BLADE
NETWORK TECHNOLOGIES

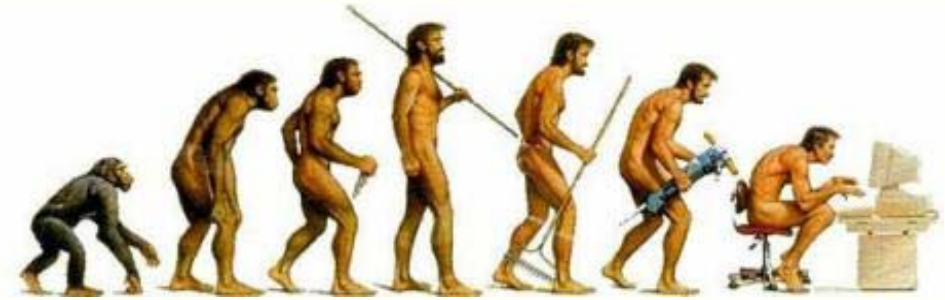

Data Center Bridging (aka CEE)

It is not just for FCoE

BLADE Background

- **Products**
  - Switches for blade servers
  - Data center top-of-rack switches
  - Virtualization & management software
- **Achievements**
  - 7.5M datacenter ports shipped
  - 16,000+ switches at one customer
  - 50+% blade switch market share
  - #3 in 10 Gigabit Ethernet (Dell'Oro Group)
  - Firsts
    - 1st FCoE for blade servers
    - 1st Junos licensee
    - 1st VM-aware networking
    - 1st CEE, Layer 2-3 & Layer 2-7 for blades
    - 1st 10G for blade servers & under $500/port
    - 1st green networking company
    - 1st with unique airflow technology
- **Customers**
  - 9000+ customers in over 330 of the Fortune 500
  - 26+ industries: #1 retailer, #1 financial services
  - Key partnerships:

- **Company**
  - Founded 2006
  - Series B 2009: $230M valuation
  - Santa Clara, CA Headquarters
  - Offices in Canada, Europe, Middle East, Australia, Japan, China, Korea, Singapore

Data Centers are addicted to FC

BLADE NETWORK TECHNOLOGIES

In the 90's FC offered shared storage and relief from parallel SCSI's distance limits

"Today Fibre Channel is the Opiate of the Data Center" Carlos Marx (SAN admin)

Addiction = "persistent compulsive use of a substance known by the user to be harmful" merriam-websters

Why is Fibre Channel so harmful ?

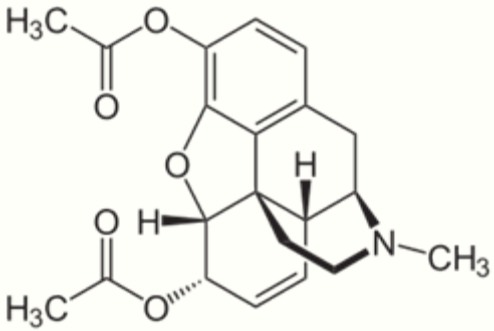

- Interoperability issues
- Uncompetitive performance
- Very expensive - run by Duopolies (aka Cartels)

Fibre Channel – Interoperability

Fibre Channel is plagued with Interoperability issues

- FC standards were “loosely” followed
  - Non-interoperability is now defacto
- Multi-Vendor FC switches only interoperate at basic levels
  - Customers have to choose 1 vendor for their entire SAN
- FC HBA, switches and arrays require interoperability testing
  - Customers need to check compatibility matrices before buying.
  - See http://www.emc.com/collateral/elab/emc-support-matrices.pdf (14060 pages !)
- Ethernet has proven to have none of these issues.
  - It just works!

Fibre Channel – Speed
(or lack thereof)

