L3HARRIS
FAST. FORWARD.

INTRODUCTION TO NETWORKING

Module 3: Applications and Security

Course Objectives

- Explain the use of network applications, virtualization, and network storage services.
- Summarize the concepts of cloud services
- Summarize the purposes of physical security devices.
- Explain authentication and access controls.
- Summarize common networking attacks.
- Explain common mitigation techniques and their purposes.

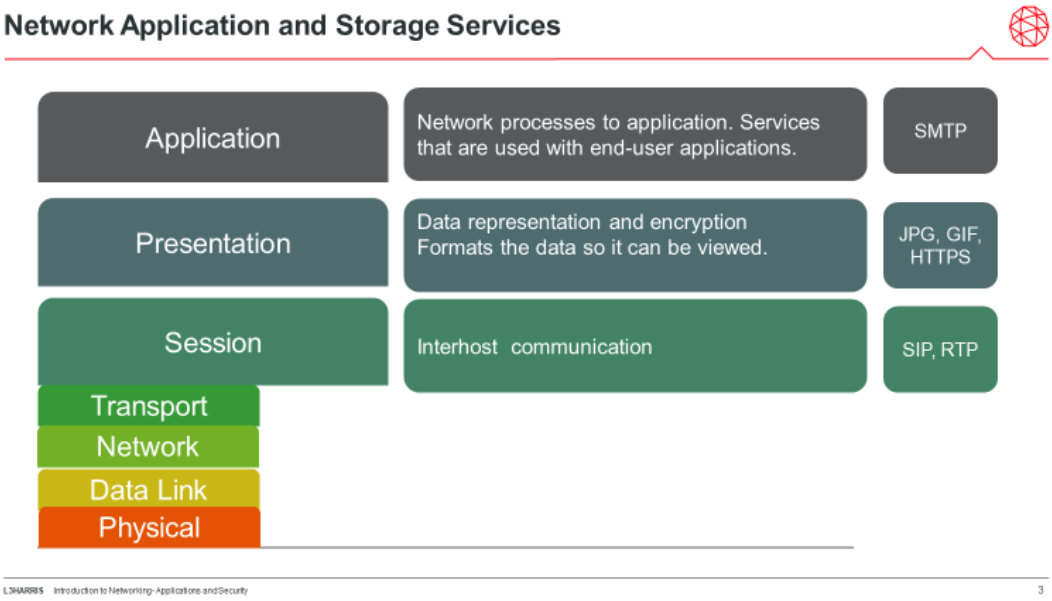

You have identified the Physical, Data Link, Network, and Transport layer technologies and protocols that enable basic connectivity between nodes. The “higher layer” of the model are mostly comprised of client-server protocols and applications. Client- server applications are based around a centralized server that stores information and waits for request from clients. (ie. Web browsing and email)

The TCP/IP protocol suite also includes application protocols that implement network services. The delivery of these services can be supported by technologies such as load balancing, virtualization, and storage networks. In this lesson, you will identify common network applications and service platforms.

Common Protocols

Protocol	Function	Common Port
HTTP	Websites and web applications	80
HTML	Standard mark up language for webpage creation	80
NTP	Enables clock synchronization	123
SMTP	Email delivery	25, 587, or 465
POP3S	Email retrieval	995
IMAP	Email retrieval	993

Hypertext Transfer Protocol (HTTP): A client connects to the server using an appropriate TCP Port and submits a request for a resource using a **Uniform Resource Locator** (**URL**). The server, which provides resources such as **Hypertext Mark Up Language (HTML**) files and other content or performs other functions on behalf of the client will acknowledge the request and responds with the data (or an error message).

Network Time Protocol (NTP) Enables the synchronization of many time dependent and time- critical applications such as authentication and security mechanisms, scheduling applications, and backup software.

Simple Mail Transfer Protocol (SMTP): Specifies how email is delivered from one system to another by making a connection from the sender's server to that of the recipient and then transfers the message.

- SMTPS- SSL/TLS Secured version of the protocol
- Can be configured on various ports (Port 25, 587, 465) dependent on business needs and accompanying protocols.

Post Office Protocol 3S (POP3S)

- Can be a different service running on the same SMTP server. Typically used for dial-up access.
- User is authenticated by username and password and contents of the mailbox are downloaded from processing on the local PC, typically deleting messages from the server. (Port 995)

Internet Message Access Protocol (IMAP)

- IMAPS for SSL/TLS secure
- Updated protocol to POP
- Default port for IMAPS is TCP port 993

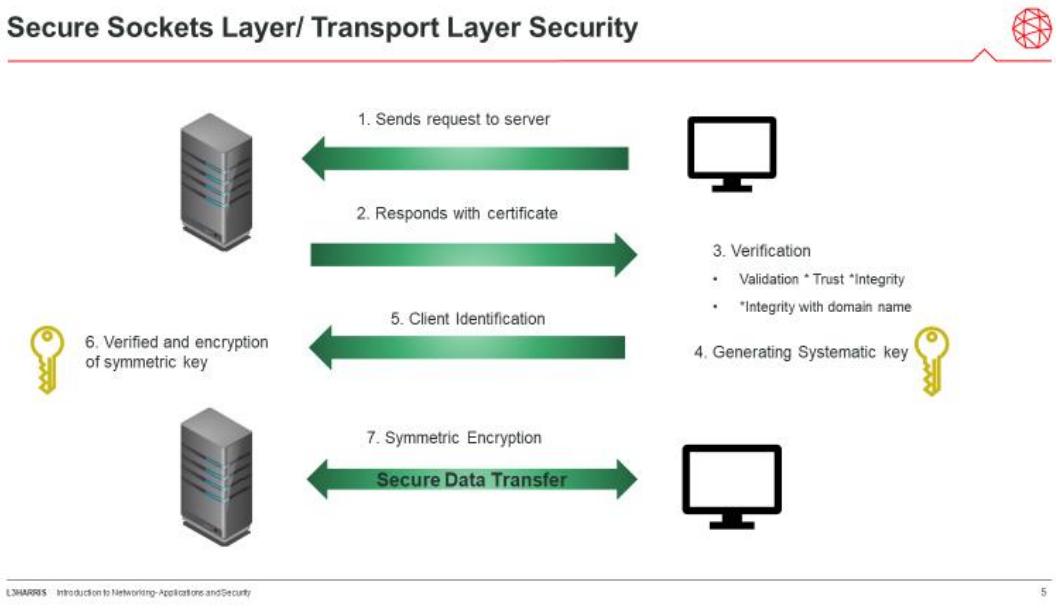

One of the critical problems for the provision of early e-commerce sites was the lack of security in HTTP. Under HTTP, all data is sent unencrypted, and there is no authentication of client or server. **Secure Sockets Layer (SSL)** was developed by Netscape in the 1990s to address these problems. **Transport Layer Security (TLS)** was developed from SSL and ratified as a standard by the IETF.

SSL/TLS works as a layer between the Application and Transport layers of the TCP/IP stack, or, in OSI terms, at the Session or Presentation layer. It's normally used to encrypt TCP connections. When it is used with the HTTP application, it is referred to as HTTPS or HTTP Over SSL or HTTP Secure, but it can also be used to secure other TCP application protocols, such as Telnet, FTP, NNTP, SMTP, and LDAP.

HTTP Secure (HTTPS) is a subset of HTTP that allows for a secure dialog between the client and server using SSL/TLS. To implement HTTPS, the web server is assigned a digital certificate by some trusted certificate authority (CA). The certificate proves the identity of the server, if the client also trusts the CA. The certificate is a wrapper for a public/private encryption key pair. The private key is kept a secret known only to the server; the public key is given to clients via the digital certificate. The server and client use the key pair in the digital certificate and a chosen cipher suite within the SSL/TLS protocol to set up an encrypted tunnel. Even though someone else might know the public key, they cannot decrypt the contents of the tunnel without obtaining the server's private key.

