

10Gbps Ethernet

10Gbps Ethernet is here...1000 times faster than 10Mbps Ethernet. With the ratification of the IEEE 802.3ae specification for 10Gbps Ethernet, enterprises can now confidently deploy 10 Gigabit Ethernet in their corporate backbones, data centers and server farms to support mission-critical applications.

The popularity of Ethernet has grown and nearly all traffic on the Internet originates or ends with an Ethernet connection. Ethernet can now step up to 10 gigabits per second, however, it remains Ethernet, including the packet format, and the current capabilities are easily transferable to the new standard. 10Gbps Ethernet does not obsolete current investments in network infrastructure.

The 10 Gigabit Ethernet Standard

Under the ISO OSI model, Ethernet is fundamentally a Layer 2 protocol. 10 Gigabit Ethernet uses the IEEE 802.3 MAC protocol, frame format, and minimum, maximum frame size.

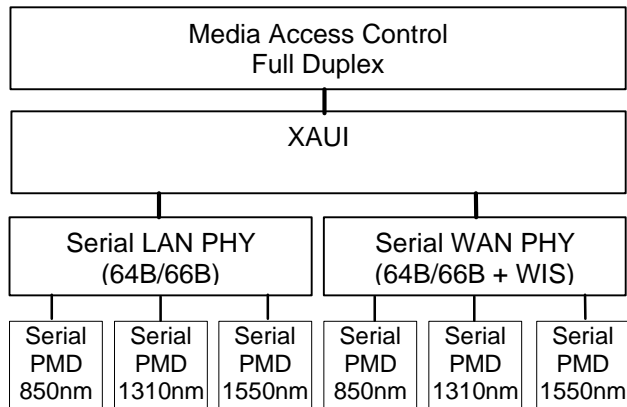
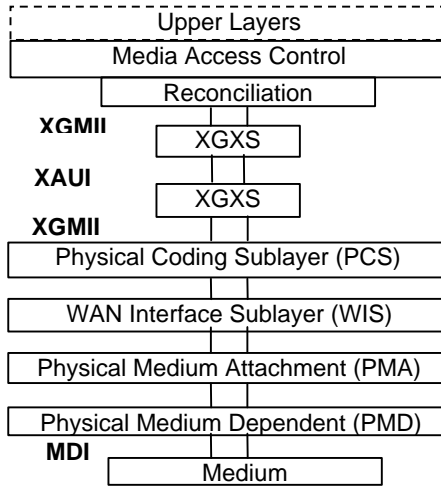


Fig.1 – Architectural components of 802.3ae



An Ethernet PHYSical layer device (**PHY**) connects the media (optical or copper) to the MAC layer. Ethernet architecture further divides the PHY into a Physical Media Dependent (**PMD**) and a Physical Coding Sublayer (**PCS**).

The 802.3ae specification defines two PHY types: the LAN PHY and the WAN PHY. The WAN PHY has an extended feature set added onto the functions of a LAN PHY.

XGMII

Between the MAC and the PHY is the XGMII, or **10Gigabit Media Independent Interface**. XGMII provides full duplex operation at a rate of 10Gbps between the MAC and PHY.

Fiber	62.5 MMF		50 MMF		SMF	
MHz*km	160	200	400	500	2000	-
SR/SW 850 nm	26m	33m	66m	82m	300m	-
LR/LW 1310 nm	-	-	-	-	-	10km
ER/EW 1550 nm	-	-	-	-	-	40km
LX4 1310nm	300m @500MHz*km		240	300m	-	10km

Four PMDs specified in IEEE 802.3ae

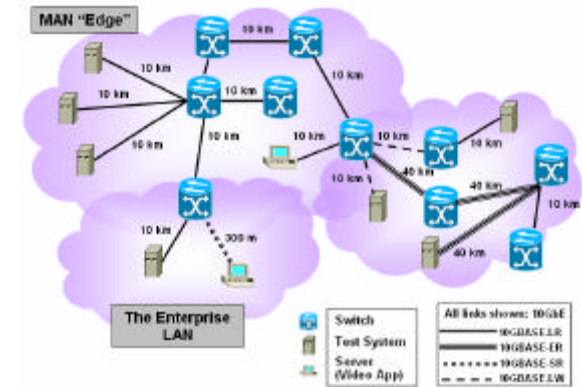
Each direction is independent and contains a 32-bit data path, as well as clock and control signals. In total the interface is 74 bits wide. While XGMII provides a 10Gbps pipeline, the separate transmission of clock and data results in significant challenge in routing the bus more than the recommended short distance of 7cm. For this reason, chip-to-chip, board-to-board and chip-to-optical module applications are not practical with this interface. Consequently, the XGMII bus puts many limitations on the number of ports that may be implemented on a system line card.

