

1. **10 Gigabit Ethernet Standard** (or 10GE or 10GbE or 10 GigE) was first published in 2002 as IEEE standard 802.3ae-2002 and is the fastest of the Ethernet standards. It defines a version of Ethernet with a nominal data rate of 10 Gbit/s, 10 times faster than Gigabit Ethernet.
2. **10 Gigabit Ethernet over UTP (802.3an)** (See “10GBase-T”)
3. **10GBASE-CX4** was the first 10G copper standard published by 802.3 (as 802.3ak-2004). It uses the XAUI 4-lane PCS (Clause 48) and copper cabling similar to that used by InfiniBand technology. It is specified to work up to a distance of 15 m (49 ft). Each lane carries 3.125 Gbaud of signaling bandwidth. Most short reach copper implementations today use SFP+ Direct Attach rather than 10GBASE-CX4.
4. **10GBASE-KX4 and 10GBASE-KR** (see “Backplane Ethernet”)
5. **10GBASE-LRM** (Long Reach Multimode) also known as 802.3aq uses the IEEE 802.3 Clause 49 64B/66B Physical Coding Sublayer (PCS) and 1310 nm lasers. It delivers serialized data over multi-mode fiber at a line rate of 10.3125 Gbit/s. 10GBASE-LRM is designed to achieve longer distances over FDDI grade optical cable (OM1) within the data center. There has not been strong adoption of 10GBASE-LRM.
6. **10GBASE-SR** (SFP+ SR “Short Range”) uses the IEEE 802.3 Clause 49 64B/66B Physical Coding Sublayer (PCS) and 850 nm lasers. It delivers serialized data over multi-mode fiber at a line rate of 10.3125 Gbit/s. This is the most common optical PHY used in the data center. Emulex “M” models support SR optics.
7. **10GBase-T** (IEEE 802.3an-2006), is a standard released in 2006 to provide 10 gbit/s connections over unshielded or shielded twisted pair cables, over distances up to 100 meters (330 ft). 10GBASE-T cable infrastructure can also be used for 1000BASE-T allowing a gradual upgrade from 1000BASE-T using auto-negotiation to select which speed to use. 10GBASE-T has higher latency and consumes more power than other 10GbE physical layers. In 2008 10GBASE-T silicon is now available from several manufacturers with claimed power dissipation of 6W and a latency approaching 1 microsecond. 10GBase-T is expected to become common in rack servers starting in 2H11.
8. **10GSFP+Cu** (SFP+ Direct Attach) is a copper interconnect using a passive twin-ax cable assembly that connects directly into an SFP+ housing. It has a range of 10 meters and like 10GBASE-CX4, is low power, low cost and low latency with the added advantage of having the small form factor of SFP+, and smaller, more flexible cabling. Emulex “X” models support SFP+ Direct Attach.
9. **Backplane Ethernet** , also known by its working group name IEEE 802.3ap, is used in backplane applications such as blade servers and routers/switches with upgradable line cards. 802.3ap implementations are required to operate in an environment comprising up to 1 meter (39 in) of copper-printed circuit board with two connectors. The standard provides for two different implementations at 10Gbit/s: 10GBASE-KX4 and 10GBASE-KR. 10Gbase-KX4 uses the same physical layer coding (defined in IEEE 802.3 Clause 48) as 10GBASE-CX4. 10GBASE-KR uses the same coding (defined in IEEE 802.3 Clause 49) as 10GBASE-LR/ER/SR. The 802.3ap standard also defines an optional layer for FEC, a backplane auto-negotiation protocol and link training, where the receiver can set a three-tap transmit equalizer. Blade servers from Cisco and HP use 10GBASE-KR; blade servers from IBM and Dell use 10GBASE-KX4.

10. **Bridges (L2 Switches)** involves segmentation of local area networks (LANs) at the Layer 2 level. A multiport bridge typically learns about the Media Access Control (MAC) addresses on each of its ports and transparently passes MAC frames destined to those ports. These bridges also ensure that frames destined for MAC addresses that lie on the same port as the originating station are not forwarded to the other ports.
11. **Broadcast Packet** means that the network delivers one copy of a packet to each destination. On bus technologies like Ethernet, broadcast delivery can be accomplished with a single packet transmission. On networks composed of switches with point-to-point connections, software must implement broadcasting by forwarding copies of the packet across individual connections until all switches have received a copy.