This means that the communications cannot be read or changed by a third party. Encrypted traffic between the client and server is sent over TCP port 443 (by default), rather than the open and unencrypted port 80. A web browser will open a secure session to a server providing this service by using a URL starting with https:// and it will also show a padlock icon in the address bar to indicate that the connection is secure. A website can be configured to require a secure session.

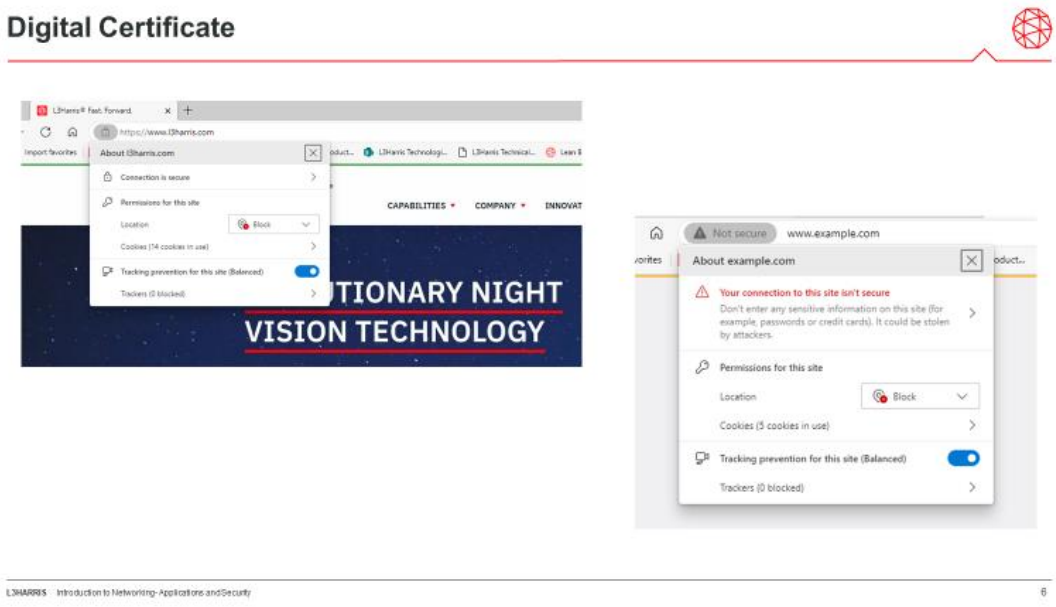

Encrypted traffic between the client and server is sent over TCP port 443 (by default), rather than the open and unencrypted port 80. A web browser will open a secure session to a server providing this service by using a URL starting with https:// and it will also show a padlock icon in the address bar to indicate that the connection is secure. A website can be configured to require a secure session.

Untrusted Certificate Issues If the certificate presented by a subject (server or user) is not trusted by the client application (such as a browser), the client will notify the user. A common reason for an untrusted certificate occurs when the certificate issuer is not trusted. Example: “More Coffees” web server receives a certificate signed by “ValidCert”. Unless VaidCert’s own certificate is stored in the browsers trusted root store, the client application will not trust “More Coffees ‘s” server.

- Users can typically choose to ignore the warning and add an exception. This should only be done if the cause of the lack of a trust relationship is understood.

Common causes of untrusted certificates:

- Different applications may have different stores of trusted certificates.
- Self-signed certificate holder means the holder is both the issuer and the subject of the certificate.
- The certificates subject name does not match the URL. This could be a configuration error, but it could indicate malicious activity. Confirm the certificate's common name and access the website by using that URL.
- The certificate is not being used for its stated purpose.
- The certificate is expired or revoked.
- Time is not correctly synchronized between the server and client.

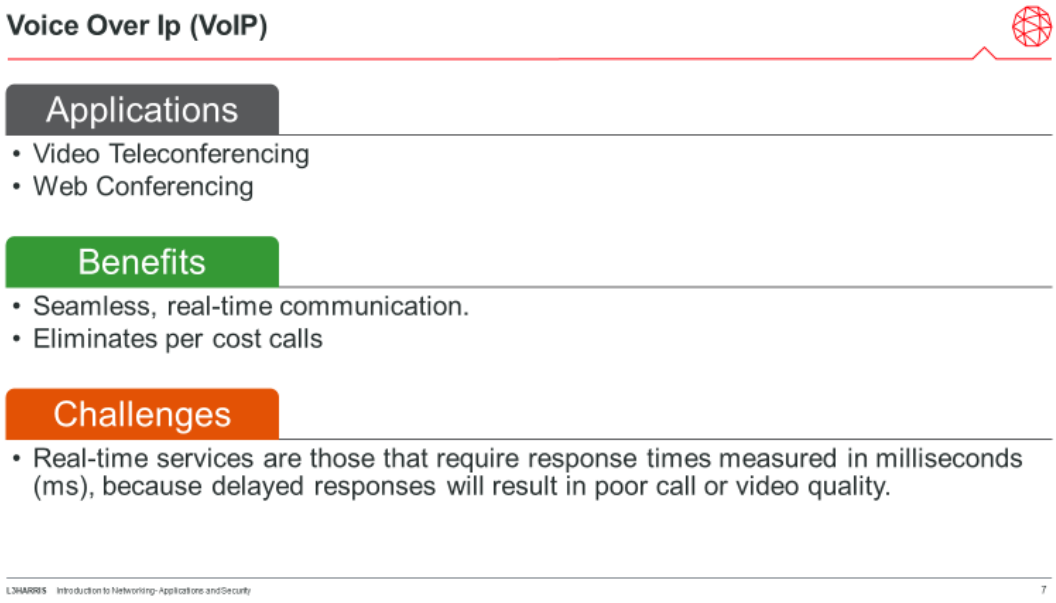

Voice over IP (VoIP) also called IP telephony, is a method and group of technologies for the delivery of voice communications and multimedia sessions over Internet Protocol networks rather than via the public switched telephone network (PSTN). The main challenges that these applications have in common is that they transfer real-time data and must create point-to-point links between hosts on different networks. Real-time services are those that require response times measured in milliseconds (ms), because delayed responses will result in poor call or video quality. This type of data can be one-way, as is the case with media streaming, or two-way, as is the case with VoIP and VTC

VoIP can provide both short-range and long-haul communications, so it can replace traditional telephone links by converting and then transmitting analog voice communications to digital signals sent over data cabling. As in a typical packet switched network, digital signals are broken down into packets, to transmit voice as data. After reassembling the packets, the digital signals are reconverted into audio signals. Because voice communications are time-sensitive, the system must ensure that packets arrive complete and in sequence.

- Voice software interfaces with an analog voice device, such as a microphone, to convert the analog voice into a data signal and to translate the dialing destination into a network address.
- When you make a telephone call, the network connection transmits signals over data networks, and transfers them to the standard phone system if the called party does not have a VoIP service.
- Conversely, when you dial a number that maps to a VoIP device, VoIP routes the call to the IP host. VoIP relies on the existing, robust infrastructure of IP networks and the near-universal implementation of IP.
- There are numerous ways of implementing VoIP. In the past, proprietary protocols specific to the software vendor were used, but most modern IP telephony solutions use Session Initiation Protocol (SIP) both for internal and public (trunk) telephone calls.

