High Speed Ethernet

Dr. Sanjay P. Ahuja, Ph.D.
Professor
School of Computing, UNF

Hubs and Switches

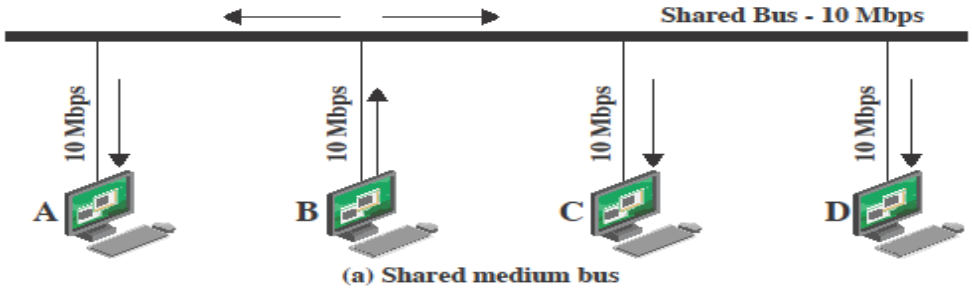

(a) Shared medium bus

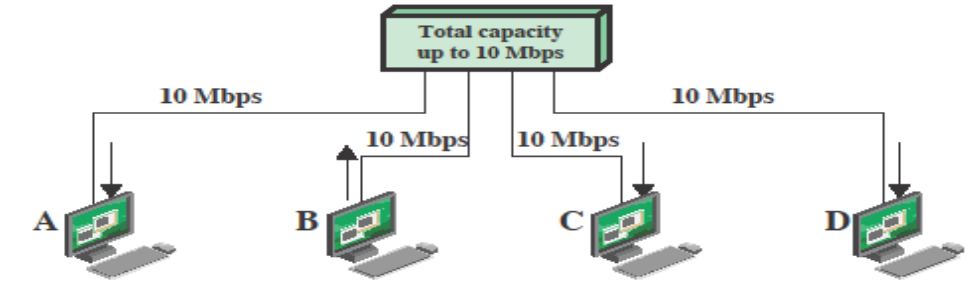

(b) Shared medium hub

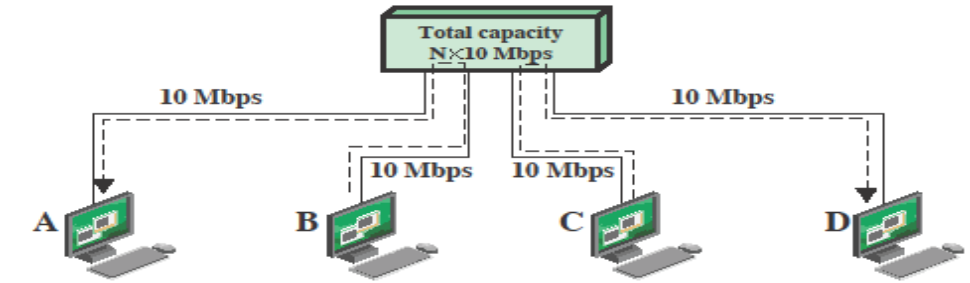

(c) Layer 2 switch

Figure 15.12 LAN Hubs and Switches

Hubs and Switches …

Shared Medium Hub

- The total capacity in the shared medium hub configuration (figure 15.12 b) is the same as that of shared bus (figure 15.12 a).
- Advantages:
  - Exploits standard building wiring practices in cable layout.
  - Maintenance is easier.
  - Uses UTP.

Switches

- Advantages:
  - No change required to software or hardware of attached devices to convert a bus LAN to a switched LAN.
  - Each attached device has a dedicated capacity equal to that of the entire original LAN.
  - Switch scales easily.