12. **Checksum (CRC)**, or (cyclic redundancy check), is a non-secure hash function designed to detect accidental changes to raw computer data, and is commonly used in digital networks and storage devices. A CRC is a "digital signature" representing data. The most common CRC is CRC32, in which the "digital signature" is a 32-bit number. FCoE packets contain a Fibre Channel checksum and Ethernet CRC. TCP/IP packets contain a TCP Checksum and Ethernet CRC. iSCSI packets optionally contain an iSCSI digest (CRC), TCP Checksum and Ethernet CRC. Offload engines and stateless offloads remove the checksum offload overhead from the host CPU.
13. **Common Internet File System (CIFS)** is a remote file system access protocol that works over IP networks to enable groups of users to work together and share documents across LANs or WANs. CIFS is an open, cross-platform technology based on the native file-sharing protocols built into the Microsoft Windows operating systems, and is also supported on other platforms.
14. **Congestion Control** for TCP uses a number of mechanisms to achieve high performance and avoid "congestion collapse," where network performance can fall by several orders of magnitude. These mechanisms control the rate of data entering the network, keeping the data flow below a rate that would trigger collapse.
15. **Congestion Notification** (IEEE 802.1Qau) provides end-to-end congestion management for protocols that are capable of transmission rates limiting to avoid frame loss. It is expected to benefit protocols such as TCP that do have native congestion management as it reacts to congestion in a timelier manner.
16. **Congestion Management (CM)** (P802.1Qau) (See "Congestion Notification")
17. **Data Center Bridging Capability eXchange Protocol (DCBX)** is responsible for configuration of link parameters for DCB function. It includes a protocol to exchange (send and receive) DCB parameters between peers, set local "operational" parameters based received DCB parameters, and resolve conflicting parameters.
18. **Direct Data Placement Protocol (DDP)**, is the main component of iWARP, which permits the actual zero-copy transmission. DDP itself does not perform the transmission; TCP does.
19. **Energy Efficient Ethernet (EEE)** is the IEEE 802.3 standard to define a mechanism to reduce power consumption during periods of low link utilization for the following PHYs: 100BASE-TX (Full Duplex), 1000BASE-T (Full Duplex), 10GBASE-T, 10GBASE-KR, 10GBASE-KX4.

20. **Enhanced Ethernet (EE)**, also known as Converged Enhanced Ethernet (CEE), is a generic term used by many vendors including HP, IBM, Dell, Brocade, and others to describe enhanced Ethernet. Data Center Ethernet (DCE) was a term originally coined and trademarked by Cisco. DCE refers to enhanced Ethernet based on the Data Center Bridging (DCB) standards, and also includes a Layer 2 Multipathing implementation based on the IETF's Transparent Interconnection of Lots of Links (TRILL) proposal. These terms generally refer to the collection of Priority Flow Control, Enhanced Transmission Selection and Data Center Bridging Capabilities Exchange Protocols.
21. **Enhanced Transmission Selection (ETS)** is the P802.1Qaz standard that specifies enhancement of transmission selection to support allocation of bandwidth among traffic classes. When the offered load in a traffic class doesn't use its allocated bandwidth, enhanced transmission selection will allow other traffic classes to use the available bandwidth. The bandwidth allocation priorities will coexist with strict priorities. It will include managed objects to support bandwidth allocation.
22. **Fibre Channel over Ethernet (FCoE)** is the encapsulation of the Fibre Channel protocol into Ethernet as defined by the INCITS T11 standards organization. This allows Fibre Channel traffic to coexist with TCP/IP traffic using a common adapter and network infrastructure.
23. **Gigabit Ethernet (GbE or 1 GigE)** is a term describing various technologies for transmitting Ethernet frames at a rate of a gigabit per second, as defined by the IEEE 802.3-2008 standard.
24. **Internet Protocol (IP)** is a protocol used for communicating data across a packet-switched internetwork using the Internet Protocol Suite, also referred to as TCP/IP. IP is the primary protocol in the Internet Layer of the Internet Protocol Suite and has the task of delivering distinguished protocol datagrams (packets) from the source host to the destination host solely based on their addresses.