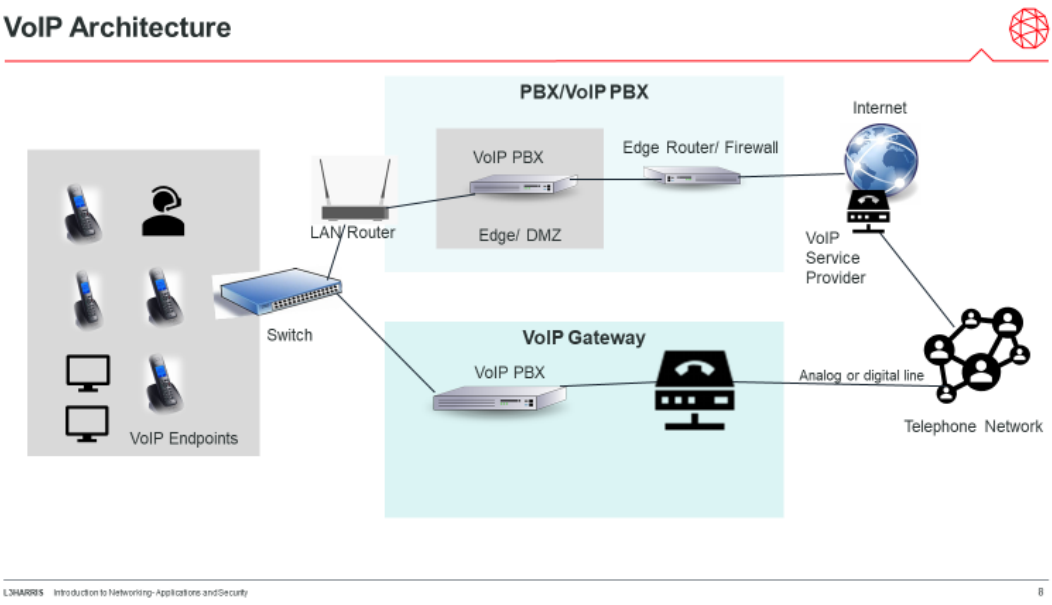

VoIP PBX: A VoIP PBX maintains a list of the internal accounts assigned to user endpoint devices. For internal calls and conferences, the PBX establishes the connection between local VoIP endpoints with data transmitted over the local Ethernet network. A VoIP PBX can also route incoming and outgoing calls with external networks. This might involve calls between internal and external VoIP endpoints, or with voice telephone network callers and receivers. A VoIP PBX will also support features such as music on hold and voice mail and can be implemented as software or hardware, depending on vendor and configuration.

VoIP Gateway: A means of translating between a VoIP system and voice-based equipment and networks, such as public switched telephone network (PSTN) lines.

- While many implementations depend on a service provider to facilitate connections between the local VoIP system and the voice telephone network, the gateway could be used if on-premise integration between data and voice networks and equipment are required.
- There are many options to deploy a VoIP gateway such as:

Hybrid or hardware-based VoIP PBX with a plug-in or integrated VoIP gateway

Separate gateway appliance

Foreign Exchange Subscriber (FXS) Gateway- to connect legacy analog handsets and fax machines to VoIP PBX

VoIP Endpoint: A VoIP/SIP endpoint can be implemented as software running on a computer or smartphone or as a dedicated hardware handset

VoIP Protocols

Real Time Services Protocols	
Session Control	Used to establish, manage, and disestablish communications sessions. They handle tasks such as user discovery (locating a user on the network), availability advertising (whether a user is prepared to receive calls), negotiating session parameters (such as use of audio/video), and session management and termination.
Data Transport	Handles the delivery of the actual video or voice information.
Quality of Service	Provides information about the connection to a QoS system, which in turn ensures that voice or video communications are free from problems, such as dropped packets, delay, or jitter.

Session Initiation Protocol (SIP) is one of the most widely used session control protocols. SIP endpoints are the end-user devices (also known as user agents), such as IP-enabled handsets or client and server web conference software. Each device, conference, or telephony user is assigned a unique SIP address known as a SIP Uniform Resource Indicator (URI).

- A VoIP/SIP endpoint can be implemented as software running on a computer or smartphone or as a dedicated hardware handset.
- Many VoIP/SIP handsets are connected to the corporate network over “normal’ data ports but assigned to separate virtual LANs (VLAN) than other data traffic.

Connection security for VoIP works in a similar manner to HTTPS. To initiate the call, the secure version of SIP (SIPS) uses digital certificates to authenticate the endpoints and establish an SSL/TLS tunnel. The secure connection established by SIPS can also be used to generate a master key to use with the secure versions of the transport and control protocols.

Real-Time Transport protocol (RTP and RTP Control Protocol (RTCP) RTP enables the delivery of a stream of media data via UDP, while implementing some of the reliability features usually associated with TCP communications.

- RTP does not guarantee reliability or real-time delivery. In fact, depending on the underlying network technology, this may be impossible to achieve.
- RTP works closely with the RTP Control Protocol (RTCP). Each RTP stream uses a corresponding RTCP session to monitor the quality of the connection and to provide reports to the endpoints.
- These reports can then be used by the applications to modify codec parameters or by the network stacks to tune Quality of Service (QoS) parameters.

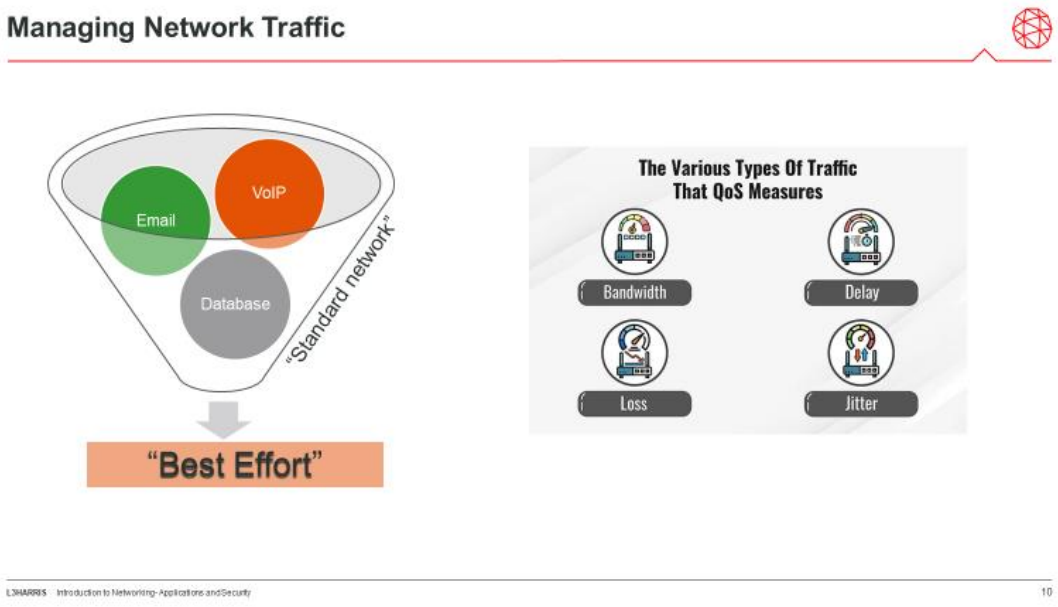

Characteristics of Network Traffic

- **Throughput** – it is the maximum amount of data per second. It is sometimes referred to as the speed or capacity of the link. This is commonly measured as bits per second (bps). On network devices, like routers and switches, the bandwidth is associated per interface. The interface bandwidth can either be Ethernet (10 Mbps), Fast Ethernet (100 Mbps), or Gigabit Ethernet (1000 Mbps).

Bandwidth requirements for voice calling can vary but allowing 100 Kbps per call upstream and downstream should be sufficient in most cases.

- **Latency/Delay** – is the time it takes for a transmission to reach the recipient, measured in milliseconds (ms). A high latency will cause a delay for traffic to arrive on the destination device and therefore causes a slower response time when establishing a connection to a specific device or application.

Latency: You can test the latency of a link using tools such as ping, pathping, and mtr. When assessing latency, you need to consider the Round-Trip Time (RTT). VoIP is generally expected to require an RTT of less than 300 ms. The link should also not exhibit more than 1% packet loss.

- **Jitter** –is defined as being a variation in the delay. Jitter manifests itself as an inconsistent rate of packet delivery. Jitter is also measured in milliseconds, using an algorithm to calculate the value from a sample of transit times. You can also use mtr to calculate jitter, Jitter should be 30 ms or less.