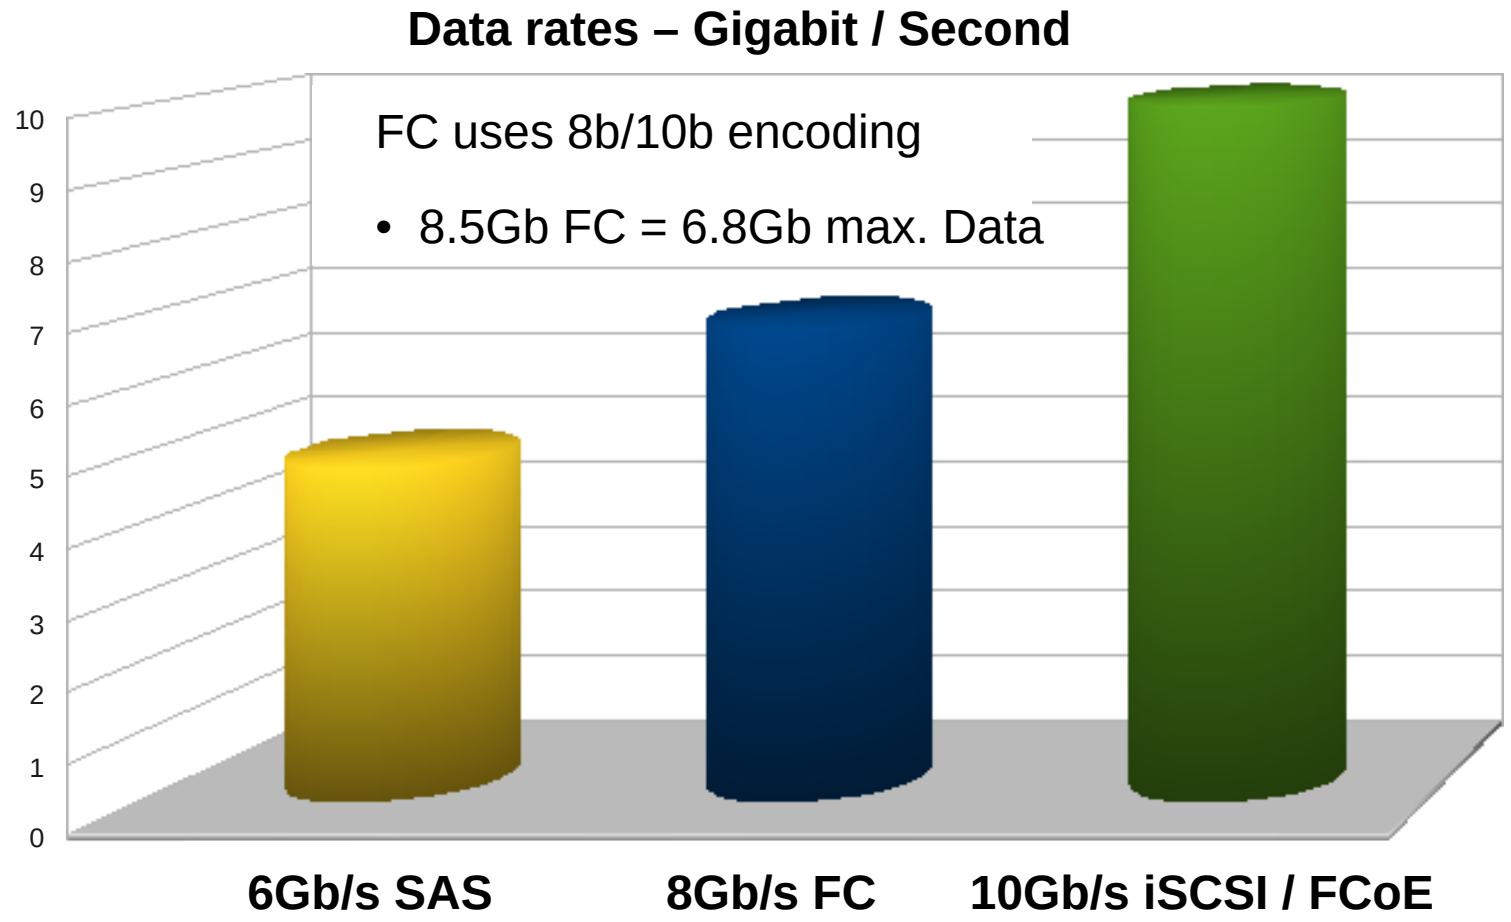

- **Ethernet has a significant performance advantages over FC**
- **When FC goes to 16Gb – Ethernet moves to 40Gb and 100Gb**

Acquisition Cost of Ethernet & FC

Example based on IBM BladeCenter

BLADE NETWORK TECHNOLOGIES

Component	10 GbE	8 Gb FC	Comments
Total Throughput Gb/sec	**480 Gb/sec** (24*10*2)	**260 Gb/sec** (20*8*2(8b/10b)	Total Gb/s bidirectional
Mezzanine Cards			
2 port 10 GE NIC Daughter card (46M6168)	$599 x 14 **$8,386**		10G NICs
2 port 8 Gb Fiber Channel HBA (44X1945)		$909 x14 **$12,726**	8Gb FC HBAs
BladeCenter IO			
10 GbE Switches for BladeCenter (46C7191)	$11,199 x 2 **$22,398**		Ethernet 24-port switch
8 Gb FC SAN Switches for BladeCenter (42C1828)		$13,999 x 2 **$27,998**	FC 20-port switch
Maintenance Costs for 3 years – 24*7 Service	$4,618	$6,108	5% of cost/year
Total	**$35,402**	**$46,832**	
Cost per Gb / sec	**$74**	**$180**	

- **FC is more than twice the cost of 10Gb Ethernet**

Getting Data Centers off FC

FCoE enables a convergence to Ethernet while leveraging FC investment in devices, tools and training

- Same FC protocols over a new wire

Is FCoE a cure or just Methadone?

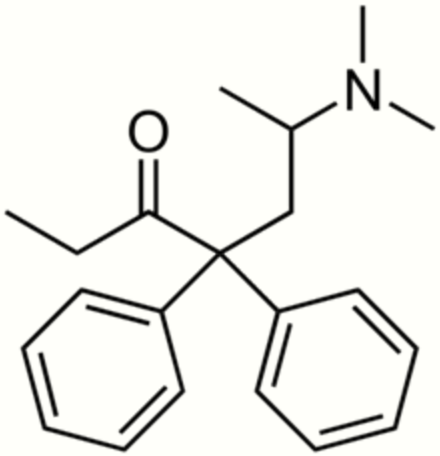

Better performance than FC ?