25. **IP Multicast** is a technique for one-to-many communication over an IP infrastructure in a network. It scales to a larger receiver population by not requiring prior knowledge of who or how many receivers there are. Multicast uses network infrastructure efficiently by requiring the source to send a packet only once, even if it needs to be delivered to a large number of receivers.
26. **iSCSI** is the storage networking standard developed by the Internet Engineering Task Force (IETF) for linking data storage over an IP-based network.
27. **Internet Storage Name Service (iSNS)** protocol allows automated discovery, management and configuration of iSCSI devices on a TCP/IP network.
28. **iSCSI Extensions for RDMA (iSER)** protocol maps the iSCSI protocol over a network that provides RDMA services (like TCP with RDMA services (iWARP or InfiniBand). This permits data to be transferred directly into SCSI I/O buffers without intermediate data copies.

- 29. iWARP (The Internet Wide Area RDMA Protocol)** is an Internet Engineering Task Force (IETF) update of the RDMA Consortium's RDMA over TCP standard. iWARP is a superset of the Virtual Interface Architecture that permits zero-copy transmission over legacy TCP. It may be thought of as the features of InfiniBand (IB) applied to Ethernet.
- 30. Jumbo Frames** are Ethernet frames with more than 1,500 bytes of payload (MTU). Conventionally, jumbo frames can carry up to 9,000 bytes of payload, but variations exist and some care must be taken when using the term. Many, but not all, Gigabit Ethernet switches and Gigabit Ethernet network interface cards support jumbo frames, but all Fast Ethernet switches and Fast Ethernet network interface cards support only standard-sized frames.
- 31. Large Send Offload (LSO)** is a technique for increasing outbound throughput of high-bandwidth network connections by reducing CPU overhead. It works by queuing up large buffers and letting the NIC split them into separate packets. The technique is also called TCP Segmentation Offload (TSO) when applied to TCP, or Generic Segmentation Offload (GSO).
- 32. Large Receive Offload (LRO)** is a technique for increasing inbound throughput of high-bandwidth network connections by reducing CPU overhead. It works by aggregating multiple incoming packets from a single stream into a larger packet buffer before they are passed higher up the networking stack, thus reducing the number of packets that have to be processed.
- 33. Lossless Ethernet** fabrics are enabled by using priority-based flow control (PFC) to pause traffic based on priority levels. This allows virtual lanes to be created within an Ethernet link, with each virtual lane assigned a priority level. During periods of heavy congestion, lower priority traffic can be paused, while allowing high-priority and latency-sensitive tasks such as data storage to continue.
- 34. Marker PDU Aligned Framing for TCP (MPA)** is required to run Direct Data Placement Protocol (DDP) over TCP to guarantee boundaries of messages.
- 35. Microsoft Chimney (TCP Offload)** architecture offloads the data-transfer portion of TCP protocol processing for one or more TCP connections to a network interface card (NIC). This architecture provides a direct connection, called a *chimney*, between applications and an offload-capable NIC.
- 36. Multi-Source Agreements (MSAs)** 10GbE Optical modules are not specified in IEEE 802.3 but by multi-source agreements (MSAs). The relevant MSAs for 10GbE are XENPAK, X2, XPAK, XFP and SFP+. The latest, smallest and lowest power is SFP+ (SFF-8431). Emulex OneConnect™ stand-up PCI adapters use SFP+ modules.
- 37. NetQueue Offload** is a performance technology that significantly improves performance in VMware ESX server deployments by queuing data to multiple receive queues, generally tied to VMs running under ESX. MSI-X is then used to signal to the specific queue being used.
- 38. Network File System (NFS)** is a distributed file system that allows a system to share directories and files with other systems over a network. NFS is most commonly used with Linux and Unix systems.

39. **Network Interface Controller (NIC)** is a hardware device that handles an interface to a computer network and allows a network-capable device to access that network. The NIC has a ROM chip that contains a unique number, the multiple access control (MAC) Address that is permanent. The MAC address identifies the device uniquely on the LAN. The NIC exists on both the 'Physical Layer' (Layer 1) and the 'Data Link Layer' (Layer 2) of the OSI model.