LAN Protocols

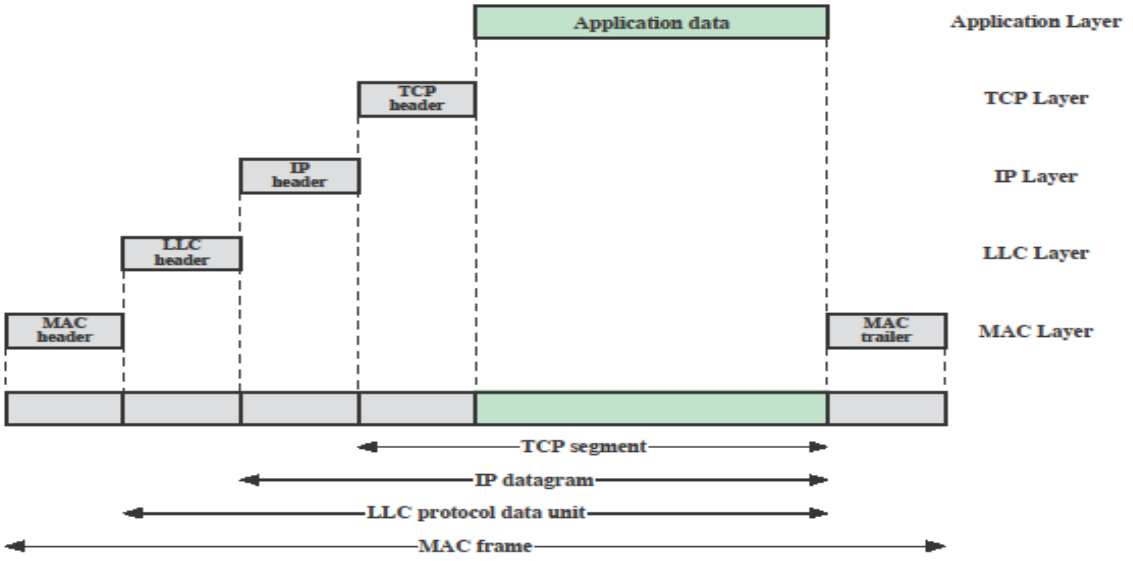

Figure 15.5 LAN Protocols in Context

IEEE 802.3 (Ethernet) Frame Format

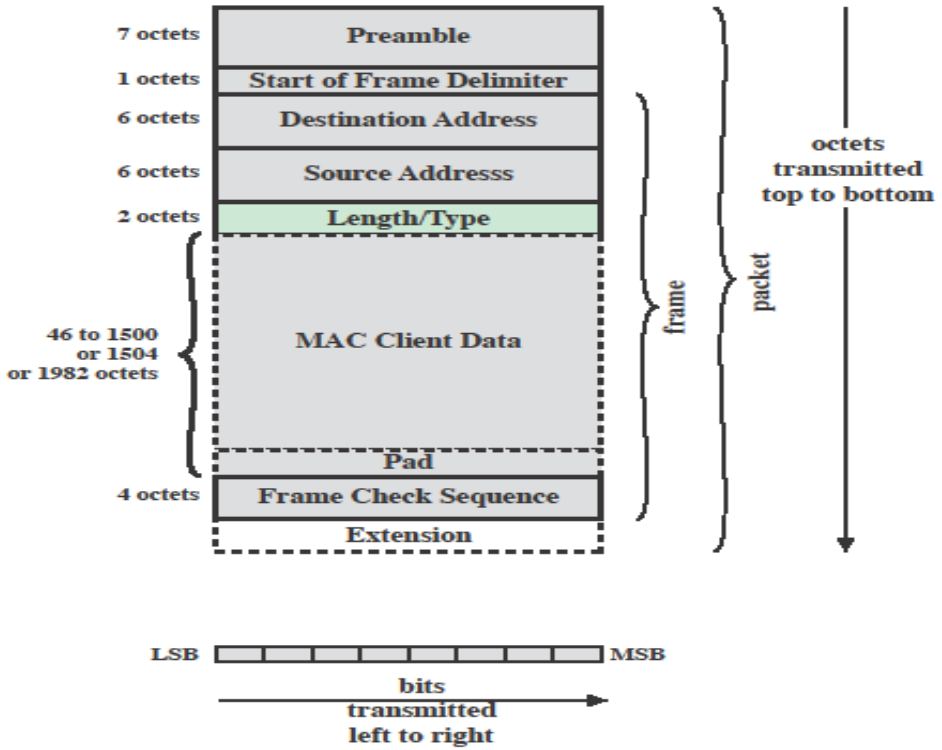

Figure 16.3 IEEE 802.3 MAC Frame Format

IEEE 802.3u (Fast Ethernet)

- Operates at 100 Mbps. It is compatible with the regular Ethernet which operates at 10 Mbps and uses CSMA/CD MAC protocol. This is done by reducing the maximum distance between two stations to 250 m as opposed to 2500 m in the Ethernet. The minimum frame size is still 64 bytes as in the Ethernet.

- **Physical Layer options of Fast Ethernet:**

- **100BASE-T** is any of several Fast Ethernet standards for twisted pair cables, including:
- **100BASE-TX** (100 Mbps over two-pair Cat5 or better cable) is the predominant form of Fast Ethernet. The segment length for a 100BASE-T cable is limited to 100 m, the same as in Ethernet and Gigabit Ethernet. A typical cat 5 cable contains 4 pairs and can therefore support two 100BASE-TX links.
- **100BASE-T4** (100 Mbps over four-pair Cat3 or better cable, now defunct). Data stream was split into 3 separate data streams, each with an effective data rate = 331/3 Mbps.

- **100BASE-FX** is a version of Fast Ethernet over optical fiber. It uses a 1300 nm light wavelength transmitted via two strands of optical fiber, one for receive (RX) and the other for transmit (TX). Maximum length is 2 kms for full-duplex over multimode optical fiber.

IEEE 802.3z (Gigabit) Ethernet Configuration

Workgroup represents high-performance workstations and servers.

The 100/1000 Mbps hub supports workstations and servers and 100 Mbps hubs via the 100 Mbps links and supports 1-Gbps links for backbone connectivity.

Gigabit Ethernet retains the CSMA/CD MAC protocol and frame format as its 10 Mbps and 100 Mbps predecessors.

