



Higher Speed Ethernet and Server Requirements

List of Supporters (1)

Jugnu Ojha	Avago	Ilango Ganga	Intel
Piers Dawe	Avago	Manoj Wadekar	Intel
Yakov Belopolsky	Bel Stewart Conn.	Schelto van Doorn	Intel
Vikram Mehta	Blade Network Techn.	Alan Flatman	LAN Technologies
Ali Ghiasi	Broadcom	Jim Perrott	Luxtera
Howard Frazier	Broadcom	Marek Tlalka	Luxtera
Pat Thaler	Broadcom	Chris DiMinico	MC Communic.
Bob Snively	Brocade	Dan Tuchler	Mellanox
Ori Aruj	Dune Networks	Ed Cady	Meritec
Joe O'Brien	Efficere Technologies	Gourgen Oganessyan	Molex
Hossein Hashemi	Emulex	Jim McGrath	Molex
Bob Thornton	Fujitsu	Brian Pawlowski	Network Appliance
Uri Cummings	Fulcrum Microsys.	Garth Gibson	Panasas
Anthony Befi	IBM	Rob Davis	Qlogic



List of Supporters (2)

Bruce Tolley	Solarflare
George Zimmerman	Solarflare
Andy Bechtolsheim	Sun
Shimon Muller	Sun
Sanjay Kasturia	Teranetics
Dean Vermeersch	Tyco
Mike Fogg	Tyco
Tom Palkert	Xilinx
Nafea Bshara	Marvell
Frank Chang	Vitesse
Arlon Martin	Kotura
Xavier Clairardin	Kotura

40 individuals, 30 companies



Introduction

- This presentation describes server networking requirements in the data center
- We acknowledge the need for 100 Gigabit aggregation links in the network core and backbones
 - IP traffic increases 100% every year
- Servers would greatly benefit from an interim speed solution between 10 Gigabit and 100 Gigabit Ethernet
 - Server/CPU performance increases “only” ~2x every 2 years

