Presentation 2

Access Networking

Access networks connect the user with the global network (Internet). For most users, "the network" refers to the access network and the service they receive from a remote web server or media streaming service. In this chapter we will review some of the technologies that are relevant — but invisible — to users at various layers.

Slide 2 **Small Office / Home Office (SOHO)**

The small office or home office user has a number of devices that connect to each other by a local area network (LAN) that includes a router for connection to the Internet. Slide 2 shows a hub that combines the functions of an Ethernet switch, a WiFi access point, an IP router, and an ADSL modem for connection to Bezeq.

Ethernet (IEEE standard 802.3) and WiFi (IEEE standard 802.11) are LAN protocols defined at the physical and data link layers. ADSL is a physical layer protocol used for fast transmission of bits from point to point, and is used to connect homes and small offices to Internet service providers.

Slides 3 – 12 **Laptop Browser to Web Server**

The simplified view of the client / server connection is shown in slide 3. The user opens a web browser on a laptop to reach a web server. The browser runs an **HTTP client** at the application layer that requests transport service from the OS TCP agent, that requests network service from the OS IP agent, that requests service from the OS WiFi agent, that transmits a request to the router over WiFi. The router (an edge router for the home AS) forwards the IP datagram over the ADSL line to the **Internet Service Provider (ISP)**, which forwards the datagram over the Internet until it reaches the server. At the server, the local router transfers the IP datagram over some protocol at the physical and data link layers to the server. The server data link agent passes the IP datagram to the IP agent, which passes the encapsulated TCP segment to the TCP agent, which passes the HTTP request to the HTTP server agent.

Slide 4 shows some features of the ADSL network. It generally multiplexes high-speed digital data with the usual analog voice channel of a telephone line (at the right side). At the telephone office the voice and data are separated — the Plain Old telephone Service (POTS) is sent to a telephone switch and the data is sent to a DSL Access Multiplexer (DSLAM). The high-speed data is routed to the Internet over the Telco's (telephone company) internal ATM network. Additional details are shown in slide 5.

The technical term for the ADSL modem is ADSL transceiver unit-remote (ATU-R). Slide 6 shows details of an older Bezeq ATU-R, along with translations of some marketing terms to the protocols that implement them.

Slide 7 shows a simplified view of the ADSL implementation by Bezeq. The client is connected by ADSL and ATM to the Bezeq telephone network. Bezeq forwards the datagrams from the client packets to the ISP over an IP network, and the ISP forwards the datagrams to the destination over the Internet backbone.

The make this system work properly a number of intermediate protocols are used. Slides 8 to 11 show the case of a campus in which each user has a separate account with the ISP.

Slide 8 shows the client connection to the router. Between the usual IP layer and Ethernet layers are two additional protocols, PPP for login management and PPPoE to allow each client to establish a private session with the ISP.

Slide 9 shows that the router places the PPPoE frame into an 802.3 Ethernet frame which is then placed inside 3 ATM protocol layers (MPOA, AAL5, and ATM). These ATM packets are sent to Bezeq over the ADSL line at the physical layer.

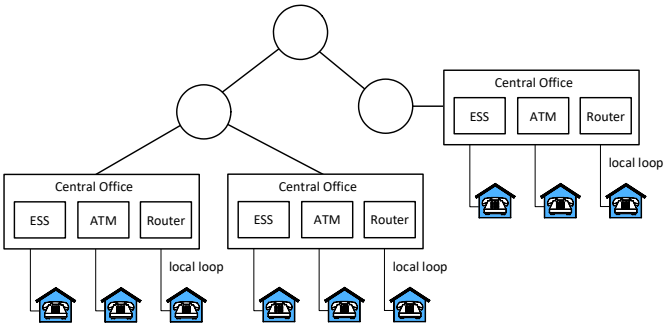
In slide 10, Bezeq forwards the 802.3 Ethernet frame (containing PPPoE and PPP frames) to the ISP over some physical channel. Although these Ethernet frames were only defined for use in a LAN, the use of PPPoE and PPP allows them to be used in the non-local network between Bezeq and the ISP. The PPPoE frames carry PPP packets from the user at home to the ISP over a virtual point-to-point channel — the user conversation with the ISP appears to be direct without intervention by Bezeq. The PPP packets contain NCP packets that allow the user to login to the ISP service with user name and password.

Slide 11 shows that after login, the ISP extracts the user's IP datagrams and forwards them over the Internet in the usual way (between the ISP and server, PPP is just a normal data link protocol).

