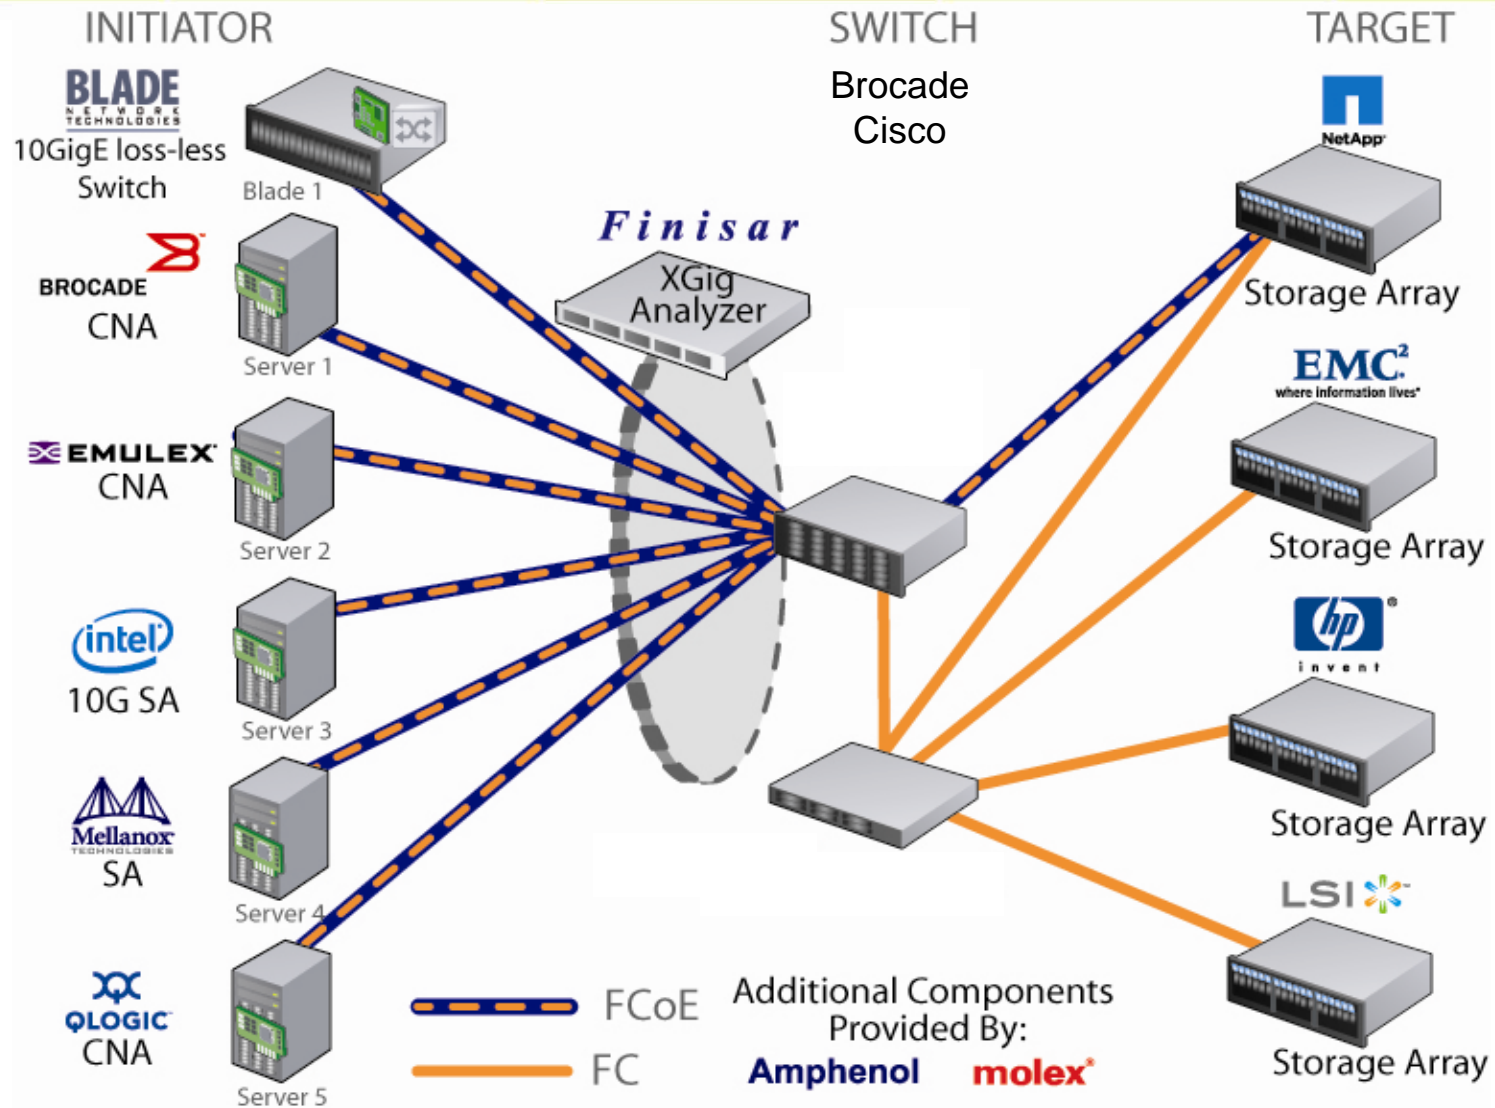
Customer 9 Before



Customer 9 After



FCoE SNW Demo Topology



FCoE Futures

- ❑ Three Previous plugfests
 - ❑ Two were FCoE specific
 - ❑ One Ethernet Focused (DCB)
- ❑ Next FCoE Plugfest
 - ❑ Nov 16-20 2009
 - ❑ University of New Hampshire Interop Lab
 - ❑ Multi-Vendor
 - ❑ Includes DCB features

- ❑ FC-BB-6 efforts began in June 2009
 - ❑ FC-BB-5 is the standard that contains the approved FCoE
 - ❑ FC-BB-6 is the standard that will contain the next version of FCoE
- ❑ FC-BB-6 plans to include
 - ❑ Direct MAC to MAC transfers
 - ❑ Probably negotiated through the FCoE switch
 - ❑ Direct End Point to End Point configurations
 - ❑ No switches (but just 2 nodes)
 - ❑ Bit Error rate examinations
 - ❑ Other good ideas we come up with
- ❑ But we're still working on it

- ❑ **FCoE is real**
- ❑ **Demoing, POCs, and pre-standard products available for a long time**
- ❑ **Standard products now available**
- ❑ **Finalization of IEEE standards will only expand the available FCoE topologies**