- **Loss** – it happens when the buffer of the router is full, and new incoming packets are being dropped. Having too much packet loss will cause the receiving device to receive an incomplete message.

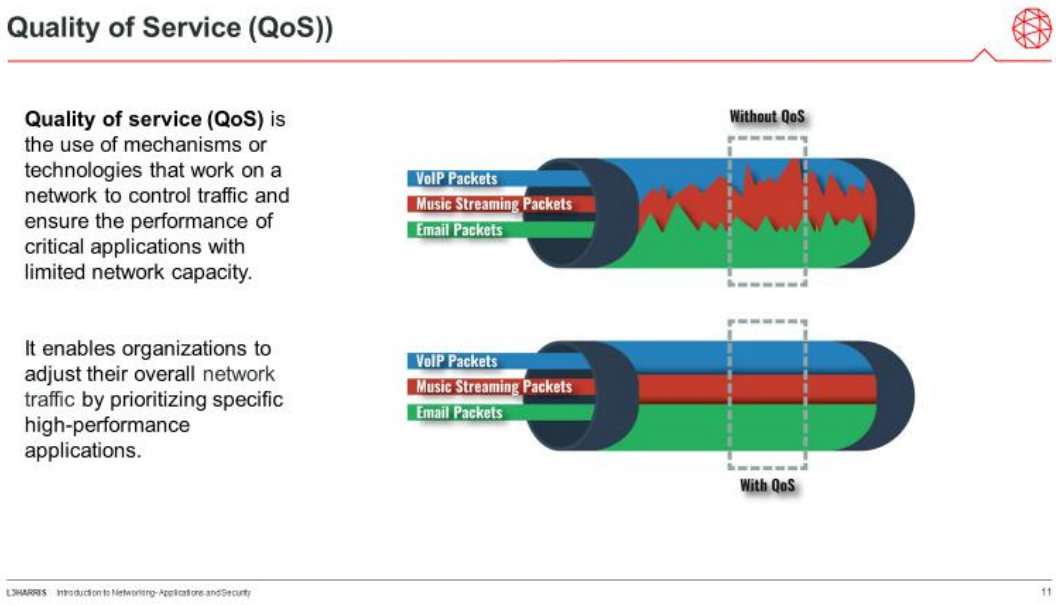

Quality of service (QoS) is the use of mechanisms or technologies that work on a network to control traffic and ensure the performance of critical applications with limited network capacity. It enables organizations to adjust their overall network traffic by prioritizing specific high-performance applications.

Quality of Service (QoS) Mechanism

- **Classification**: applied to the router's interface and classifies if the packet requires QoS implementation or not.
- **Marking**: it marks the packets based on classification. It puts a value on the packet header so that the packet can be easily recognized throughout the network based on its classification.
- **Congestion Management**: prioritizes the transmission of each packet by queuing on each interface.
- **Congestion Avoidance**: drops packets early to avoid congestion.
- **Queuing**: accommodates temporary congestion on a network device's interface by storing excess packets in buffers until bandwidth becomes available.
- **Policing**: enforces rate limit by dropping down or marking the packets.
- **Shaping**: enforces rate limit by delaying the packets and store them in the router's buffer for a certain amount of time.

Class of Service (CoS)

Traffic shaping provides a means to control:

- Volume of traffic being sent into a network in a specified period (bandwidth throttling), or
- Maximum rate at which the traffic is sent (rate limiting), or
- More complex criteria such as generic cell rate algorithm.

CoS	Application
7	Network Control
6	Internetwork Control
5	Voice
4	Video
3	Call Signaling
2	Transactional Data
1	Bulk Data
0	Best Effort

Class Of Service Layer2 Example

Class of Service (CoS) CoS mechanisms such as DiffServ and 802.1p just categorize protocols into groups that require different service levels and provide a tagging mechanism to identify a frame or packet's class.

- **IEEE 802.1p** (Layer 2) protocol that classifies and prioritizes traffic passing over a switch or wireless access point. 802.1p defines a tagging mechanism within the 802.1Q VLAN field
- **DiffServ** (Layer 3) protocol classifies each packet passing through a device. Router policies can then be defined to use the packet classification to prioritize delivery. DiffServ is an IP (layer 3) service tagging mechanism. It uses the Type of Service field in the IPv4 header (Traffic Class in IPv6) and renames it the Differentiated Services field. The field is populated with a 6-byte DiffServ Code Point (DSCP) by either the sending host or by the router. Packets with the same DSCP and destination are referred to as Behavior Aggregates and allocated the same Per Hop Behavior (PHB) at each DiffServ-compatible router

Traffic Shaping: Bandwidth management technique which delays some or all datagrams to bring them into compliance with a desired traffic profile. Traffic shapers delay certain packet types—based on their content—to ensure that other packets have a higher priority. This can help to ensure that latency is reduced for critical applications. This traffic management function can take place on the originating host, or source, but is more usually implemented at the network edge—that is, the perimeter network. This control can be accomplished in many ways and for many reasons; however, traffic shaping is always achieved by delaying packets.

Load Balancer and Multilayer Switching

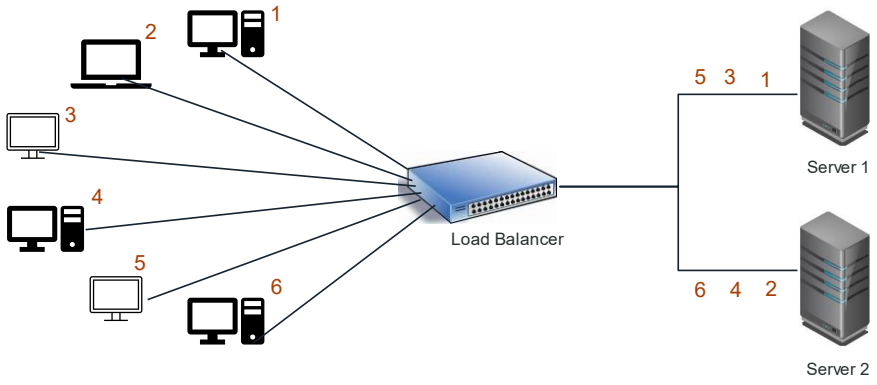

Load Balancers: optimize performance by distributing client requests across available server nodes in a “farm” or “pool”. Clients use the single name/IP address of the load balancer to connect to the servers in the farm.

There are two main types of load balancers:

- **Layer 4 load balancer**—Early instances of load balancers would base forwarding decisions on IP address and TCP/UDP port values this type of load balancer is stateless; it cannot retain any information about user sessions
- **Layer 7 load balancer** (content switch)—As web applications have become more complex, modern load balancers need to be able to make forwarding decisions based on Application-level data, such as a request for a set of URLs or data types like video or audio streaming. This requires more complex logic, but the processing power of modern appliances is sufficient to deal with this.

Multilayer Switches:

Content switches (or multilayer switches) provide switching functionality at upper layers 4 or 4-7 and are used for load balancing applications, typically for the web (HTTP and HTTPS) or SSL based VPNs, although they can be used to switch for any specified TCP or UDP port.

In a common scenario, a multilayer switch would be the interface for a server farm. The switch facilitates connections between clients and servers to optimize performance. This load balancing can take place using defined metrics and rule sets.

- A layer 4 switch applies these rules by inspecting the TCP segment while a layer 4-7 switch can be configured with rules relating to the headers and possibly content of application-layer packets.
- Layer 7 switching allows more fine-grained control but is consequently more difficult to configure, slower, and requires more expensive hardware.