Slide 12 shows the usual arrangement in which the user's computer runs the standard protocol stack (TCP / IP / WiFi) and the Bezeq router handles the intermediate protocols required to perform transparent IP routing of user datagrams, while maintaining the user connection with the ISP using PPP over PPPoE over 802.3 Ethernet.

Slides 13 – 18 Telephone Network

The largest, oldest and most available wide area network is the telephone network.

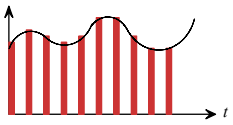


The traditional telephone network is a hierarchy of central offices and higher level switching centers. The **local loop** (so-called "last mile") is the standard analog phone line connecting the Telco central office to nearly every home and office. These lines were designed to carry analog voice and low speed data. The analog voice channel is converted to a 64 kbps bit stream (**DS0**) in the central office.

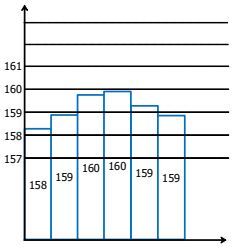
The **PDH / SDH** digital hierarchy combines DS0 digital voice streams for lines operating at data rates of 1.544 Mbps (T1) to 40 Gbps (STM-256). DS0 voice streams are switched between user phones in the **ESS7** switching system, a hierarchical tree of central office and higher-level switches. The more complex (and more expensive) ATM system is a general mesh of packet switches operating up to 2.5 Gbps for combined voice and data streams.

The conversion of **analog voice to DS0 data** streams is performed in three steps:

1. The analog voice is restricted to a bandwidth from 300 Hz to 3300 Hz and sampled every 125 μsec (producing 8000 voice samples per second).



2. The sample stream is rounded off to whole integer values in the range 0 to 255. These values can be represented by an 8-bit number.



3. The bits are transmitted sequentially.
   (8000 samples / second) × (8 bits / sample) = 64 kbps

The DS0 data streams are combined for transmission between central offices in the PDH / SDH **multiplexing hierarchy** shown on slide 15. This means that a single optical fiber running at 155 Mbps can carry 2048 phone calls from Tel Aviv to Haifa at one time.

The **cellular voice network** connects wireless phones to the telephone network using wireless transmission. Geographical areas are divided into **clusters** and clusters are divided into **cells**. In any given cell, a phone connects to a base station by digital radio transmission. The digitized voice (digitized in the phone) is received by the base transceiver, forwarded to a cell controller, then to a cluster controller, and from the cluster to a **mobile switching center (MSC)**. The voice streams are generally sent to the regular Telco network for forwarding to remote locations.

In order to manage **mobility**, the MSC maintains a database (HLR / VLR) that identifies the cluster and cell in which each mobile device is operating. This is required to permit calls to be routed to any given mobile user.

Data traffic is separated from voice traffic at the cluster controller and sent to a Serving GPRS Support Node (**SGSN**), which operates as an MSC for data. The data is forwarded over the General Packet Radio Service (**GPRS**) network, based on IP. The data is sent to a Gateway GPRS Support Node (**GGSN**) and from there to the Internet. The GPRS network manages user mobility for incoming Internet traffic. Slide 17 shows the relationship of controllers in **3G** terminology.

From an economic point of view, the existing telephone network represents an enormous investment in existing equipment, including a global network of hardware nodes and transmission lines developed to provide many services besides Internet. Most services we developed long before the Internet:

- Telegraph — 1794
- Telephone — 1876
- Teletype modem — 1943
- Digital telephone — 1962
- Internet opened to public — 1992

As a practical matter, existing infrastructure is adapted to new uses, and newer models replace older models when older equipment is retired.

Slide 19 Changes in Telephone Business: 2000 — 2008

By the year 2000 it was clear that while business revolves around the data center (data access, storage, processing, and service), employees still conduct business by talking. Mobility had become a standard part of workflow and while video call volume was growing, voice calls still cost money.

Network infrastructure providers (Telcos) saw that their existing infrastructure was designed for voice calls, but their main source of profits was no longer in voice calls, but rather in leasing their telecommunication lines to Internet service providers, private corporate and institutional WANs, and mobile backhaul (the interconnection of cell sites and other cellular infrastructure).

Telco business strategy developed around a plan to scrap the existing PSTN (the hierarchical tree of central offices), sell off the expensive real estate devoted central offices, and use the proceeds to build new a data-oriented mesh network that would support voice calls as just another form of media streaming. This program proceeded smoothly until the worldwide economic crisis of 2008, which began with an international crash of real estate prices.