- Gb/s Yes ; IOPs close

Interoperable?

- No, same FC issues in FCoE

Lower cost than FC ?

- Yes, but savings limited to the server edge

The Holy Grail of Convergence - or what are we trying to achieve ?

- Lower power consumption
- Lower purchase cost
- Lower management cost

SAN, LAN and Clustering are converging to Ethernet

Cost Saving via Fabric Convergence

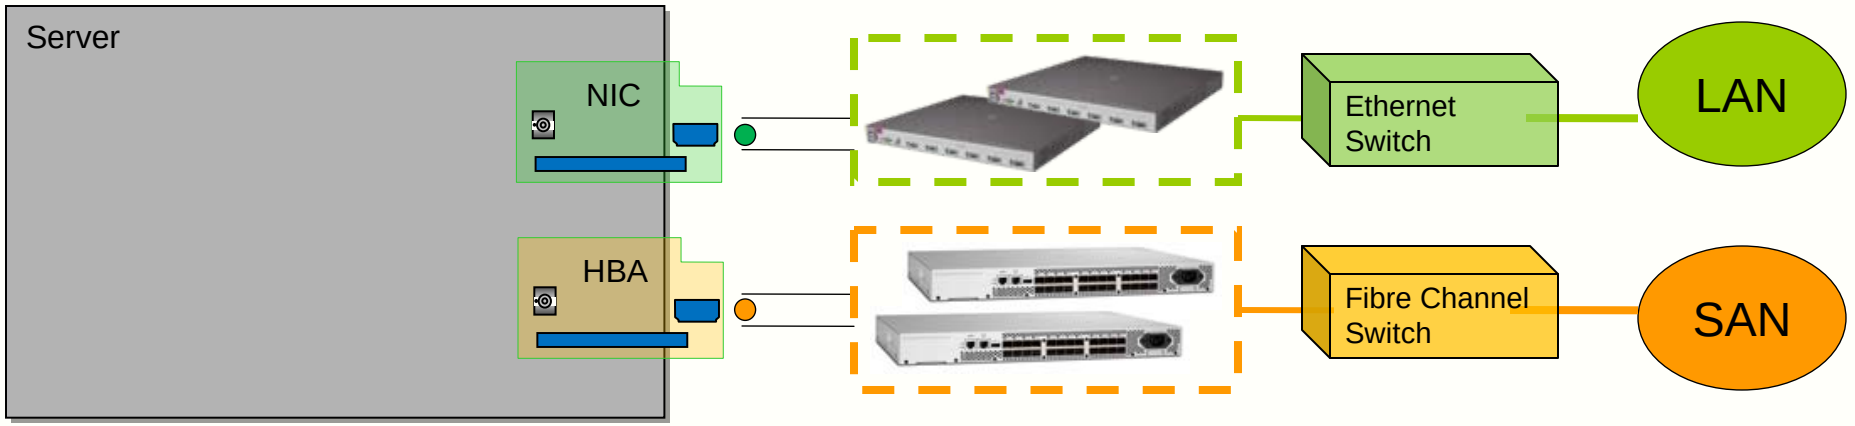

Before Convergence

Separate LAN & SAN Adapters and Switches

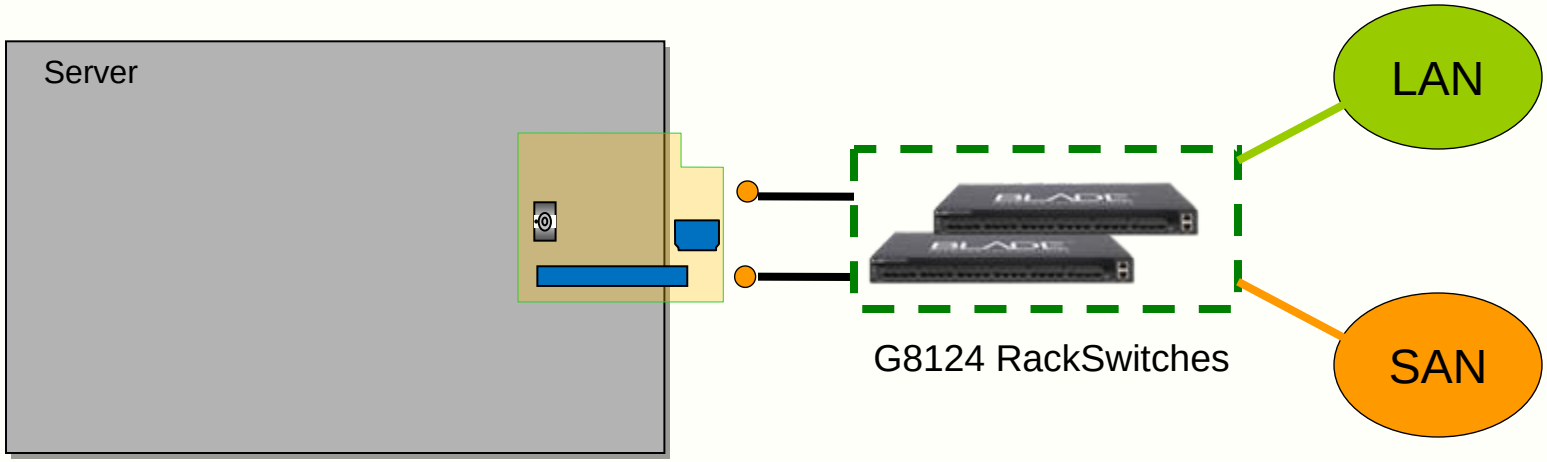

After Convergence

50% Reduction in switches & adapters, lowers cost & power

Converge SANs to FCoE or iSCSI ?

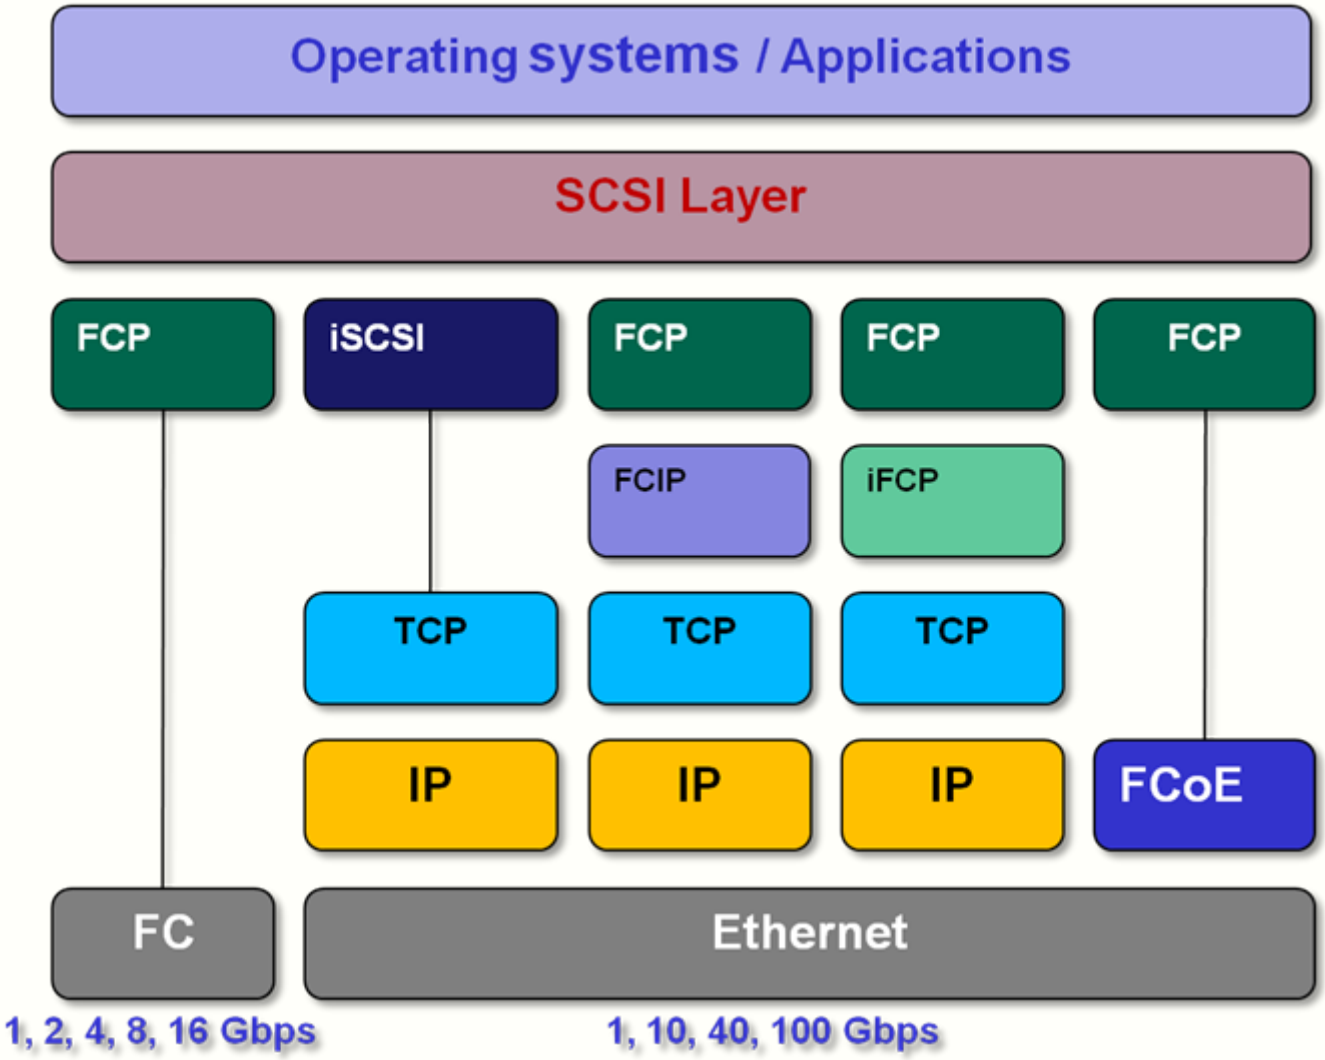

CEE/DCB improves iSCSI (& NFS)

iSCSI uses the Transport Control Protocol (TCP)