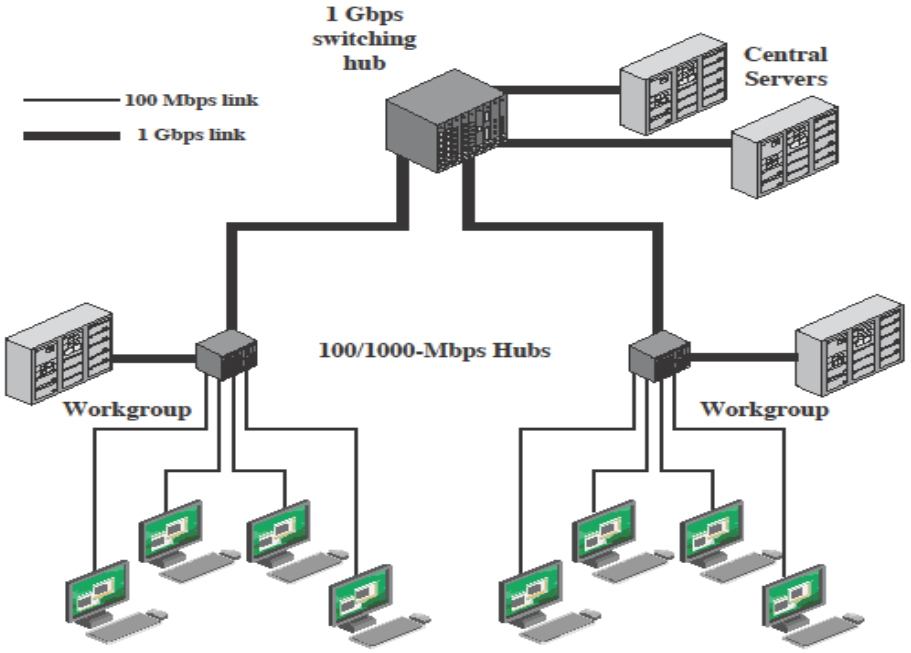

Figure 16.4 Example Gigabit Ethernet Configuration

Gigabit Ethernet Physical Layer Options

1000BASE-LX uses long wavelength laser (1,270–1,355 nm) and is specified to work over a distance of up to 5 km over 10 µm single-mode fiber. This is preferred for campus backbones and similar applications. It can also run over all common types of multi-mode fiber with a maximum segment length of 550 m.

1000BASE-SX is a fiber optic gigabit Ethernet standard for operation over multi-mode fiber using a 770 to 860 nanometer light wavelength. The standard specifies a distance capability between 220 m (62.5 µm fiber) and 550 m (50 µm fiber). This standard is highly popular for intra-building links in large office buildings.

1000BASE-T (also known as IEEE 802.3ab) is a standard for gigabit Ethernet over copper wiring. Each 1000BASE-T network segment can be a maximum length of 100 m, and must use Cat 5 cable at a minimum. Cat 5e cable or Cat 6 cable may also be used.

1000BASE-CX is a standard for gigabit Ethernet connections with maximum distances of 25 m using shielded twisted pair (STP). It is used in equipment rack room.

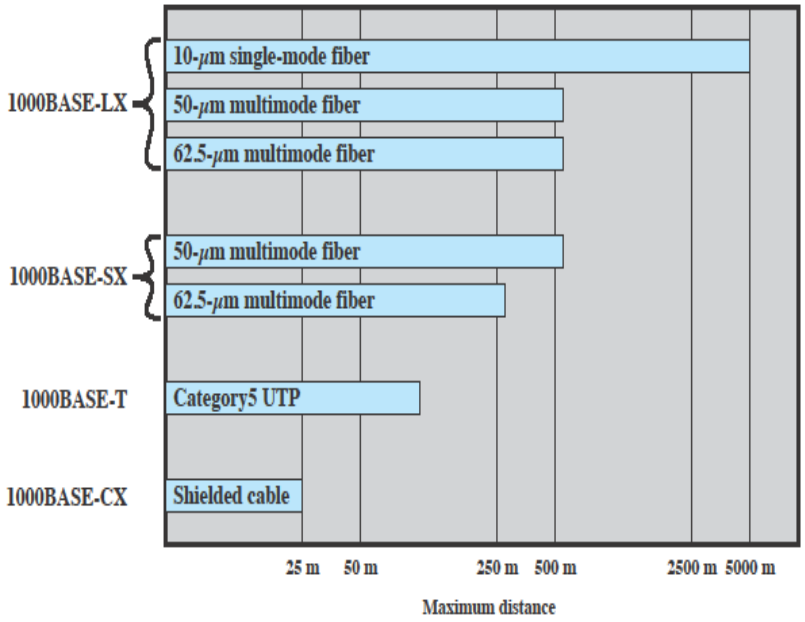

Figure 16.5 Gigabit Ethernet Medium Options (log scale)

Gigabit Ethernet

- Gigabit Ethernet supports two modes of operation: Full-duplex and half-duplex. In full-duplex mode, a switch is used and there is no CSMA/CD. In half-duplex mode, a hub is used, CSMA/CD is needed and collisions can occur.

- In Ethernet, the minimum frame size is 64 bytes at 2500 m and 64 bytes at 250 m for Fast Ethernet. But in Gigabit Ethernet to retain 64 bytes minimum frame size would require a 25 m maximum LAN length. This is too small. So IEEE 802.3z committee added two new features:

- 1. Carrier Extension: The hardware (NIC card) adds its own padding to extend the frame to 512 bytes. Since this padding is added by the sending hardware and removed by hardware, the software is unaware of it and thus no changes are required to existing MAC software. But efficiency is low: 46 bytes / 512 bytes = 9% line efficiency (46 bytes is the payload of the minimum frame size of 64 bytes).

