

40G Ethernet Market Potential

IEEE 802.3 HSSG Interim Meeting

April 2007

Presenters:

Robert Hays, Intel

Howard Frasier, Broadcom



IEEE 802.3 HSSG April 2007 Interim Meeting




Market Trends Affecting Server Networking


- Opportunity: Server I/O bandwidth generators
 - Moore's Law processing improvements
 - Data center virtualization
 - Networked storage
 - Clustered servers
 - Internet applications (e.g. IPTV, VoIP, Web2.0, Finance)
- Constraints: Purchase criteria for x86 servers
 - Performance (MIPS)
 - Cost (CapEx)
 - Power (Watts)
 - Density (m²)
 - Complexity (OpEx)

Vision: A Single, Flexible Ethernet Server I/O

Separate LAN & SAN Fabrics (Today)



Ethernet LAN & SAN Fabric (Tomorrow)

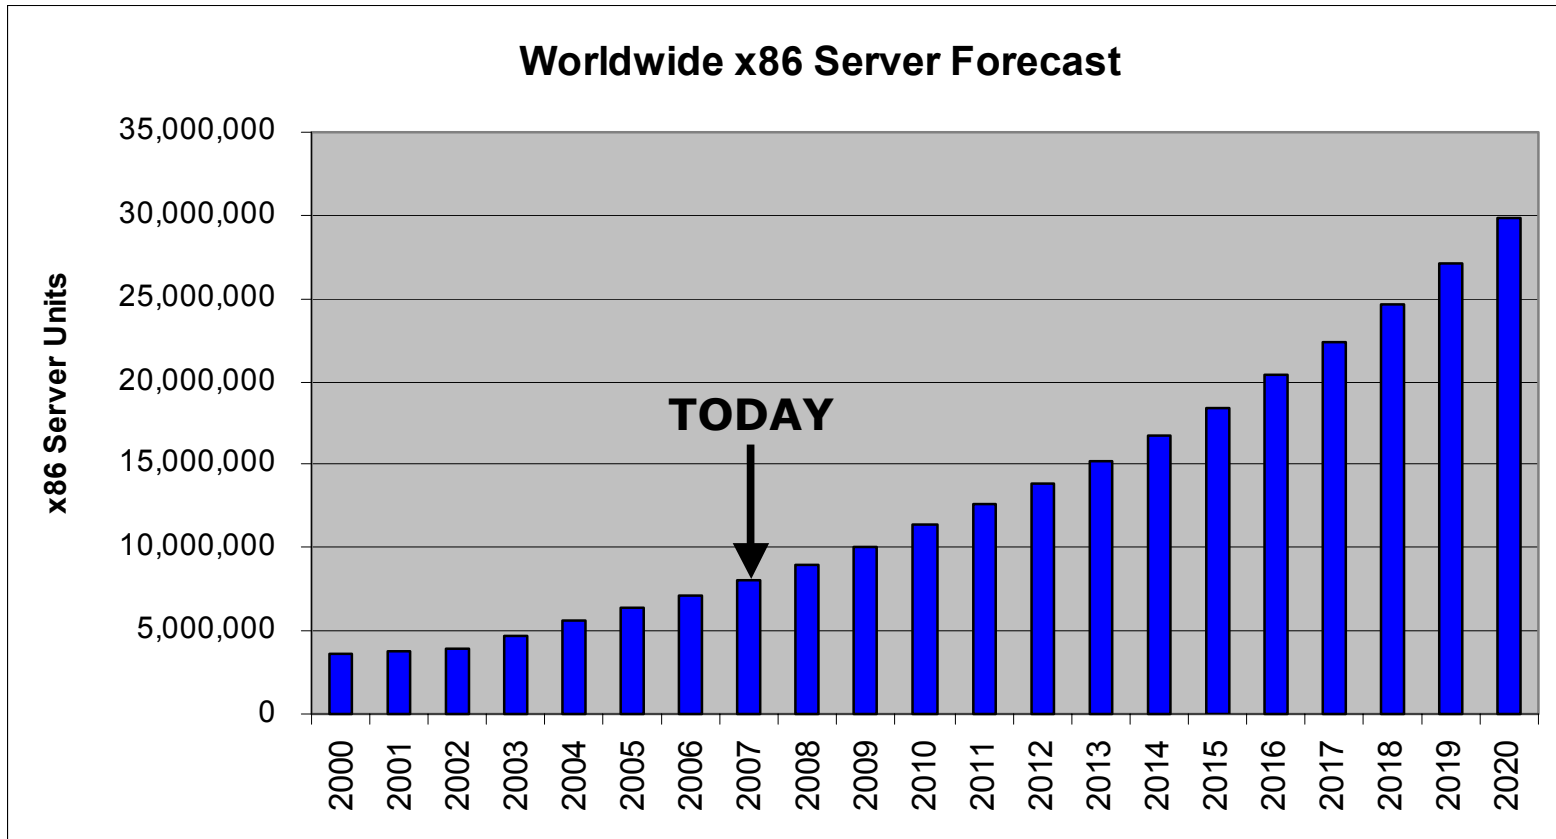


Goal: Reduce the cost, size, power, and complexity of servers by eliminating the need for multiple I/O connections & fabrics

Challenges in Realizing Our Vision

- Today server I/O is fragmented
 - GbE performance doesn't meet current server I/O requirements
 - 10GbE is still too expensive
 - 2/4G Fibre Channel & 10G Infiniband are being deployed to meet the growing I/O demands for storage & HPC clusters
- In 2-4 Years, convergence on 10GbE looks promising
 - 10GBASE-T availability will drive costs down
 - 10GbE will offer a compelling value-proposition for LAN, SAN, & HPC cluster connections
- What about 7-10yrs? Is the Ethernet value-prop sustainable?
 - 10GbE will no longer meet the performance requirements
 - 100GbE may be too expensive for wide spread deployment
 - The market could turn once again to alternative I/O solutions to fill the gap (e.g. 40G Infiniband, shared PCI-Express)