Slide 20 Next Generation Networks (NGN)

The Telco strategy described in slide 19 was formalized as the **Next Generation Networks (NGN)**, an initiative of the International Telecommunications Union (ITU) for long-term network planning. The **ITU** is an organization within the United Nations that coordinates standards for telephone systems. NGN standardizes the Telco view of technology convergence as **ITU-T Recommendation Y.2001 (12/2004)**:

The NGN is conceived as an all-IP network, in which every service is defined to operate over IP, and all physical infrastructure is defined so to provide efficient support for IP. The NGN is seen as a universal gateway point that dissolves traditional service boundaries (voice versus data, mobile versus fixed, and so on). NGN guarantees universal mobility and promises every service to every user. In order to support different services over IP, extensive new QoS control mechanisms are required, especially IPv6, MPLS, and SIP (service initiation protocol).

Slides 21 – 22 Virtualization in Networking

As seen in slide 17, communication networks were traditionally built from dedicated hardware components, each designed to perform a specific task. For example an IP router was designed and optimized to perform IP forwarding as fast as possible. A virtual system performs the same task as a dedicated system but is implemented in software running on a generic programmable device. Virtual systems must be transparent — it must appear to be a dedicated device to any system to which it provides service and any system from which it requests service.

One example is a virtual LAN (vLAN) in which a programmable LAN switch appears as one or more LAN switches interconnected by a router. The vLAN can segregate LAN segments so that local traffic on one segment does not interfere with local traffic on another, unless the IP addresses require forwarding.

Another example is a Virtual Private Network (VPN) in which a private network is implemented over a public network by isolating private traffic from other users by IP address.

Slide 23 Traditional Networking Functions

Networks can be described as involving two basic functions — forwarding and control.

The forwarding function performs the transfer of data, implementing the network protocols at layers 1, 2, and 3, of the OSI model. In order to achieve highest performance, forwarding is usually performed by dedicated hardware designed specifically for the purpose — Ethernet switches and so on.

The control function involves management of the forwarding function, that is, configuration of the network topology and operating policies, supervision of the running network, measurement, and maintenance. Control is traditionally implemented in the dedicated network hardware by setting mechanical switches, rearranging cable connections, and programming the hardware at a console interface.

Slides 24 — 25 Software Defined Networking (SDN)

A Software Defined Networking (SDN) is a virtual network built from generic programmable hardware. This requires fast implementation of networking protocols and very high I/O capacity. The hardware includes a standardized control interface for configuration by a system programmer, to implement the traditional networking functions — modes, connections, protocols, and policies.

The overall hardware that can be configured to implement an SDN is called the fabric. A networking function implemented in the SDN is known as a logical device.

The data plane (user plane) refers to the aspects of the SDN directly involved in forwarding user data. The control plane handles network configuration data, along with security, maintenance, and other management responsibilities.

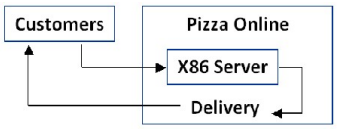
In a Software Define Network, an SDN Controller provides applications with an abstracted view of the logical network assembled as the SDN. Based on this information, SDN application programs communicate their network requirements to the Controller. The Controller translates these requirements to SDN Datapaths — the assembly of physical hardware that implements data forwarding. The Datapath may consist of a single forwarding device, such as a switch or a router, or may be a logical device defined from an internal network of forwarding devices. The hardware assembled to form an SDN is called a fabric.

Slides 26 – 32 Cloud Computing

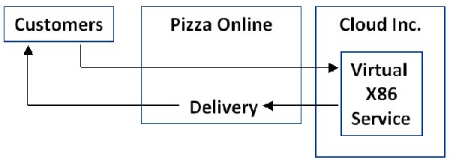
The rapid expansion of **cloud computing** was not anticipated when the NGN framework was being written, but it promotes similar goals. Although cloud computing requires various specialized support technologies, it is more a **business model** than a new form of computing.

The cloud is an outsourcing service model in which a service provider offers computing service to the user as specified in a Service Level Agreement (SLA) that guarantees the service details. In particular the service provider handles the complex tasks associated with **operations, administration, maintenance (OAM)**.

For example, consider a pizza parlor that offers its customers online ordering. The restaurant operator must arrange for a web server, which may require purchase of hardware and hiring of a programmer to maintain the system.



As an alternative, the restaurant operator can arrange the web server to be operated by a cloud service provider who assumes responsibility for the web service in return for monthly payment.