  2. Frame bursting: Allows a sender to transmit a sequence of multiple frames in a single transmission. If total burst < 512 bytes, then the hardware pads it again. If enough frames are waiting for transmission, this scheme is highly efficient and preferred over scheme 1.

Gigabit Ethernet Discussion

- In practice, no one is going to spend on Gigabit Ethernet cards to get high performance and then connect the computers with a hub to simulate classic CSMA/CD with its collisions.

- While hubs are cheaper than switches, Gigabit Ethernet cards are still expensive. To then economize by buying a cheap hub and slash performance of the new system is foolish. So both Carrier Extension and Frame Bursting are moot since no one will use Gigabit Ethernet hubs. Only switches will be used. But IEEE 803.z committee put them in for backward compatibility.

IEEE 802.3ae (10 Gbps) Ethernet Configuration

10 Gbps Ethernet backbone pipes will help relieve congestion for workgroup switches, where Gigabit Ethernet uplinks can easily become overloaded, and for server farms, where 1-Gbps NICs are already in widespread use.

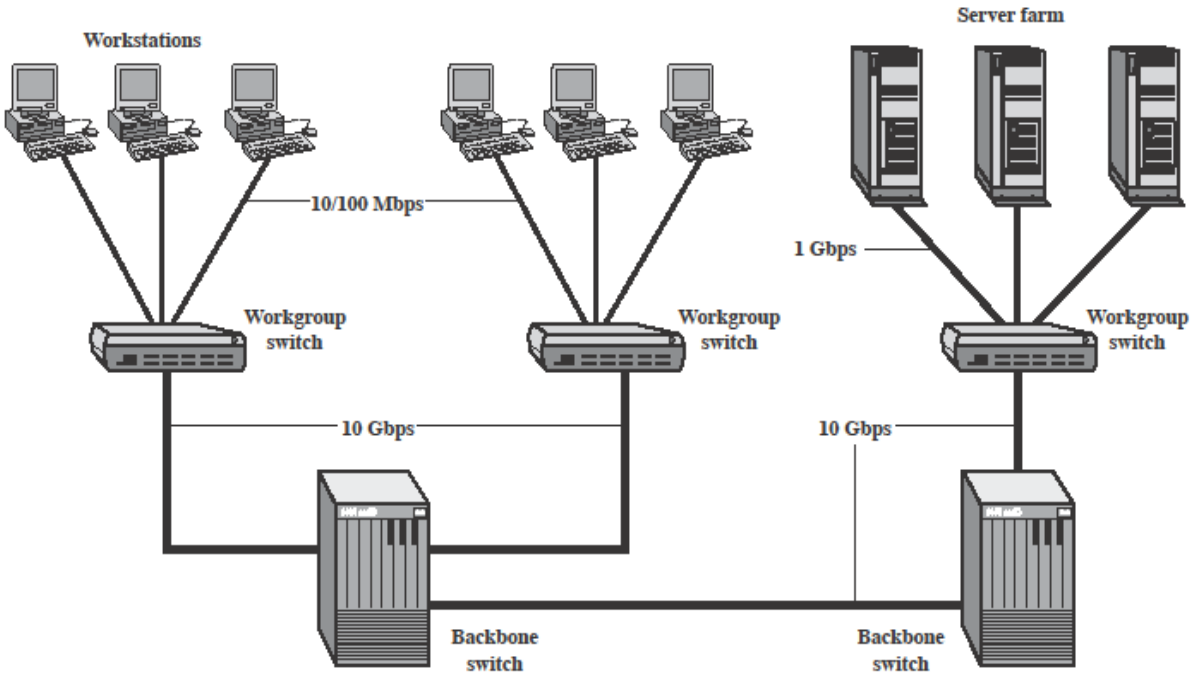

Figure 16.6 Example 10 Gigabit Ethernet Configuration

10 Gbps Ethernet

- **Principle driving requirement for 10 Gbps Ethernet is the increase in Internet and intranet traffic.**