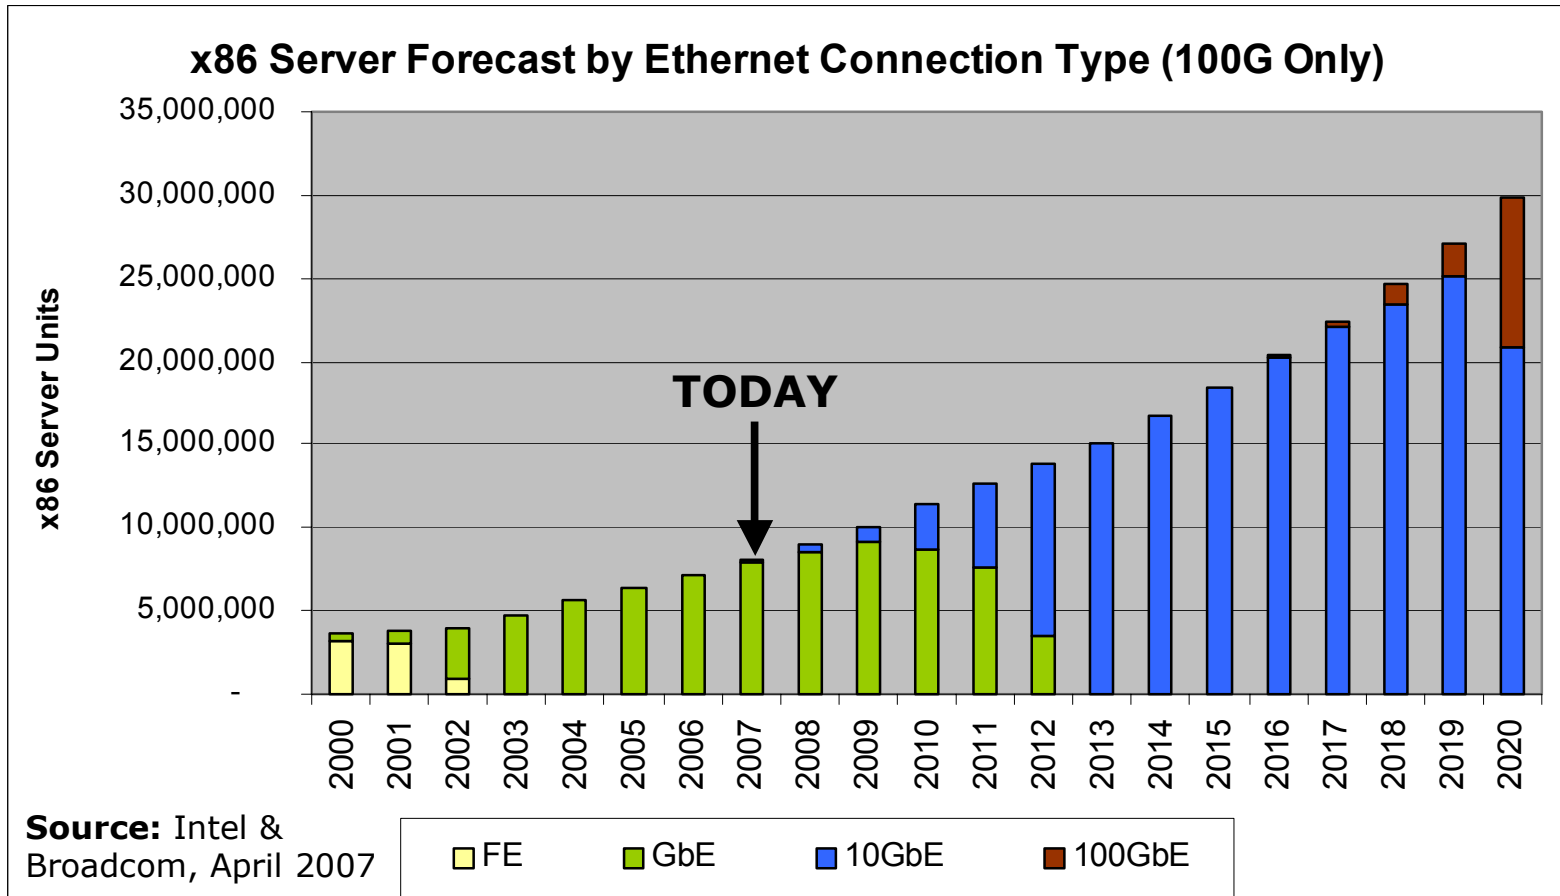
x86 Server Volume Estimate (2000-2020)




Source: IDC 2000-2010 Server Forecast, Sept 2006; 10% CAGR 2011-2018

NOTE: This x86 Server TAM scenario will be used as a baseline to analyze the potential Ethernet server connection market opportunity.

x86 Server Ethernet Connection Speeds with 100GbE Only



x86 Server Ethernet Connection Speeds with 40GbE & 100GbE



Market Potential for 40GbE Servers & Switches

Scenario 1: 40GbE & 100GbE	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	TOTAL
40GbE Server Connections (Kports)	23	51	139	306	1,008	2,218	12,201	31,317	47,736	51,427	47,637	
40GbE Switch Downlinks (Kports)	23	51	139	306	1,008	2,218	12,201	31,317	47,736	51,427	47,637	
40GbE TOTAL (Kports)	46	101	278	611	2,017	4,437	24,403	62,634	95,472	102,854	95,275	388,126
100GbE Switch Uplinks (Kports)	2	4	12	25	84	185	1,017	2,610	3,978	4,286	3,970	
100GbE Server Connections (Kports)	-	-	-	-	-	37	122	447	1,476	2,707	11,909	
100GbE Server Downlinks (Kports)	-	-	-	-	-	37	122	447	1,476	2,707	11,909	
100GbE TOTAL (Kports)	2	4	12	25	84	259	1,261	3,505	6,931	9,699	27,789	49,570