- TCP is designed to work over unreliable fabrics
  - If a packet is dropped, TCP will:
    - Request the packet is retransmitted
    - Reorder the Packets
    - Slow down transfer rates to reduce further packet loss

Loss-Less Ethernet eliminates packet drops therefore improves TCP performance and predictability

DCB and iSCSI performance

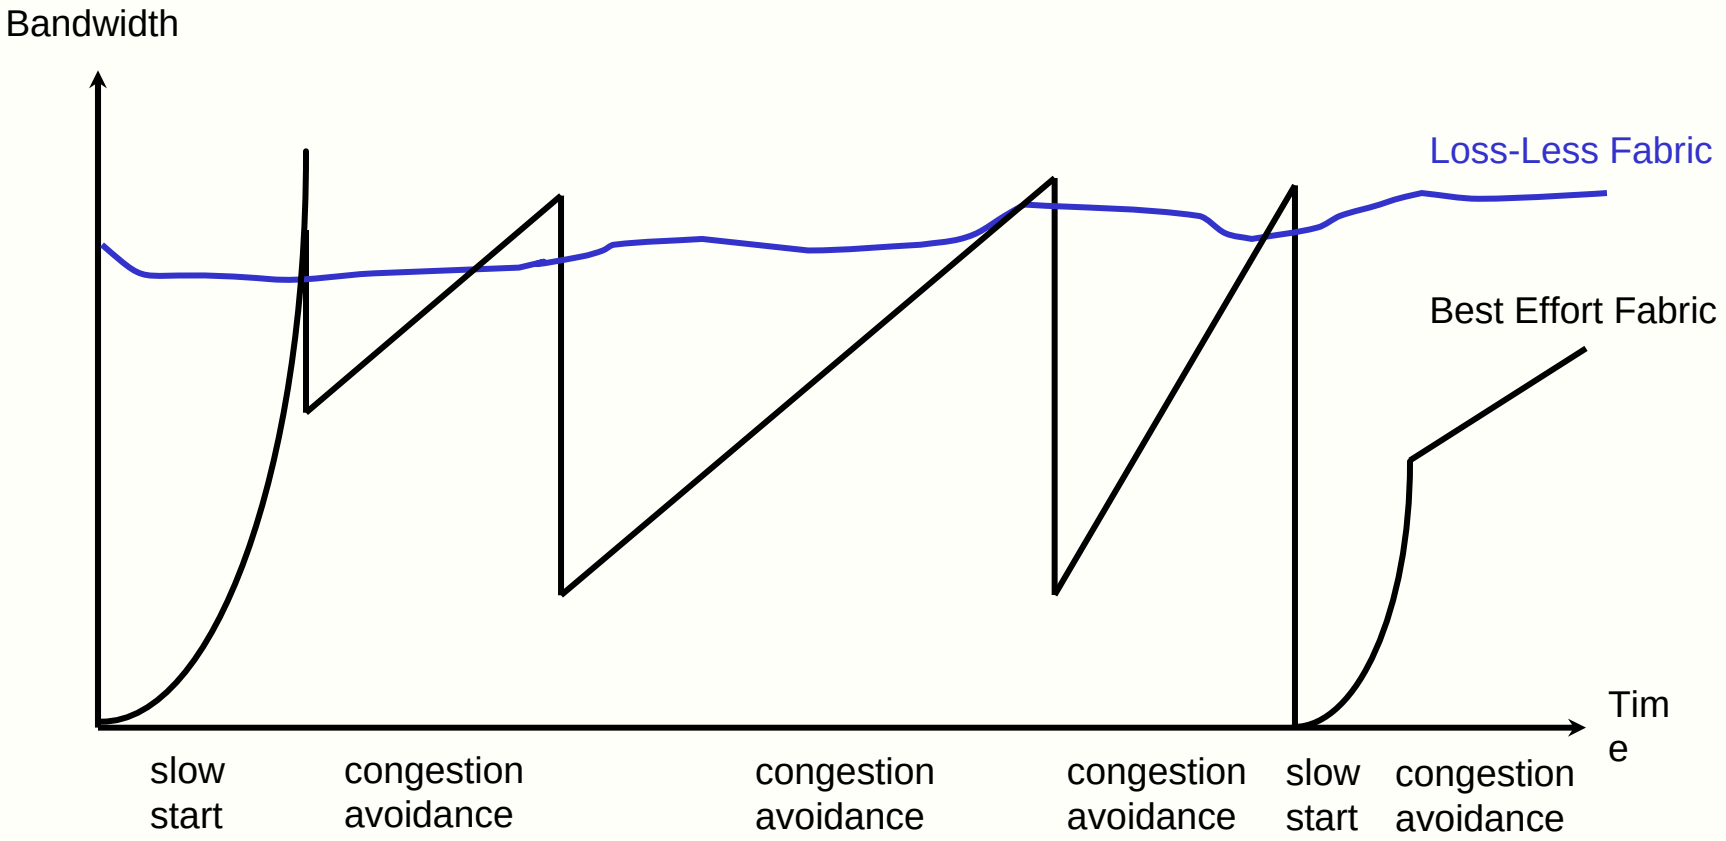

Loss-Less Ethernet increases the performance, reliability and predictability of iSCSI and NAS by eliminating retransmissions

FCoE or iSCSI ?

If you use FC storage today – choose FCoE

- Preserves investment in FC storage, knowledge and management tools
- Requires lossless CEE capable switches - such as BLADE G8124
- Requires CNAs and FCF (Gateways) – both are expensive
- Layer 2

For a new Data Center – choose iSCSI

- iSCSI works well today @ 1Gb and 10Gb
- Wide storage vendor support (NetApp, EMC, HP, IBM)
- Lower cost, NICs and switches
- Much simpler to deploy and manage
- Is routable across Layer 3 networks
- Benefits from lossless Ethernet, but is not dependent on it

What do customers think?

From a recent “Benefits Of SAN/LAN Convergence” study by the Forrester Research – Jan 2010

66% of respondents stated interest in SAN/LAN convergence

Protocol of choice for SAN/LAN convergence ?