  Factors contributing to this growth are:

- An increase in the number of network connections
- An increase in the connection speed of each end-station (e.g. 10 Mbps users moving to 100 Mbps)
- An increase in the bandwidth-intensive apps such as high-quality video.
- An increase in Web hosting and app hosting traffic

- The **10 gigabit Ethernet** standard was first published in 2002 as IEEE Std 802.3ae. It defines a version of Ethernet with a nominal data rate of 10 Gbps, ten times as fast as gigabit Ethernet.

- 10 gigabit Ethernet supports only full duplex links which can be connected by switches. Half duplex operation and CSMA/CD (carrier sense multiple access with collision detection) are not supported in 10GbE.

- **Initially network managers will use 1- Gbps Ethernet to provide high-speed, LAN backbone connectivity between large capacity switches. As the demand for bandwidth increases, 10 Gbps Ethernet will be deployed throughout the network to include server farms, backbone, and campus wide connectivity.**

10 Gbps Ethernet …

- The technology also allows the construction of MANs that connect geographically dispersed LANs between campuses. Thus, Ethernet begins to compete ATM and other WAN technologies. Where the customer requirement is data and TCP/IP transport, 10 Gbps Ethernet provides substantial value over ATM transport for both network end users and service providers:

- No expensive, bandwidth consuming conversion between Ethernet packets and ATM cells is required; the network is Ethernet end-to-end.
- The combination of IP(v6) and Ethernet offers QoS and traffic policing capabilities that approach those provided by ATM.
- A wide variety of standard optical interfaces (wavelengths and link distances) have been specified for 10-Gbps Ethernet, optimizing its operation and cost for LAN, MAN, and WAN applications.

  The goal for maximum link distances covers a range of applications: from 300 m to 40 km.

10 Gbps Ethernet Physical Layer Options

- **10 GBASE-S (short):** For 850 nm transmission on multi-mode fiber; distance up to 300 m. (10GBASE-SR and 10GBASE-SW)

- **10 GBASE-L (long):** For 1310 nm transmission on a single-mode fiber; distance up to 10 km. (10GBASE-LR and 10GBASE-LW)

- **10 GBASE-E (extended):** For 1550 nm transmission on a single-mode fiber; distance up to 40 km. (10GBASE-ER and 10GBASE-EW)

- **10 GBASE-LX4:** For 1310 nm transmission on a single-mode or multi-mode Fiber; distance up to 10 km. This uses WDM to Multiplex the bit stream across 4 light waves.

Note: The first three of these have two sub-options: "R" and "W". The R implementations are designed for use over *dark fiber,* i.e. a fiber that is not in use and not connected to any other equipment. The W implementations are used to connect to SONET equipment (WANs).

10GBASE-SW, 10GBASE-LW, 10GBASE-EW are varieties that use the WAN (W) option are designed to interoperate with OC-192/SONETequipment using A SONET frame running at 9.953 Gbps.