XAUI

XAUI (pronounced “zowie”), overcomes these issues. XAUI (**10Gigabit Attachment Unit Interface**) is a full duplex interface that uses four self-clocked serial differential links in each direction to achieve 10Gbps data throughput. Each serial link operates at 3.125Gbps to accommodate both data and the overhead associated with coding.

The self-clocked nature eliminates skew concerns between clock and data, and extends the functional reach of the XGMII by another 50 cm. Thus, the 74 pin wide XGMII interface is reduced to a XAUI interface consisting of 8 differential pair or 16 pins.

Interoperability Demos



World's Largest 10 Gigabit Ethernet Interoperability Demo

One of the keys to Ethernet's success is the widespread interoperability between vendors. The 10 Gigabit Ethernet Alliance (10GEA) was established in order to promote standards based 10 Gigabit Ethernet technology. The 10 GEA hosted the world's largest 10 Gigabit Ethernet Interoperability Network comprised of products from 23 vendors. The end to end 10GbE network was over 200 kilometers long and showcased five PMD port types specified in the IEEE 802.3ae. 12 companies showed chip to chip communication over the IEEE 802.3ae XAUI interface in the demo.

Features of 802.3ae

- ? Full- duplex only (no CSMA/ CD)
- ? Fiber only (802.3ae)
- ? WAN (SONET- friendly) PHY
- ? Mapping to OC- 192 carrier
- ? Rate adaptation to SONET payload capacity
- ? New line code (64b/ 66b)


Applications for 10 Gigabit Ethernet

Similar to Gigabit technology, the 10gps standard supports both Single Mode Fiber (SMF) and Multi Mode Fiber (MMF) mediums. With appropriate interfaces, optical transceivers and single mode fiber, service providers will be able to build links reaching 40km or more in a Network. 10Gbps Ethernet will provide infrastructure for both Network Attached Storage (NAS) and Storage Area Networks (SAN). 10Gbps Enet with WAN PHY will also allow the construction of WANs that connect geographically dispersed LANs between campuses.

iSCSI

Internet Small Computer Systems Interface (iSCSI), is an end-to-end protocol for transporting storage I/O block data over an IP network. The protocol is used on servers (initiators), storage devices (targets), and protocol transfer gateway devices. iSCSI uses standard Ethernet switches and routers to move the data from server to storage.

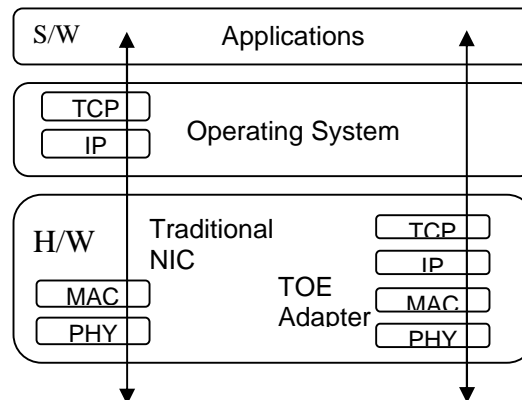
It also enables IP and Ethernet infrastructure to be used for expanding access to SAN storage and extending SAN connectivity across any distance. The technology is based on SCSI commands used in storage traffic today and IP protocols for networking.



Building iSCSI SANs with 10GbE


Applications like synchronous mirroring demand low latency and file serving needs high bandwidth. By using a host bus adapter (HBA), which supports both the network protocols and the iSCSI protocols, both SAN and NAS environments can be consolidated. 10GbE networks facilitate the high band-width and low latency required in this environment, thereby resulting improved application response time.

TCP/IP Offload Engine (TOE)



TCP/IP offload Engine (TOE) is one of the technologies that can reduce the amount of TCP/IP processing handled by CPU and server I/O subsystem, and thus ease server networking bottleneck. The processing of TCP/IP over Ethernet is traditionally accomplished by software running on the CPU of the server. TOE is emerging as a solution to

limit the processing required by CPUs for network links. A TOE may be embedded in a Network Card or Host Bus Adapter (HBA). Deployment of TCP/IP offload in conjunction with high-speed Ethernet technologies enables applications to take full advantage of the networking capabilities.



• Intel Pro/10GbE LR server adapter card based on Intel® 82597EX 10 Gigabit Ethernet Controller

Closing Thoughts

In the future, network managers will be able to use 10 Gigabit Ethernet as a corner stone for network architectures that encompass LANs, MANs and WANs using Ethernet as the end to end, Layer 2 transport method. IEEE started working on 10GBaseT standard – which enables 10Gbps Ethernet over UTP Cables. Ethernet bandwidth can then be scaled from 10Mbps to 10Gbps – a ratio of 1 to 1000 — without compromising intelligent network services such as Layer 3 routing and layer 4 to layer 7 intelligence.

All Trademarks belong to the respective companies