Overall results

- 56% selected 10G iSCSI
- 27% selected FCoE
- 17% selected NAS – NFS or CIFS

Result from users with legacy FC SANs

- 44% selected FCoE
- 30% still selected 10G iSCSI
- 26% selected NAS

Ethernet Storage High Level View

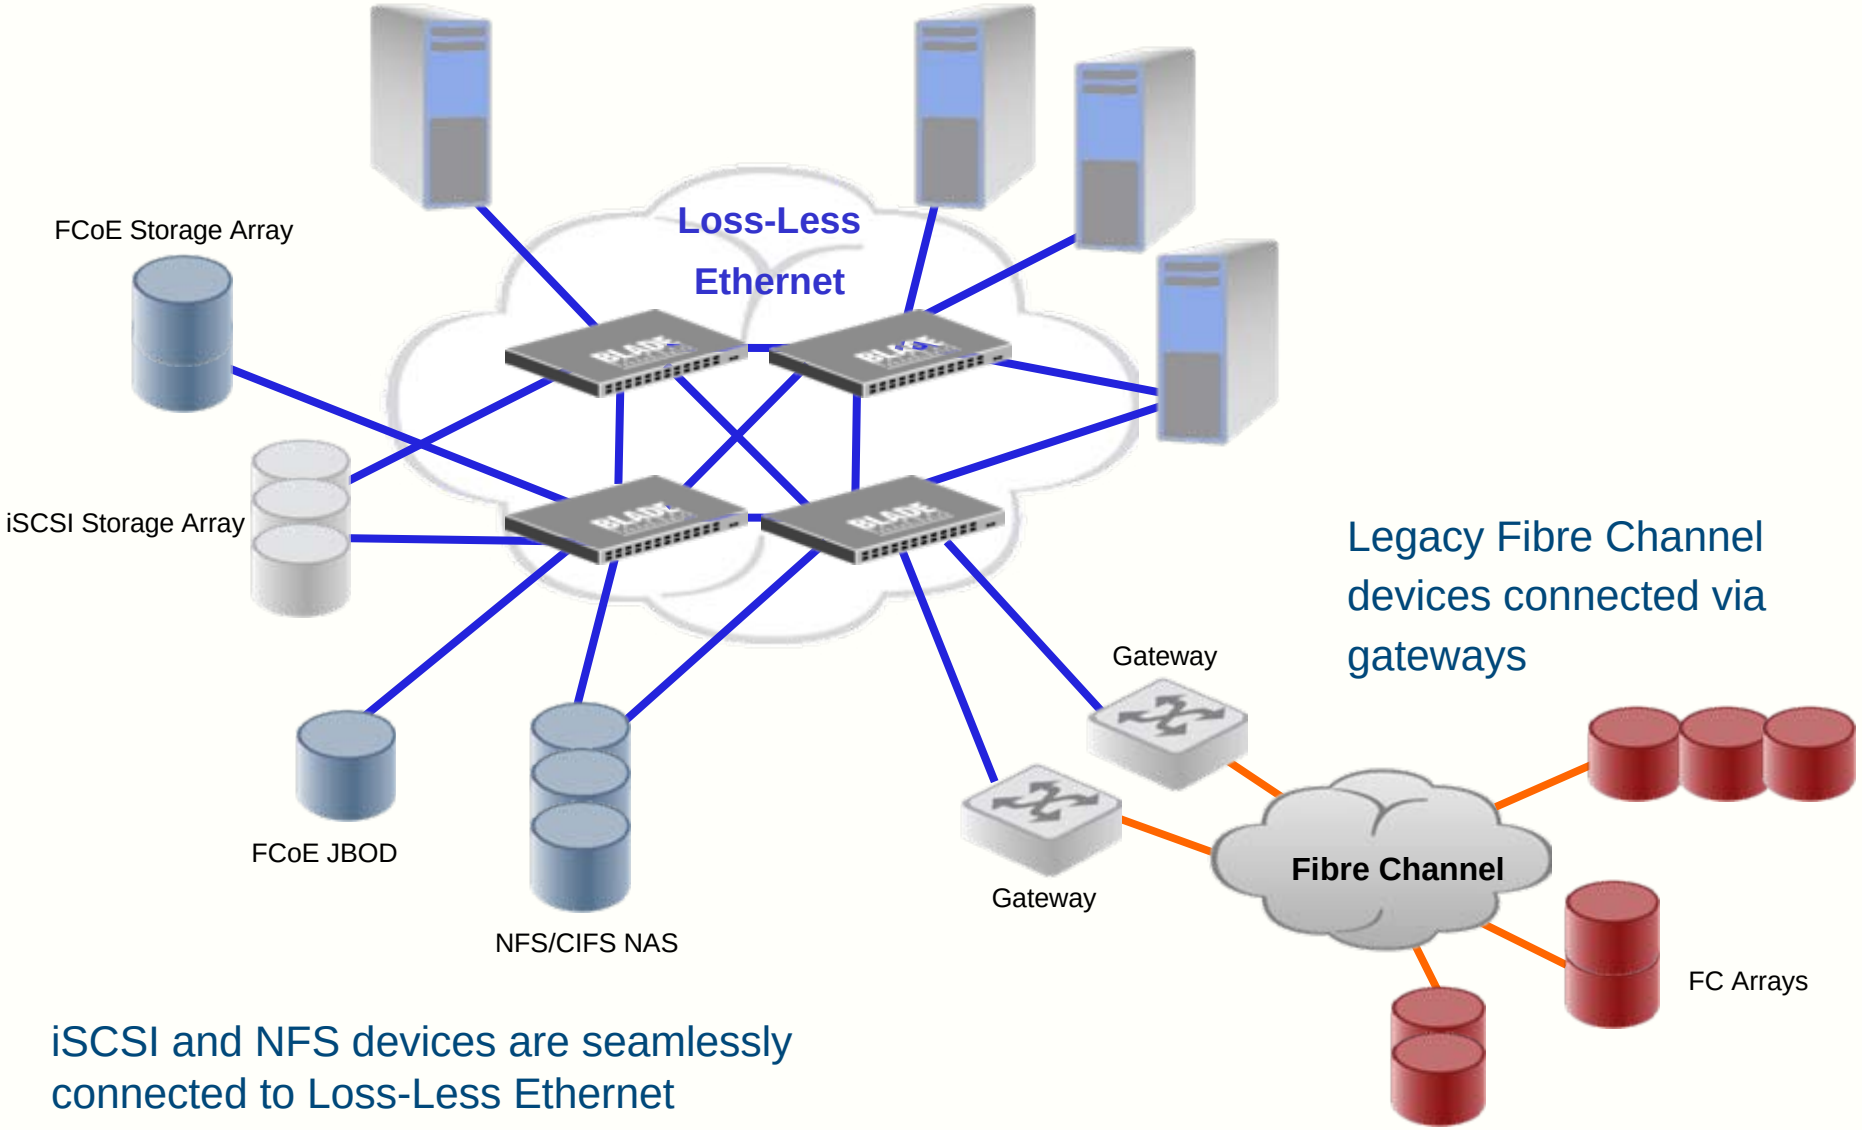

Legacy Fibre Channel devices connected via gateways

iSCSI and NFS devices are seamlessly connected to Loss-Less Ethernet

RoCEE – RDMA over CEE

RDMA (Remote Direct Memory Access) transmits data from the memory of one computer directly to the memory of another