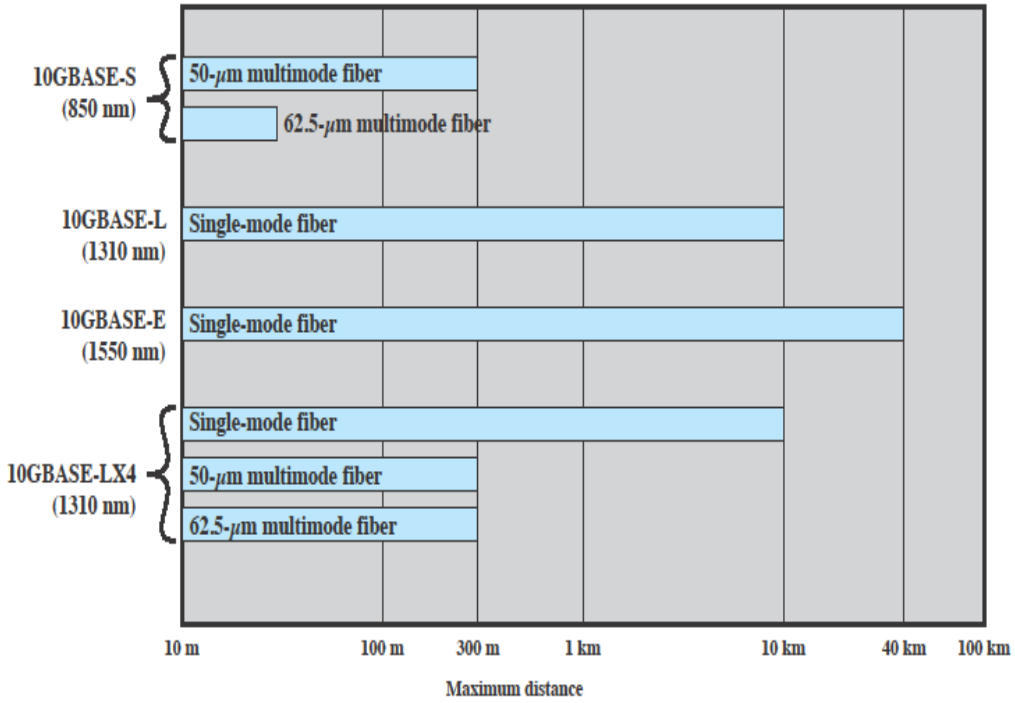

Figure 16.7 10-Gbps Ethernet Distance Options (log scale)

IEEE P802.3ba (100 Gbps) Ethernet

- **40 Gigabit Ethernet**, or **40GbE**, and **100 Gigabit Ethernet**, or **100GbE**, are Ethernet standards developed by IEEE P802.3ba Ethernet Task Force which was ratified in June 2010. These standards support sending Ethernet frames at 40 and 100 gigabits per second over multiple 10 Gbps or 25 Gbps lanes. Two distinct speeds were chosen to serve different applications (network aggregation and end-station requirements).

- Provides a significant increase in bandwidth while maintaining maximum compatibility with the installed base of 802.3 interfaces.

- Use either WDM in 1310 nm wavelength region with four 25 Gigabit or four 10 Gigabit channels, or parallel optics with four or ten optical fibers per direction.

IEEE P802.3ba (100 Gbps) Ethernet

- **Market drivers for 100 Gbps-Ethernet:**

- **Data Center/Internet media providers:**

  To support the growth of Internet multimedia content and web apps, content providers have been expanding data centers, pushing 10-Gbps to its limits. Likely to be high-volume early adopters of 100 Gbps Ethernet.

- **Metro-video/service providers:**

  Video on demand has been driving a new generation of 10-Gbps Ethernet MANs. Likely to be high-volume adopters in the medium term.

- **Enterprise LANs:**

  Continuing growth in convergence of voice/video/data is driving up network switch demands. Adoption of 100-Gbps is likely to be slow with main reliance on 1-Gbps/10-Gbps mix.

- **ISP core routing:**

  With the massive amount of traffic flowing through these nodes, these installations are likely to be early adopters of 100 Gbps Ethernet.

100 Gbps Ethernet Configuration

Trend at large data centers, with substantial bank of blade servers, is the deployment of 10-Gbps ports on individual servers to handle massive amounts of multimedia traffic provided by these servers. This stresses on-site switches needed to connect large numbers of servers. A 100 Gbps rate provides the bandwidth to handle the increased traffic load.

100 Gbps will be deployed in switch uplinks inside the data center as well as providing interbuilding, intercampus MAN and WAN connections for enterprise networks.