40. **Partial Offload** (see “Microsoft Chimney”)
41. **PCI Express (Peripheral Component Interconnect Express)**, abbreviated as **PCIe** or **PCI-E**, is a computer expansion card standard designed to replace the older PCI, PCI-X, and AGP standards. PCI Express is used as a motherboard-level interconnect (to link motherboard-mounted peripherals) and as an expansion card interface for add-in boards.
42. **Priority-based Flow Control (PFC)**, IEEE 802.1Qbb, provides a link level flow control mechanism that can be controlled independently for each Class of Service (CoS), as defined by 802.1p. The goal of this mechanism is to ensure zero loss under congestion in DCB networks.
43. **Remote Direct Memory Access Protocol (RDMA or RDMAP)** is a direct memory access from the memory of one computer into that of another without involving either one's operating system. This permits high-throughput, low-latency networking, which is especially useful in massively parallel computer clusters.
44. **Receive-Side Scaling (RSS)** is a technology that enables packet receive-processing to scale with the number of available computer processors, by dynamically load-balancing inbound network connections across multiple processors or cores.
45. **Retransmission** is the resending of packets which have been either damaged or lost. It is a term that refers to one of the basic mechanisms used by protocols operating over a packet switched computer network to provide reliable communication (such as that provided by a reliable byte stream, for example, TCP).
46. **SACK processing** stands for Selective Acknowledgement and is an advanced TCP feature. With SACK, the receiver explicitly lists which packets, messages, or segments in a stream are acknowledged (either negatively or positively). Positive selective acknowledgment is an option in TCP.
47. **SFP+ Direct Attach** (See “10GSFP+Cu”)
48. **SFP+ SR “Short Reach”** (See “10GBASE-SR”)
49. **Single Root I/O Virtualization (SR-IOV)** allows a PCIe device to appear to be multiple separate PCIe devices. The SR-IOV specification includes physical functions (PFs) and virtual functions (VFs). PFs are full-featured PCIe functions. VFs are lightweight functions that do not support configuration resources. With SR-IOV, virtual machines can share adapter ports using virtual functions to optimize performance.

50. **Spanning Tree Protocol (STP)** is defined in the IEEE Standard 802.1D and creates a spanning tree within a mesh network of connected layer-2 bridges (typically Ethernet switches), and disables those links that are not part of the tree, leaving a single active path between any two network nodes.
51. **Stateful Offload** describes the type of offload done by a TCP Offload Engine, where connection state is passed to the device.
52. **Stateless Offloads** do not require state to be stored in the offload engine. These are operations like checksum offload, LSO, and LRO.
53. **TCB (Transport Control Block)** is a data structure in a TCP connection that contains information about the connection state, its associated local process, and feedback parameters about the connection's transmission properties.
54. **TCP Offload Engine or TOE**, is a technology used in network interface cards (NICs) to offload processing of the entire TCP/IP stack to the network controller. It is primarily used with high-speed network interfaces, such as gigabit Ethernet and 10GbE, where processing overhead of the network stack becomes significant. The term TOE is often used to refer to the NIC itself, although it more accurately refers only to the integrated circuit included on the card which processes the TCP headers. TOEs are used to reduce the overhead associated with protocols like iSCSI.
55. **SFP+ form factor (SFF-8431)** is a specification for compact, hot-pluggable transceiver that interfaces a device to a fiber optic or copper networking cable. SFP transceivers are designed to support Gigabit Ethernet and Fibre Channel and have expanded to SFP+ to support data rates up to 10.0 Gbit/s (including data rates for 8 gigabit Fibre Channel, and 10GbE.)
56. **User Datagram Protocol (UDP)**, sometimes called the Universal Datagram Protocol, is one of the core members of the Internet Protocol Suite, the set of network protocols used for the Internet. With UDP, computer applications can send messages, in this case referred to as datagrams, to other hosts on an Internet Protocol (IP) network without requiring prior communications to set up special transmission channels or data paths.
57. **Unicast transmission** is the sending of information packets to a single network destination. The term "unicast" is formed in analogy to the word "broadcast" which means transmitting the same data to all destinations.
58. **VLAN (virtual LAN)** is a group of hosts with a common set of requirements that communicate as if they were attached to the Broadcast domain, regardless of their physical location. A VLAN has the same attributes as a physical LAN, but it allows for end stations to be grouped together even if they are not located on the same network switch.