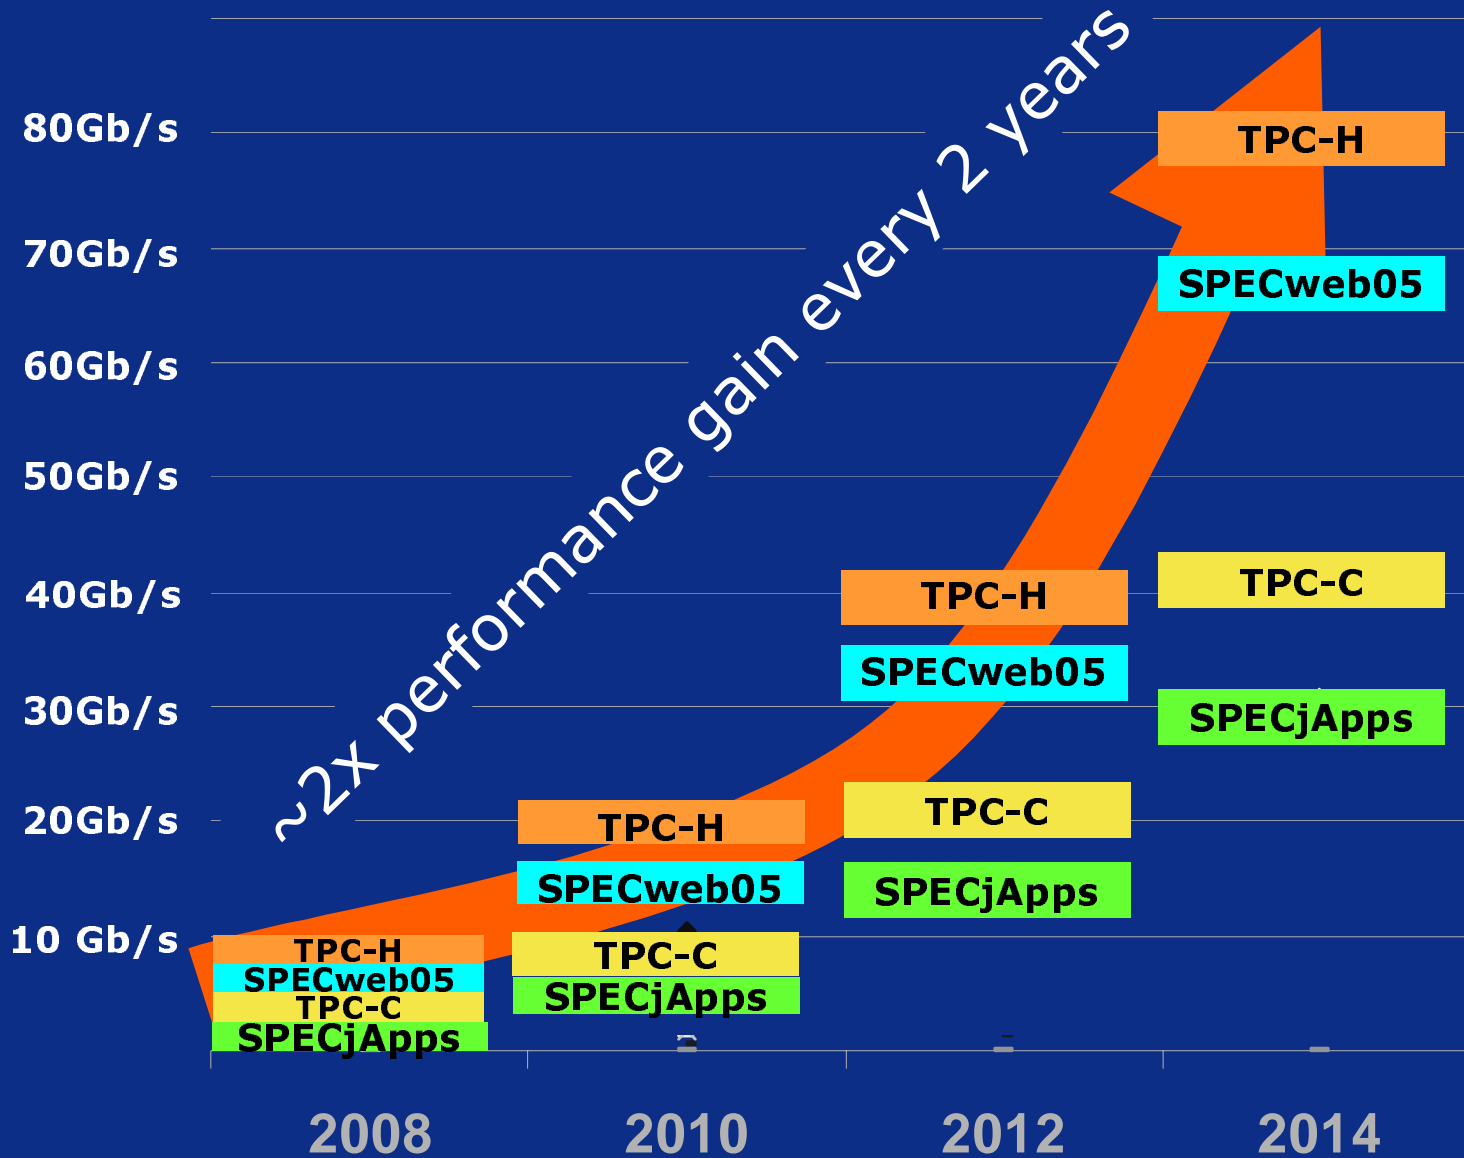


Background

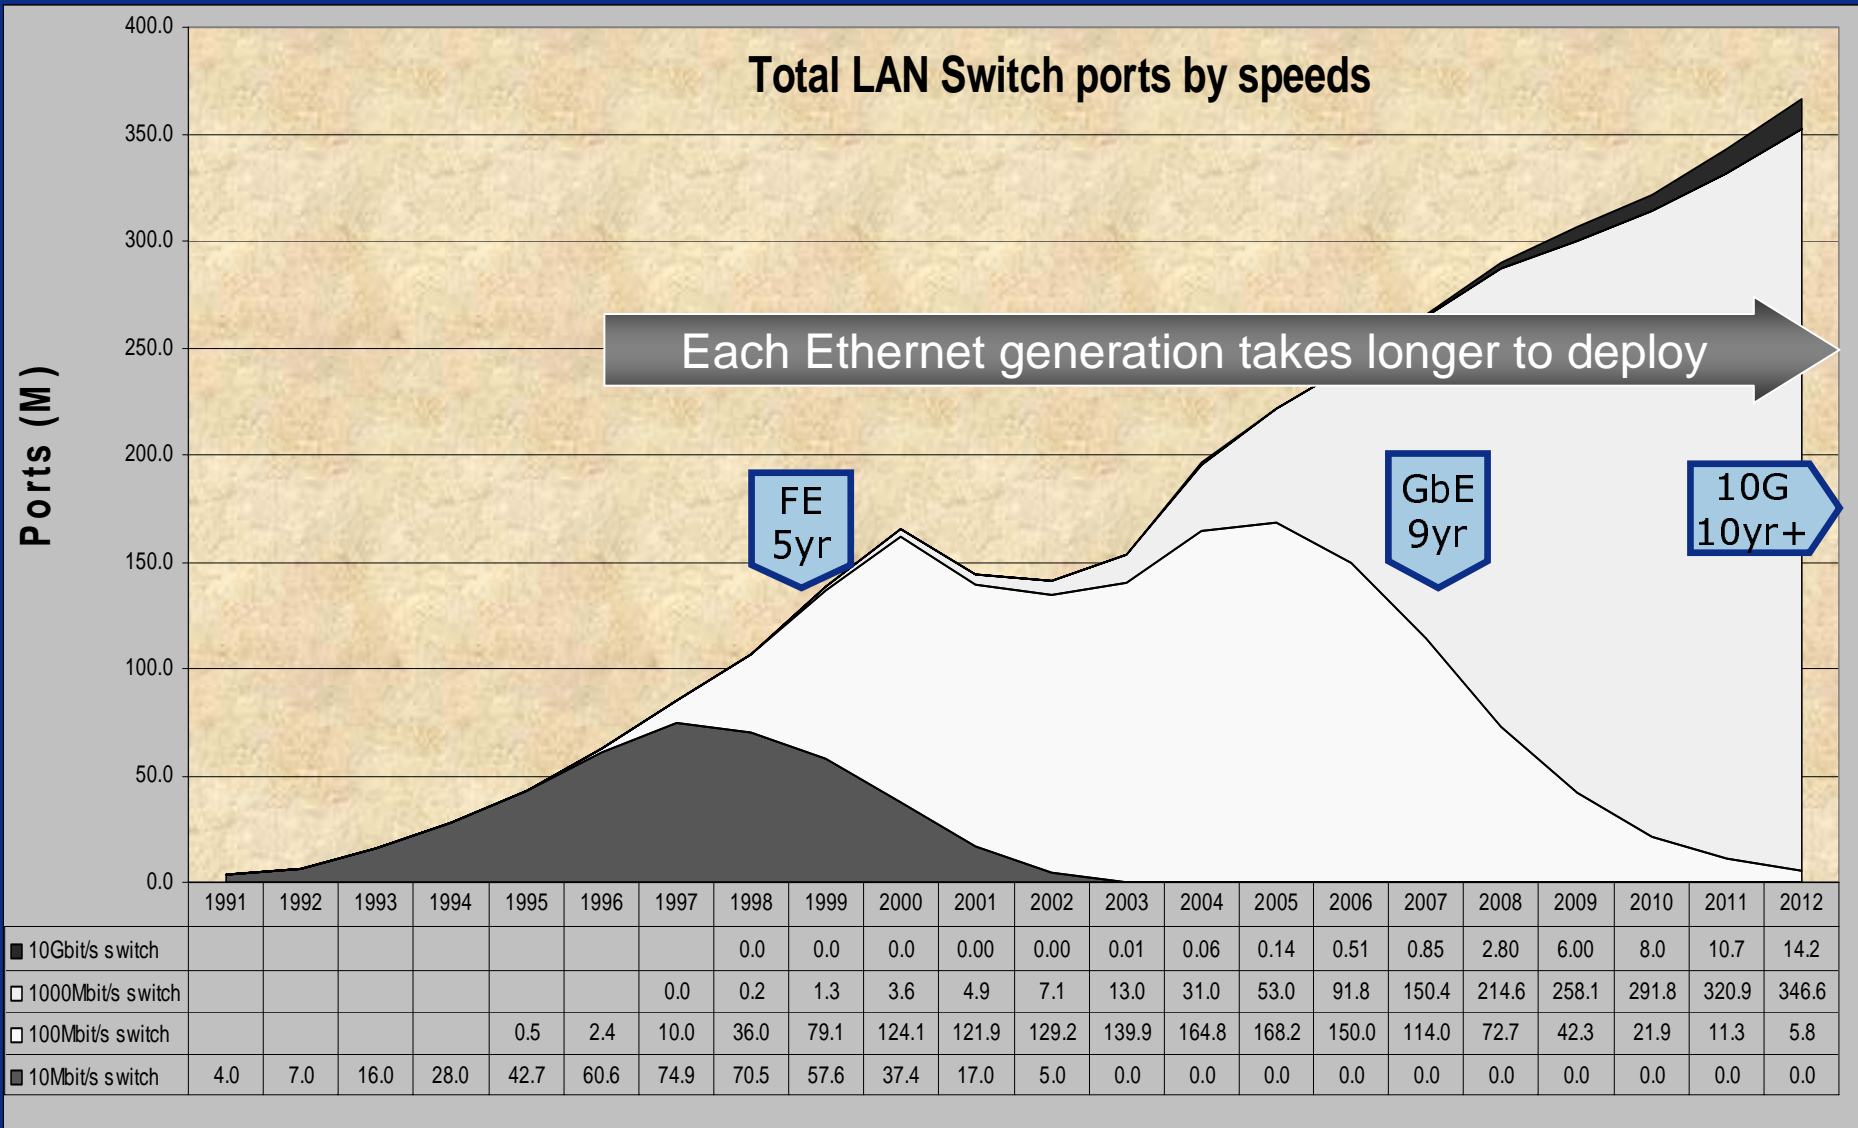
- Ethernet's origin was primarily as an end-station to end-station network
- Fast Ethernet was timed for immediate rollout in higher-end workstations and servers
- Full duplex and switching fully decoupled end-station bandwidth from aggregated bandwidth requirements
- Gigabit Ethernet rollout was initially on aggregated links and servers
- 10 Gigabit Ethernet is still primarily an aggregation technology, years after its standardization, because it is still too expensive for volume servers
- It will be many more years before high performance servers need 100 Gigabit Ethernet