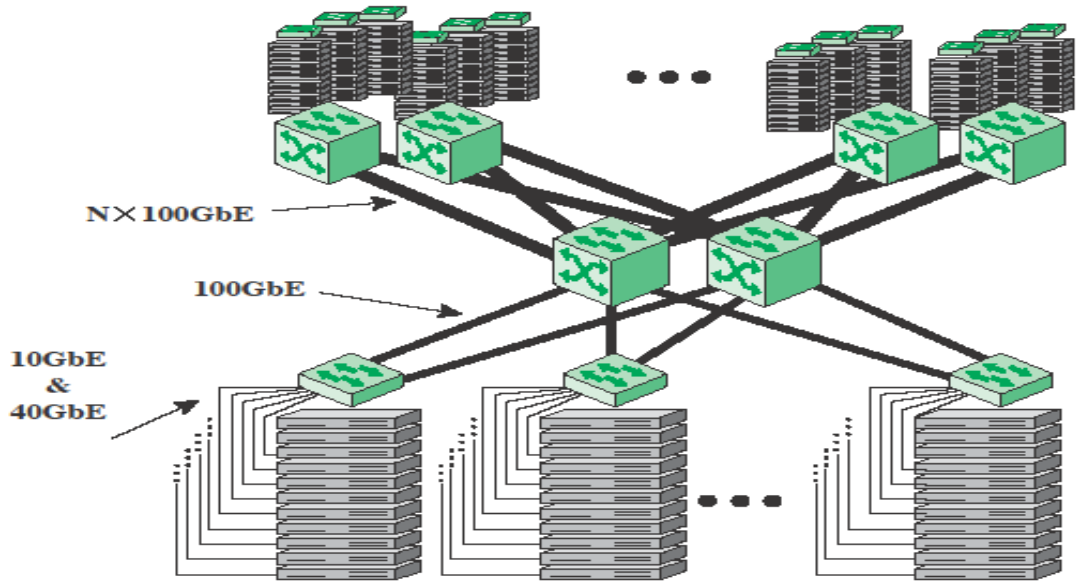

Figure 16.8 Example 100-Gbps Ethernet Configuration for Massive Blade Server Site

Media Options for 40 Gbps and 100 Gbps Ethernet

(Table 16.4)

PHY	40 Gigabit Ethernet	100 Gigabit Ethernet
1 m backplane	40GBASE-KR4	
10 m over copper cable	40GBASE-CR4	100GBASE-CR10
100 m over Multimode Fiber	40GBASE-SR4	100GBASE-SR10
10 km over Single Mode Fiber	40GBASE-LR4	100GBASE-LR4
40 km over Single Mode Fiber		100GBASE-ER4

Multilane Distribution for 100 Gbps Ethernet

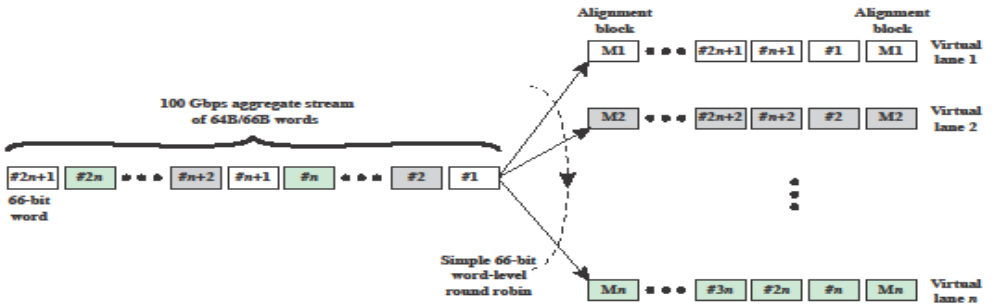

(a) Virtual lane concept

1	0	Frm1	Frm2	reserved	reserved	reserved	reserved	~VL#	VL#

(b) Alignment block

Figure 16.9 Multilane Distribution for 100-Gbps Ethernet