Scenario 2: 100GbE Only	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	TOTAL
100GbE Server Connections (Kports)	-	-	-	-	-	37	163	671	2,215	4,060	17,864	
100GbE Server Downlinks (Kports)	-	-	-	-	-	37	163	671	2,215	4,060	17,864	
100GbE TOTAL (Kports)	-	-	-	-	-	74	325	1,342	4,429	8,120	35,728	50,019

Source: Intel & Broadcom, April 2007

- **>350 Million** port opportunity for 40GbE
 - Assumptions: 2 connections/server, 1:1 server connections to switch downlink ports, 12:1 switch downlinks to uplinks
- 40GbE will drive earlier demand for 100GbE uplinks
- Without 40GbE, 40G IB/FC/PCI-E may displace 10GbE, reducing the long-term 100GbE opportunity in the data center



Summary

- >350M port market potential for 40GbE connections
 - 1st server NICs ~2010, LOM transition ~2016-2017
 - 40GbE increases the demand & TTM for 100GbE aggregation
- 100GbE server deployment will be 4-5 years behind 40GbE
 - 1st server NICs ~2015, LOM transition >2020
- 10X/10yr steps don't keep pace with server I/O needs
 - Cost/performance gaps leave room for alternative technologies to get established and threaten Ethernet as the I/O fabric of choice
 - 40GbE closes the gap between 10G & 100G Ethernet and keeps the Ethernet convergence momentum going

Request: **Modify** the HSSG Objectives to...

- Support full-duplex operation only
- Preserve the 802.3/Ethernet frame format at the MAC Client service interface
- Preserve minimum and maximum FrameSize of current 802.3 Std
- Support a BER better than or equal to $10E-12$ at the MAC/PLS service interface
- Support a speed of 100Gb/s at the MAC/PLS interface
- Provide a family of physical layer specifications for 100Gb/s operation:
 - Support at least TBD meters on Cu cable
 - Support at least 100 meters on OM3 MMF
 - Support at least 10km on SMF
- Support a speed of 40Gb/s at the MAC/PLS interface
- Provide a family of physical layer specifications for 40Gb/s operation:
 - Support at least TBD meters on Cu cable
 - Support at least 100 meters on OM3 MMF

Supporters of a 40GbE Objective

Lars Thon (Aeluros)
Jugnu Ojha (Avago)
Piers Dawe (Avago)
Yakov Belopolsky (Bel Stewart Connector)
Vikram Mehta (Blade Network Technologies)
Ali Ghiasi (Broadcom)
Howard Frazier (Broadcom)
Pat Thaler (Broadcom)
Bob Snively (Brocade)
Ori Aruj (Dune Networks)
Hossein Hashemi (Emulex)
Bob Thornton (Fujitsu)
Uri Cummings (Fulcrum Microsys.)
Anthony Befi (IBM)
Ilango Ganga (Intel)
Manoj Wadekar (Intel)
Shelto Van Doorn (Intel)
Xavier Clairardin (Kotura)
Arlon Martin (Kotura)
Alan Flatman (LAN Techn.)
Hansen Henning (Leoni)

Jim Perrott (Luxtera)
Marek Tlalka (Luxtera)
Nafea Bshara (Marvell)
Chris DiMinico (MC Communic.)
Michael Kagan (Mellanox)
Dan Tuchler (Mellanox)
Ed Cady (Meritec)
Gourgen Oganessyan (Molex)
Jim McGrath (Molex)
David Martin (Nortel Networks)
Garth Gibson (Panasas)
Rob Davis (Qlogic)
Bruce Tolley (Solarflare)
George Zimmerman (Solarflare)
Andy Bechtolsheim (Sun)
Shimon Muller (Sun)
Sanjay Kasturia (Teranetics)
Dean Vermeersch (Tyco)
Mike Fogg (Tyco)
Frank Chang (Vitesse)
Tom Palkert (Xilinx)