Virtualization

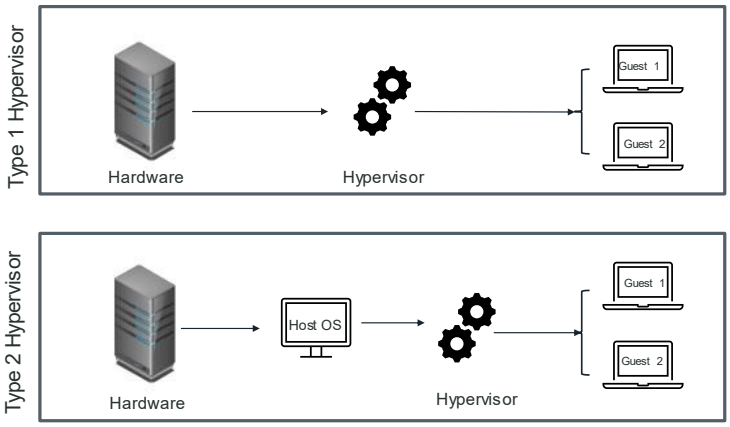

Virtualization: Multiple operating systems can be installed and run simultaneously on a single computer. Benefits include:

- Sever consolidations
- Quickly and easily deploy network functions
- Virtual Desktop Infrastructure (VDI)—Provision client desktop instances as VMs.

A virtual platform requires at least three components:

- **Host(s)—**The platform that will host the virtual environment. Optionally, there may be multiple computers networked together.
- **Hypervisor or virtual machine monitor (VMM**)—Manages the virtual environment and facilitates interaction with the computer hardware and network.
- **Guest operating systems or virtual machines (VMs)—**Operating systems installed under the virtual environment. The number of operating systems is generally only restricted by hardware capacity.

Type 1 Hypervisor (Bare Metal): installed directly onto the computer and manages access to the host hardware without going through a host OS

Type 2 Hypervisor (Host- Based) installed onto a host operating system

virtual NIC (vNIC) will look exactly like an ordinary network adapter and will be configurable in the same way. (Ex. protocols and services can be bound to it, and it can be assigned an IP address.) In other words, a virtual NIC functions identically to a physical NIC for data transmission

Virtual Switch (Vswitch): Functionality is like a physical switch

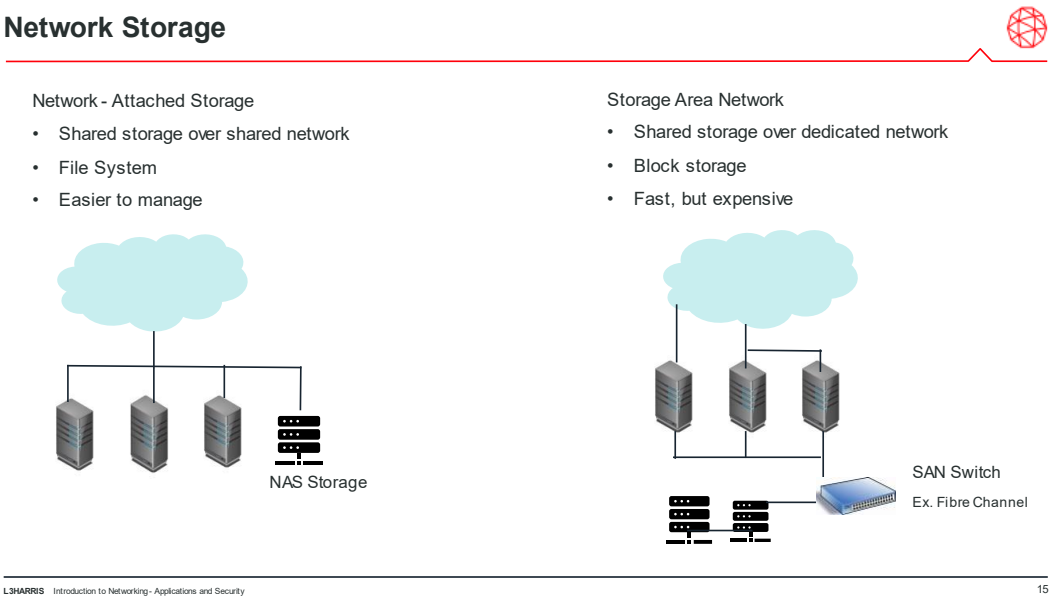

Another element in a virtual platform is **storage virtualization**. In a virtual storage platform, a software layer is inserted between client OSes and applications and the physical storage medium—a **storage area network (SAN).**

- Reduces data deduplication as each user reference to a file can point to the same physical file location (without the user having to track where this might be).
- Simplify operations such as backup, replication, and migration by consolidating data storage in one physical location
- Storage virtualization also assists the implementation of tiered storage hierarchies.

– **Offline Storage** medium might require physical interaction to access the data, such as putting a tape into a drive

– **Nearline storage** refers to technology such as tape loaders or "slow" hard disk media that can operate in low-power states.

Network Attached Storage (NAS) is a data storage mechanism that uses special devices connected directly to the network media. These devices are assigned an IP address and can then be accessed by clients via a server that acts as a gateway to the data, or in some cases allows the device to be accessed directly by the clients without an intermediary.

Storge Area Networks (SANs): A network of storage devices that are connected to each other and to a server, or cluster of servers, which act as an access point to the SAN. In some configurations a SAN is also connected to the network. SAN's use special switches as a mechanism to connect the devices. These switches, which look a lot like a normal Ethernet networking switch, act as the connectivity point for SAN's. Making it possible for devices to communicate with each other on a separate network brings with it many advantages.

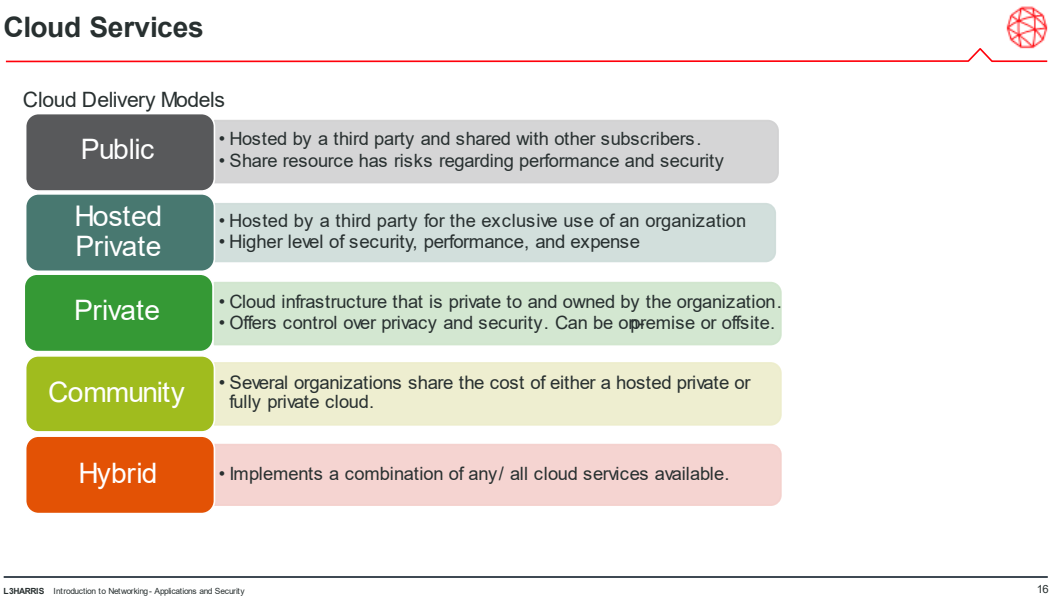

From the consumer point of view, cloud computing is a service that provides OnDemand resources—server instances, data storage, databases, or applications—over a network, typically the Internet. The service is a cloud because the end user is not aware of or responsible for any details of the procurement, implementation, or management of the infrastructure that underpins those resources. The end user is interested in and pays for only the services provided by the cloud.

Cloud Service Types As well as the ownership model cloud services are often differentiated on the level of complexity and pre-configuration provided.