In this sense, the restaurant operator "purchases" computing as a commodity resource just like of water and electricity. The main business advantage is economies of scale, meaning that a large service provider can provide service at a lower cost. Cloud computing also cuts labor and capital costs from the user's balance sheet, which makes investors happy (regardless of the business case).

Cloud service is organized from conventional resources — mainframe computers or server farms with large high I/O capacity running a mix of standard operating systems in native or virtualized environments. The service provider offers a menu of services, a mix of hardware, software, and network capacities. The unique technological issues include:

- Service reliability — the provider is committed to the SLA as a minimum level of service
- The provider must be able to minimize the configuration cost required for the SLA
- The user must be able to identify the SLA requirements to minimize contract costs

Cloud computing is often divided into a hierarchy of service types:

- In **Infrastructure as a service (IaaS)**, the user sees a virtual hardware environment, which may be implemented in real hardware or a virtual machine. The user installs an OS and software to run jobs.
- In **Platform as a service (PaaS)**, the user sees a virtual OS environment, which be implemented as a single OS on a single hardware platform or as a virtual OS. The user installs software to runs jobs.
- In **Software as a service (SaaS)**, the user sees a virtual application software environment consisting of applications running on a private OS or "sandboxed" (isolated) on a shared OS. A sandbox is a private execution environment per application instance.
- In popular **Storage as a service (STaaS)** systems (Dropbox, iCloud, OneDrive, and others), the User sees a virtual mounted storage device. This may be seen as a subset of SaaS.

Among the considerations in cloud computing are cost, reliability, and agility.

In choosing between a setting up a private server or using a cloud service, the user must decide whether operating costs (capital + OAM costs) associated with a private server will be more or less than the operating debt associated with commitment to a cloud contract. The advantages of guaranteed service and simplified growth must be evaluated as a business case.

The provider must find ways to achieve economies of scale to lower the cost per compute job. This can be a difficult task when considering the need for redundant infrastructure for continuity

disaster recovery, and the need to reconfigure service and infrastructure to provide growth, load balancing, and time-zone serving. These needs are simplified through centralized management of OAM, security, and performance. Aspects of centralization are the virtualization that allows multiple users to be served on one physical server, and the multitenancy that provides multiple sandboxed application instances on OS.

In public cloud computing services, the service provider acts as a public utility selling (renting) computing service. The initial service providers, Amazon, Microsoft, Google, and IBM, were leveraging large existing infrastructure, offering a menu of services at fixed prices.

Many organizations have organized their data centers in the form of a private cloud, managed internally or outsourced. The goal is to isolate service developers from implementation issues by establishing a standard development platform. A large corporation can establish a mainframe computer that provides a large number of virtual Windows OS environments for the development of software services offered to the public. In this way, the developers and the public see a standard and familiar software environment, and are isolated from OAM and growth issues in the mainframe infrastructure. This may permit the business to lower costs by maintaining a small group of experts in the infrastructure details and a much larger number of developers who will spend their careers generating new products for the same Windows environment without retraining, promotion or significant salary increases. In Israel, a number of workplace disputes have arisen in connection with this trend.

Cloud computing can be seen as another stage in a **ping-pong** of developments in the computing world over its history:

In the 1950s and 60s, computing was dominated by centralized mainframe computers offering multiple OS instances over a hypervisor. A timesharing OS was used to serve multiple users, so that the user sees a single OS environment via a dumb terminal (now called a thin client).

In the 1970s, user applications were offloaded to a collection of minicomputers offering timesharing services, so that the user would see a timeshared OS environment via dumb terminal.

In the 1980s, user applications were further offloaded to personal workstations (PC) on which the user sees a single-user OS environment running locally.

In the 1990s, it became common (possible) to network single user workstations, so that the user would see a single-user OS environment running locally, but allowing file sharing with other users.

In the early 2000s, centralized control of local the OS environment on PCs was transferred to IT departments, making individual users again dependent on system managers.

In the 2010s, the cloud along with the netbook, tablet, and smart phone have largely reproduced the computing environment of the 1950s, where the old dumb terminal has been replaced by a high-resolution GUI and some limited local computing power.

It remains to be seen how the pendulum swings between centralization and decentralization.

Slide 32 shows a brief history of the changing attitudes toward centralization /decentralization over the past 50 years.