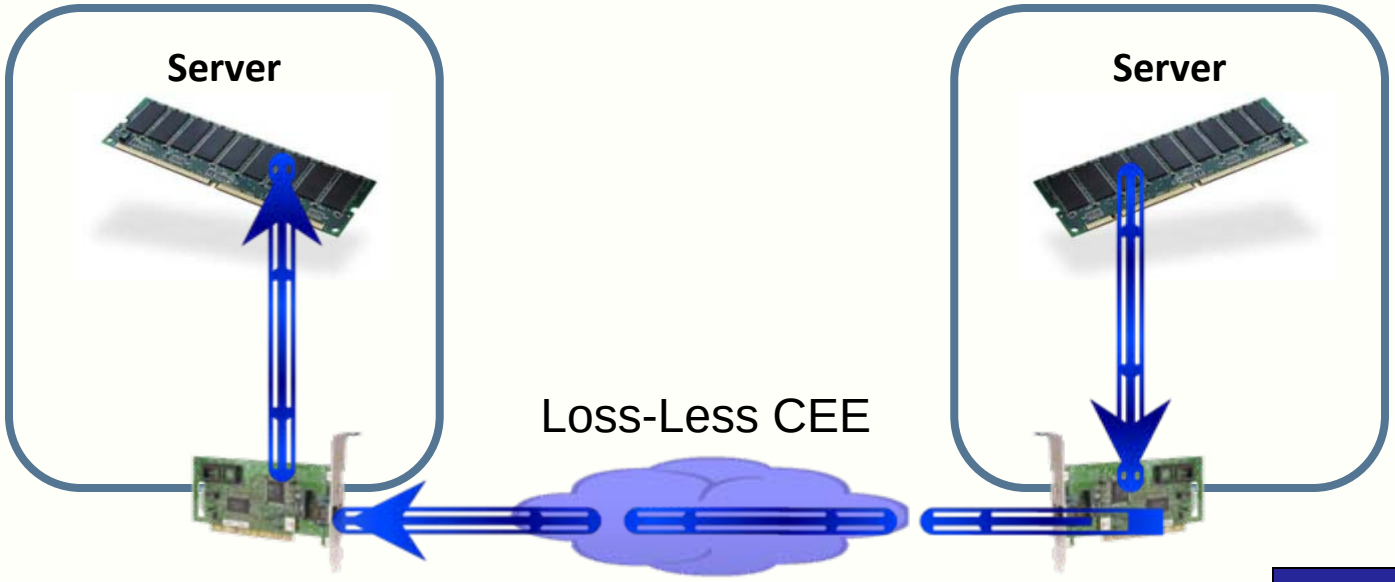

RoCEE is a proposal to run RDMA over standard CEE
Leverage IBTA RDMA transport layer

iWarp is also available for RDMA over tcp/ip

- Lowers latency and CPU & Memory overhead (zero copy)
- Important for High Performance Computing clusters
- Faster storage with FC-RDMA, NFSoRDMA, SRP and iSER

Ethernet Everywhere

BLADE NETWORK TECHNOLOGIES

#1 choice for Networking

#1 choice for Storage

- **NFS / CIFS**
- **iSCSI**
- **FCoE**

#1 choice for HPC

- **>56% of top500 already use Ethernet**
- **RDMA via iWarp and RoCEE**
- **Lower latency NICs and Switch**
- **10G prices <$500 / port and falling fast**

Rest in Peace all others

Disclaimer

- BLADE Network Technologies Confidential.
- The information contained in this presentation is proprietary to BLADE Network Technologies (BLADE) and is offered in confidence, subject to the terms and conditions of a Non-Disclosure Agreement.
- BLADE makes no warranties regarding the accuracy of this information. BLADE does not warrant or represent that it will introduce any product to which the information relates. It is presented for discussion purposes only.
- All statements regarding BLADE's future direction and intent are subject to change or withdrawal without notice, at BLADE's sole discretion.
- Prices, delivery dates and technical specifications are subject to change without notice.
-

BACKUP

Data Center Bridging (DCB) aka CEE

BLADE NETWORK TECHNOLOGIES

Ethernet enhancements geared for data center I/O consolidation driven by IEEE 802.1 DCB working group.

Emerging standards related to DCB:

- **Priority-based Flow Control (IEEE 802.1Qbb)** – enables PAUSE capability on the user priorities that are defined by IEEE 802.1p specification
- **Enhanced Transmission Selection (IEEE 802.1Qaz)** – specifies enhancement of transmission selection to support allocation of bandwidth amongst traffic classes
- **Congestion Notification (IEEE 802.1Qau)** – pushes congestion to the edge of a network
- **DCBX (Data Center Bridging Exchange) Protocol** – A discovery and capability exchange protocol that is used for conveying capabilities and configuration of the above features between neighbors to ensure consistent configuration across the network, leveraging functionality provided by 802.1AB (LLDP).

FCoE & DCB standards – current status

FCoE is defined the FC-BB-5 working group of T11

- On June 3rd 2009, the group completed its work and approved forwarding rev 1.03 of the FC-BB-5 standard to INCITS for further processing as an ANSI standard.

Data Centre Bridging (aka CEE)

- Priority based flow control (IEEE 802.1Qbb)
  - Currently at draft 1.0 (Feb 2009)
  - Draft is stable and is part of CEE rev 0.
- Enhanced Transmission Selection (IEEE 802.1Qaz)
  - Currently at draft 0.3 (June 09)
  - Draft is stable and is part of CEE rev 0.
- Congestion Notification (IEEE 802.1Qau)
  - Currently at draft 2.1 (June 09)
  - Not part of CEE rev 0. - no vendor can provide compliant implementation
- DCBX (extension to LLDP)
  - Stable and is part of CEE rev 0.

Note. CEE rev 0. defines the pre-standard versions of the DCB specifications that multiple vendors have agreed to conform to.

We expect to see products from BLADE, Brocade, Cisco, Qlogic, Emulex, NetApp and EMC in 2H09, conforming to FCoE rev 1.03 + CEE rev.0