- Infrastructure as a Service (IaaS) Ren I.T. components as needed from service providers data center (Example: Amazon Web Server)
- Platform as a Service (PaaS) Servers and storage infrastructure combined with multi-tier web applications and database platforms. (Example: Azure)
- Software as a Service (SaaS) Software hosted on suppliers' servers with pay as you go and on demand options. (Example: Google Docs)

Connectivity

- **INTERNET/VIRTUAL PRIVATE NETWORK (VPN) The** simplest way of interfacing with a cloud service is to use the provider's website or application programming interface (API) over the Internet.
- **DIRECT/PRIVATE CONNECTION/CO-LOCATION:** Co-location within a data center offers a higher bandwidth solution by providing a direct or private link. The customer establishes infrastructure within a data center supported by the cloud provider or provisions a direct link from his or her enterprise network to the data center, possibly using a service provider's MPLS network.

Network Security

- Define network security concepts and best practices
- Recognize common network threats.
- Identify mitigation hardware, protocols and services.

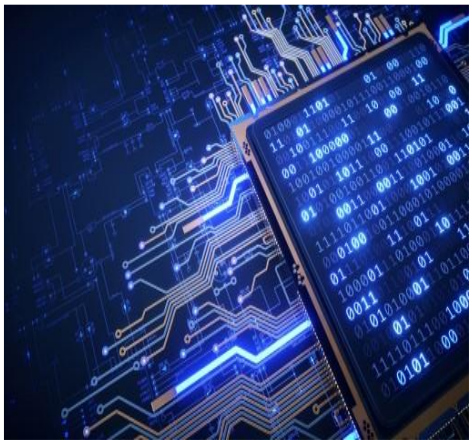

Notes:

CIA Triad

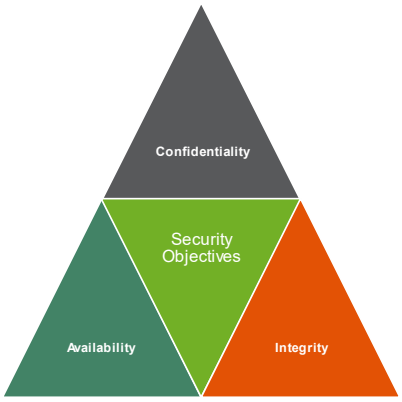

One of the foundational principles of computer security is that the systems used to store, transmit, and process data must demonstrate three properties, often referred to as the CIA Triad.

Confidentiality: Certain information should only be known to certain people.

- Access controls: Selectively restrict access to a resource.
- Encryption: Encode messages so only certain people can read it
- Steganography: Conceal information with another piece of information.

Integrity: The data is stored and transferred as intended, and any modification is authorized.

- Hashing: Map of data of an arbitrary length to data of a fixed length.
- Digital signatures: Mathematical scheme to verify the integrity of data
- Certificates: Combine with a digital signature to verify an individual.
- Non- Repudiation: Provides proof of integrity, can be asserted to be genuine.

Availability: Information is accessible to those authorized to view or modify it. Many tools and techniques are available to ensure that network systems demonstrate these three key properties.

- Redundancy: Build services that will always be available
- Fault Tolerance: System will continue to run, even when failure occurs
- Patching: Ensure stability and close security holes

Security policies ensure that an organization has evaluated the risks it faces and has put **security controls** in place to mitigate those risks. Making a system more secure is also referred to as **hardening.** Different security policies should cover every aspect of an organization's use of computer and network technologies, from procurement and change control to acceptable use.

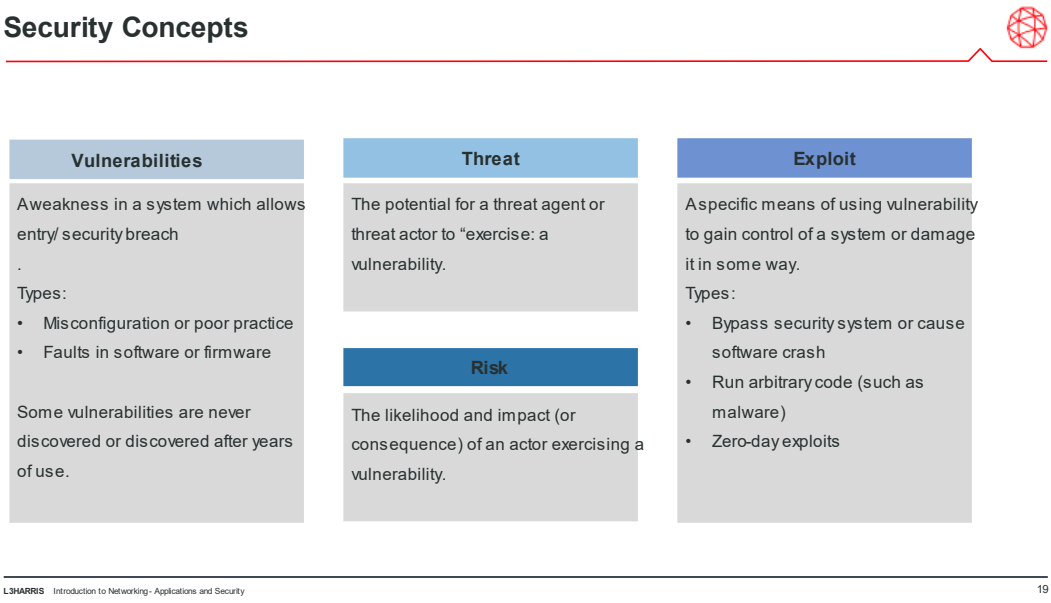

Vulnerabilities can exist because of misconfigurations or poor practice, but many people understand the term to mean faults in software specifically. This type of software vulnerability is a design flaw that can cause the application security system to be circumvented or that will cause the application to crash.

The code or method by which an attacker uses a vulnerability is called an **exploit**. Typically, software vulnerabilities can be exploited only in quite specific circumstances, but because of the complexity of modern software and the speed with which new versions must be released to market, almost no software is free from vulnerabilities.

A vulnerability that is exploited before the developer knows about it or can release a patch is called a **zero-day exploit**. These can be extremely destructive, as it can take the vendor a lot of time to develop a patch, leaving systems vulnerable for days, weeks, or even years.

A greater threat is the large number of unpatched or legacy systems in use. An unpatched system is one that its owner has not updated with OS and application patches; a legacy system is one where the software vendor no longer provides support or fixes for problems.

- U.S. National Vulnerability Database (NVD)
  - Sponsored by the U.S. Department of Homeland Security and Cybersecurity and Infrastructure Security Agency
  - Https://nvd.nist.gov/

Note: This issue does not just affect PCs. Network appliances can also be vulnerable to exploits. The risks to embedded systems have become more obvious over the last few years, and the risks posed by unpatched mobile devices and the Internet of Things is likely to grow.

Common Networking Attacks

A network can be attacked by many kinds of intruders or adversaries for many different reasons. The goals of most types of adversaries will either be to steal (exfiltrate) information from the network to misuse network services (for fraud, for instance), or to damage it. Insider threat-type attacks may be launched with privileged access to the network, while external threats must find some way of accessing the network, perhaps by installing malware on a host system. Network attacks can then proceed in several ways.

- **Foot printing** is a process of information gathering, in which the attacker attempts to learn about the configuration of the network and security systems.
- **Port scanning** specifically aims to enumerate the TCP or UDP application ports that are "open" on a host.
- **Eavesdropping (or sniffing**) refers to capturing and reading data packets as they move over the network.
- **Spoofing:** can include any type of attack where the attacker disguises his or her identity, or in which the source of network information is forged to appear legitimate.
- **Social engineering:** (or hacking the human) refers to means of getting users to reveal confidential information.
- **IP spoofing attack**: the attacker changes the source address recorded in the IP packet.
- **A man-in-the-middle (MitM**): attack is a specific type of spoofing attack where the attacker sits between two communicating hosts and transparently intercepts and relays all communications between them.