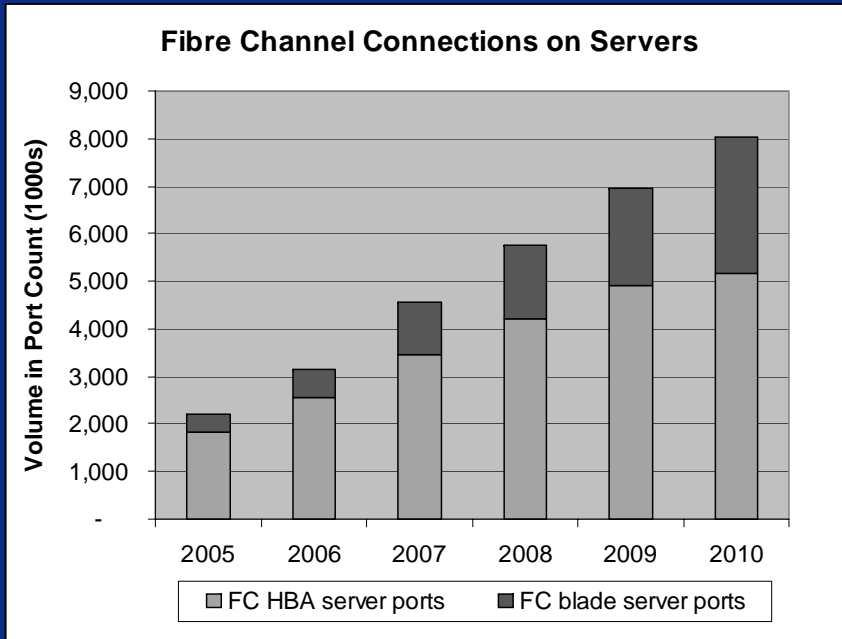
x86 Server I/O Bandwidth Projections



Ethernet Transition Periods are Too Long



Long Gaps in Ethernet Generations Require Alternative Interim Solutions for Servers



Source: Worldwide Fibre Channel Forecast, IDC, Dec'06

- The GbE→10GbE transition is taking longer than the industry planned
- HPC clustering and network storage are driving I/O bandwidth demand
- 10GbE solution cost has been too high, so the server industry has turned to alternatives to fill the gap – e.g. Fibre Channel and Infiniband
- 10Gb→100Gb leap causes us to continue to support multiple I/O types
- Multiple I/O technologies increase cost and complexity of data centers

40 Gb/s could be a compelling speed

- It's easy to implement at NIC level
 - It aligns with the server I/O requirements
 - i.e., any server with a 16x PCIe2 slot will be able to support it
- It's easy to implement at PHY level
 - We can leverage existing PHY technologies
 - 40G components are readily available today
 - QSFP MSA is done
 - Supports active copper cables
 - Supports ribbon short reach fiber,
 - Supports WDM long haul for 4X QDR Infiniband
 - 4 x 10G (XAUI like) for server backplanes

Conclusion

- The Server market would greatly benefit from an Ethernet speed step between 10Gb/s and 100Gb/s, as a complement to 100Gb/s for backbones
- Based on differing bandwidth requirements, it is reasonable to assume that different speed solutions would better serve end-stations and aggregated links in the backbone
- Recommend the HSSG objectives include a more appropriate speed for Servers and end stations.