Slide 33 Enterprise Networks

It is unsurprising that large organizations maintain large private networks. These are often called enterprise networks, less because they are managed by large organizations that because the word "enterprise" is intended to suggest grandiosity. In any event, an enterprise network is characterized as:

- Larger, more complex, more expensive, and more "awesome" than a SOHO network
- Combines both LAN and WAN (wide area network) technologies
- Operates at a higher level of network security
- Has a very high overall traffic volume

Slides 34 – 39 Traditional Network Infrastructure

Slide 34 shows an illustration from the 2013 **Cisco Smart Business Architecture (SBA)**, which indicates a number of commercial systems and technologies for organizing a corporate or institutional headquarters connected with branch offices. We will discuss these technologies one by one.

The architecture shows 4 **regions**:

- Headquarters (right side)
- A regional site or main branch (on the top left)
- A remote site or smaller branch or partner (middle left)
- A remote worker

These regions are connected by three types of WAN:

- The Public Switched Telephone Network (PSTN)
- The Internet
- An Multiprotocol Label Switching (MPLS) WAN forming a Virtual Private Network (VPN)

The **headquarters** is further divided into several areas:

- The data center (top left) accessed through main local switches and dial-up PSTN
- Core switches controlling access to the data center
- WAN aggregation (lower left) providing access to the core switches from the Internet and the MPLS VPN
- Internet edge (middle left) providing gateway routers for access to WAN aggregation through a firewall
- Internal users (lower right)

The **data center** contains the main corporate servers, indicated as Cisco Unified Computing System (UCS) rack-mounted servers. This UCS system is advertised by Cisco as "Managing computing, networking, management, virtualization, storage access into a single integrated architecture," providing "a unique architecture that enables end-to-end server visibility, management, and control in both bare metal and virtual environments." On the left is the Wide Area Application Services (WAAS) central manager. Also visible is the Application Control Engine (ACE). These management systems are intended to provide IT managers a view of the

application environment, which is physically implemented across a large number of machines, operating systems, and environments.

The main data switches in the data center are Nexus 5500 — the main switch — and Nexus 2000, which permits expansion of the 5500 switch.

In the WAN aggregation region the MPLS VPN connects directly to the switch because the VPN is inherently secure. The general Internet access is connected through a firewall system or into a Remote Access (RA) VPN, which connects individual users to the data center over the Internet.

Slide 35 shows the conventional model of layered networking for **layer 3 routing** and **layer 2 switching**. Layer 3 (network layer) routing permits universal access to any node by next-hop IP datagram forwarding. Layer 2 (data link) switching is implemented in N x N non-blocking switches and permits access to nodes on the same physical network. Examples of layer 2 switches are:

- Ethernet MAC physical port-to-port switching
- ESS7 64 kbps voice telephone switch
- ATM cell switching for switched and permanent virtual circuits (SVC and PVC)
- Frame Relay switching for permanent virtual circuits (PVC)

Layer 2 switching is faster than layer 3 routing because the routes are set up in advance and there is no need to examine the IP addresses in the datagrams. Large organizations may need to partition their internal networks into **subnets** for administrative and security reasons. In the conventional model, subnets are defined by the IP address mask and routers are required to isolate subnets.

In a virtual LAN (vLAN) the LAN switch is configured in software to partition nodes into subnets with separate broadcast domains. There is no need to insert layer 3 routers to divide the local network into subnets.

In a Virtual Private Network (VPN), encrypted protocols are used to implement a private network on public infrastructure. The VPN creates a "tunnel" between widely separated LAN segments, so that a user in one location can access a node at a remote location as if it were on the same LAN. This is similar to the use of PPPoE in slide 11 to permit a home user to communicate with the ISP over Ethernet as if the user were connected to a LAN at the ISP.

In **Multiprotocol Label Switching (MPLS)** a **label header** is added to each IP datagram. The label can be used to identify a preconfigured route and implement an end-to-end **virtual circuit**. This system requires MPLS-enabled routers, which choose the next hop by the MPLS label (similar to a VC label in a virtual circuit network) instead of examining the IP addresses of each datagram. This saves the time associated with layer 3 processing: accessing the datagram, IP routing, TTL update and so on.

In general, the boundary between switches and routers is becoming less rigid, permitting vLANs and VPNs to be defined quickly in software. An MPLS VPN of the type shown in slide 20 uses MPLS to configure routers to provide private routes — details of the customer's routers are hidden from other customers and the routers in the Internet backbone are not aware of the end-to-end VPN route. Generally, the data traffic is encrypted using IPsec to provide another level of security. **IPsec** is a layer 3 protocol that creates an "encrypted tunnel" in which data and identifying characteristics of the host endpoints are hidden.