Common Networking Attacks Cont.

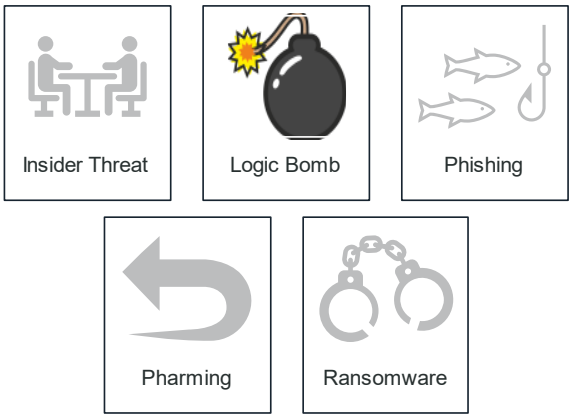

- **Insider threat** means attacks launched by the organization's own trusted users (employees, partners, or contractors).
- A **logic bomb** is a type of malware that executes in response to a system or user event. A typical example is a disgruntled system administrator who leaves a scripted trap that runs when his or her account is deleted or disabled
- **Phishing** is a combination of social engineering and spoofing (disguising one computer resource as another).
- **Spear phishing** refers to a phishing scam where the attacker has some information that makes an individual target more likely to be fooled by the attack.
- **Pharming** is another means of redirecting users from a legitimate website to a malicious one. Rather than using social engineering techniques to trick the user, pharming relies on corrupting the way the victim's computer performs Internet name resolution, so that they are redirected from the genuine site to the malicious one.
- **Ransomware** is a type of malware that tries to extort money from the victim. One class of ransomware will display threatening messages, such as requiring Windows to be reactivated or suggesting that the computer has been locked by the police because it was used to view child pornography or for terrorism.

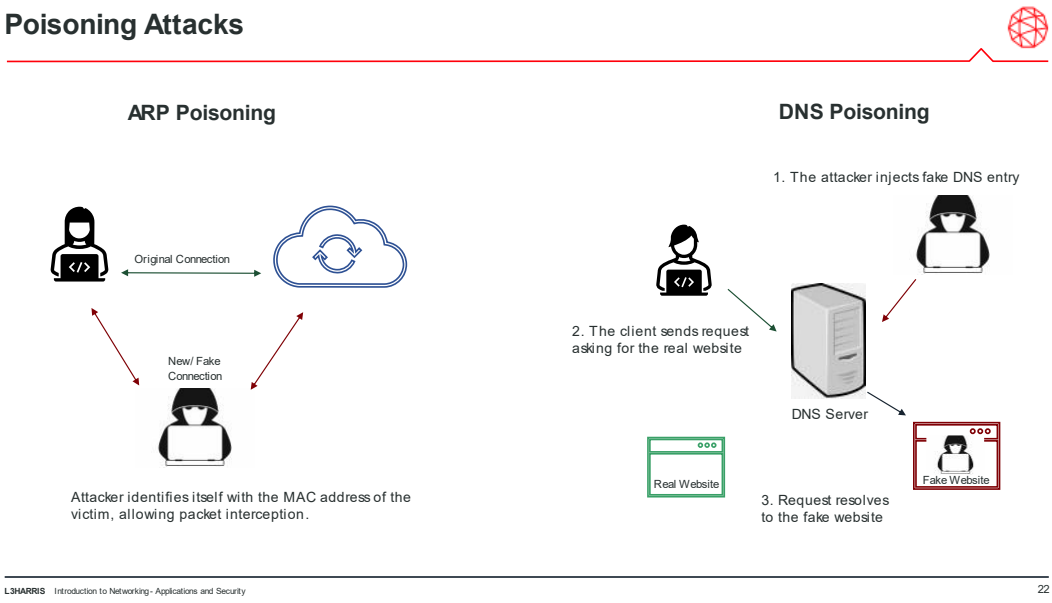

ARP Poisoning:

An ARP cache poisoning attack is a common means of perpetrating a MitM attack.

- It works by broadcasting unsolicited ARP reply packets with a spoofed source address.
- Because ARP has no security, the receiving devices trust this communication and updates their MAC: IP address cache table with the spoofed address.
- An ARP spoofing attack can be launched by running software such as Dsniff, Cain and Abel, or Ettercap from a host attached to the same broadcast domain as the target.

DNS Poisoning Attacks:

DNS poisoning is an attack that compromises the name resolution process.

- Typically, the attacker will replace the valid IP address for a trusted website, such as mybank.com, with the attacker's IP address.
- The attacker can then intercept all the packets directed to mybank.com and bounce them to the real site, leaving the victim unaware of what is happening (referred to as pharming).
- Alternatively, DNS spoofing could be used for a DoS attack by directing all traffic for a particular FQDN to an invalid IP address (a black hole).
- One way to attack DNS is to corrupt the client's name resolution process by changing the servers used for resolving queries, intercepting, and modifying DNS traffic, or polluting the client's name cache (by modifying the HOSTS file, for instance).
- DNS server cache poisoning (or pollution) is another redirection attack, but instead of trying to subvert the name service used by the client, it aims to corrupt the records held by the DNS server itself.

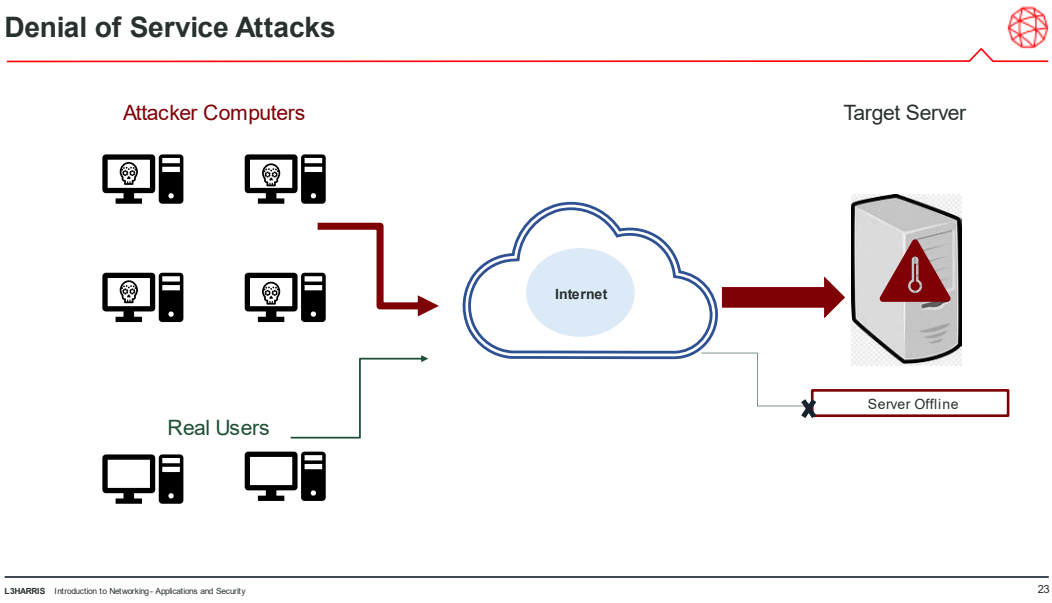

Denial of Service (DoS) Attacks

A denial of service (DoS) attack causes a service at a given host to fail or to become unavailable to legitimate users. Typically, DoS attacks focus on overloading a service by using up CPU, system RAM, disk space, or network bandwidth (resource exhaustion). It is also possible for DoS attacks to exploit design failures or other vulnerabilities in application software.

A common method of attack is when the attacker sends several requests to the target server, overloading it with traffic. These service requests are illegitimate and have fabricated return addresses, which mislead the server when it tries to authenticate the requestor. As the junk requests are processed constantly, the server is overwhelmed, which causes a DoS condition to legitimate requestors.