Slide 37 shows examples of the core switches shown in slide 34. The **Nexus 7000** switches lines carrying 100 Gbps Ethernet with a total capacity of 15.76 Tbps (15,760 Gbps). The **Nexus 2000** is a fabric extender (FEX) — it can add physical ports to the Nexus 7000 from a remote location.

Slide 38 shows examples of commercial routers. The **XR 12000** can handle 1280 Gbps of incoming traffic, managing all standard IP versions and routing protocols. It can accept infrastructure traffic (physical and data link) from optical fiber (SONET/SDH), Ethernet, ATM, or multiplexed data on copper cable (DS-3/E3). Smaller routers for branch offices are shown below.

Slides 39 – 41 Cisco Software-Defined Access Design

Slide 39 shows a simplified system in the Cisco SDN model. An **underlay network** is the physical hardware used to implement the SDN. An **overlay network** is the logical network model defining the SDN. A database of user accounts (ISE) and hardware descriptions (DNA) provides the system with information about the end uses and capabilities of the network. **Edge nodes** connect user equipment to the SDN and **border nodes** provide the interface with the SDN as an autonomous system.

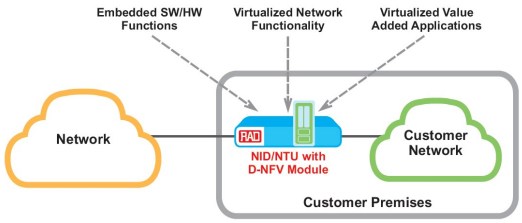
Slide 40 shows a Layer 2 overlay. Here, a physical configuration of (layer 2) switches is used to build an SDN that provides a virtual layer 2 system of a slightly different configuration.

Slide 41 shows a Layer 3 overlay. Based on the same physical configuration of (layer 2) switches on slide 37, this SDN provides a virtual network of switches connected through an IP router that forwards information based in IP addresses within the message section of layer 2 packets.

Slides 42 – 49 What Network Access Providers Do

The activities of network access providers can be seen by looking at the catalog of RAD, a large Israeli company that provides infrastructure for these providers. The key to these activities is management of virtual systems and services. Virtualization permits highly configurable network hardware supporting **Software Defined Network (SDN)** functions.

RAD calls this **Distributed Network Functions Virtualization (D-NFV)**.



The customer network is connected to the WAN (Internet or other) through a **Network Interface Device (NID)** for a general network or a **Network Termination Unit (NTU)** specific to **Carrier Ethernet**. The D-NFV provides various integrated networking and processing functions.

Carrier Ethernet shown in slide 44 is a set of methods for bridging Ethernet LAN segments over a WAN forming a VPN. The combined Ethernet segments, which may be separated over a wide geographical area can operate as a single Ethernet broadcast domain.

We recall that Ethernet (802.3) is a shared medium local area network defined up to about 100 meters with data rates of 10 Mbps to 100 Gbps on copper or optical fiber.

Carrier Ethernet types include:

- E-Line — a point-to-point connection between 2 LAN segments
- E-LAN — a general mesh among several LAN segments

- E-Tree — a hierarchical tree of LAN segments

Security is a major concern for network operators as shown in slide 43. Network access is controlled at the data link layer using the **Network Access Control protocol (IEEE 802.1X)** and **MANsec — MAC layer security standard (IEEE 802.1AE)**. At other layers **Supervisory Control and Data Acquisition (SCADA)** provides firewall protection (FW), intrusion protection (IPS), and anomaly protection.

Slide 46 shows a configuration for connecting a corporate headquarters with remote offices over standard Telco PDH leased lines. Some abbreviations are:

- ETH — Ethernet
- GbE / 10-GbE — Gigabit Ethernet (1 Gbps) / 10 Gigabit Ethernet (10 Gbps)
- DSL — digital subscriber line
- PDH — Plesiochronous Digital Hierarchy (standard telco multiplex hierarchy)
- L2/L3 — layer 2 (data link) and layer 3 (network)
- GPON — Gigabit Passive Optical Network
- ETX — Carrier Ethernet terminal equipment
- POP — point of presence

Slide 47 shows a Carrier Ethernet E-LAN mesh in which the headquarters, remote data center, and branch offices are organized into a single virtual LAN, using Carrier Ethernet over standard Telco PDH leased lines.

Slide 47 shows a configuration for connecting a corporate headquarters with remote offices over a public packet switched network such as IP or MPLS.

Slide 48 shows a configuration for **cellular backhaul** — connecting cell sites with the Mobile Switching Center.