- A **SYN flood** occurs when an attacker sends a request to connect to the target server but does not complete the connection. The incomplete handshake leaves the connected port in an occupied status and unavailable for further requests. An attacker will continue to send requests, saturating all open ports, so that legitimate users cannot connect.

Distributed DoS (DDoS) Attacks and Botnets: Most bandwidth-directed DoS attacks are distributed. This means that the attacks are launched from multiple, compromised computers.

- **Distributed Reflection DoS (DRDoS):** In this attack, the adversary spoofs the victim's IP address and attempts to open connections with multiple servers.
- **Smurf Attack**: the attacker sends Internet Control Message Protocol broadcast packets to a number of hosts with a spoofed source Internet Protocol (IP) address that belongs to the target machine. The recipients of these spoofed packets will then respond, and the targeted host will be flooded with those responses.

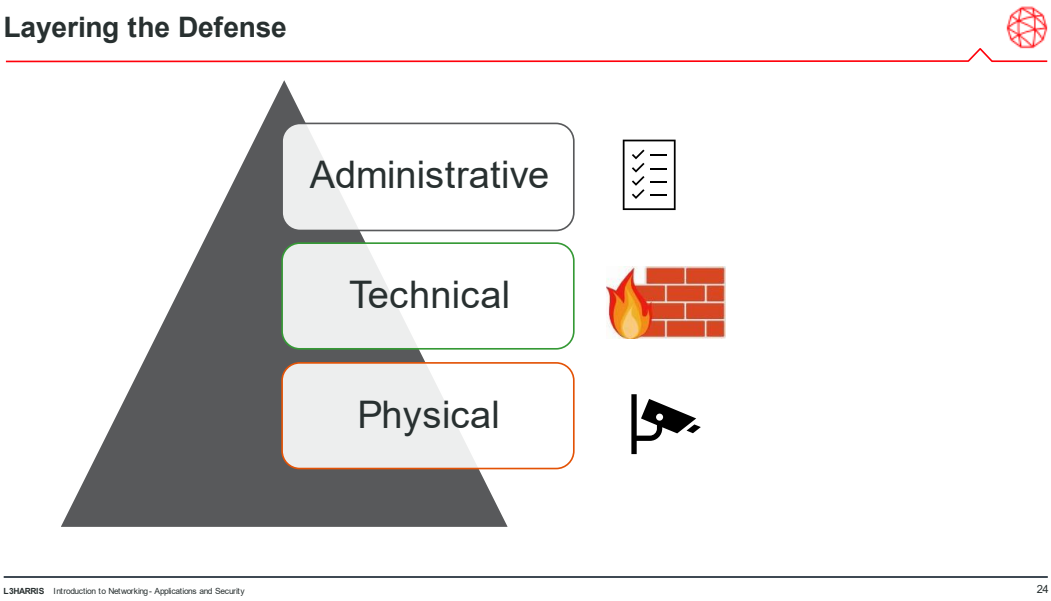

There are many ways in which networks can be attacked and just as many ways for making networks more secure. This requires a layered defensive approach.

Physical controls:

- Keep people away from technology
- Door locks, fences, rack locks, cameras

Technical controls:

- Hardware and software to keep things secure
- Firewalls, active directory authentication, disk encryption, VPN, antivirus/ anti malware

Physical Segmentation:

- vLAN, screened subnet (DMZ), network access control
- Administrator physically disable unused ports

Administrative controls:

- Policies and procedures
- Back up media handling
- Separation of Duties

Network Segmentation

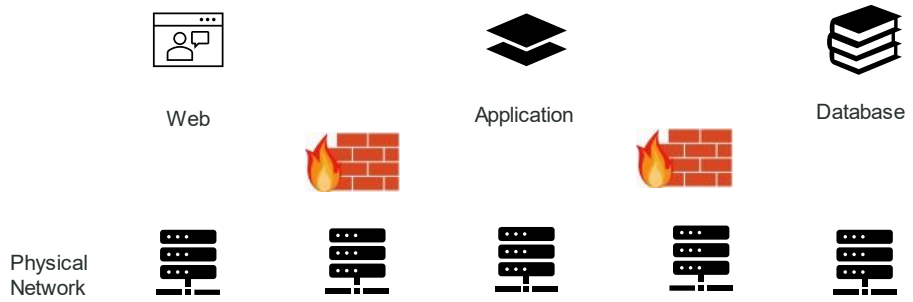

Modern local networks are built using switches. In its default configuration, every port on a switch will be in the same local segment or, put another way, in the same broadcast domain. Any host within a broadcast domain can contact any other host using the same logical addressing scheme (IP subnet) and by hardware/MAC addressing. With too many hosts attached to the same switch, broadcast traffic can become excessive and reduce performance. Also, nodes within the same broadcast domain can be vulnerable to attacks such as ARP spoofing and worm malware. Breaking up the broadcast domains helps to limit the scope of these kinds of attacks.

Network Segmentation: The process of determining which parts of the network are accessible to other parts. Segmentation can mitigate an attack by restricting it to a smaller group of network hosts. The technologies that can be used to enforce network segmentation include virtual LANs (VLANs), subnets, virtual private networks (VPNs), and host virtualization.

All these types of systems enforce network segmentation by deploying access controls of different types. The networks are not physically separate, but they are logically separate. For example, when a host is assigned to a VLAN, the switch restricts it to communications designated for that VLAN. To communicate outside the VLAN, the host must use a router, and a router equipped with a firewall can apply additional rules to what it allows in and out.

- Segment network by creating zones with different security/access configurations
- “Containers” can be created using Virtual LAN (VLAN), subnets, Virtual Private Networks (VPN), or host virtualization
- Traffic between “containers” subject to Access Control Lists (ACL), typically enforced by firewalls
- Air gapped network or host

Virtual Local Area Network (VLAN)

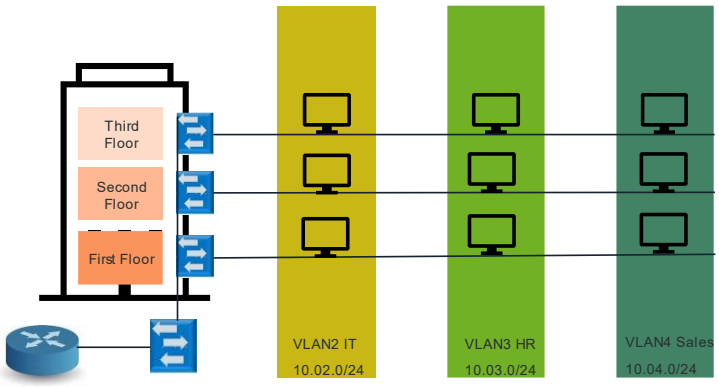

VLAN: Allows different groups of computers on the same cabling and attached to the same switches can appear to be in separate LAN segments. Nodes in each VLAN are in separate broadcast domains. As well as reducing the impact of broadcast traffic, from a security point of view, each VLAN can represent a separate zone. These zones would typically be configured to protect the integrity and confidentiality of different departments within the organization.

VLAN Benefits

- If something like a virus or worm were introduced in one VLAN, it should not be able to spread to other VLANs
- Prevent unauthorized users from accessing the data in a particular VLAN.
- Option to create a “null” VLAN that is non-routable in the network to be used for any physical ports that do not have authorized connected equipment.
- Separate nodes based on traffic type and the need for Quality of Service.
- Overcome limitations imposed by the physical location of a host and assign them to the appropriate logical group of hosts.

The simplest way to configure a VLAN is by configuring the port interface on the switch.

- Each VLAn is typically configured with its own subnet address and IP address range.
- Communications between VLANs must go through an IP router
- Port based assignment is described as a static VLAN
- Nodes or hosts can also be assigned to dynamic VLANs using some feature of the host, such as a its MAC address or authentication credentials supplied by the user.

VLAN Trunking

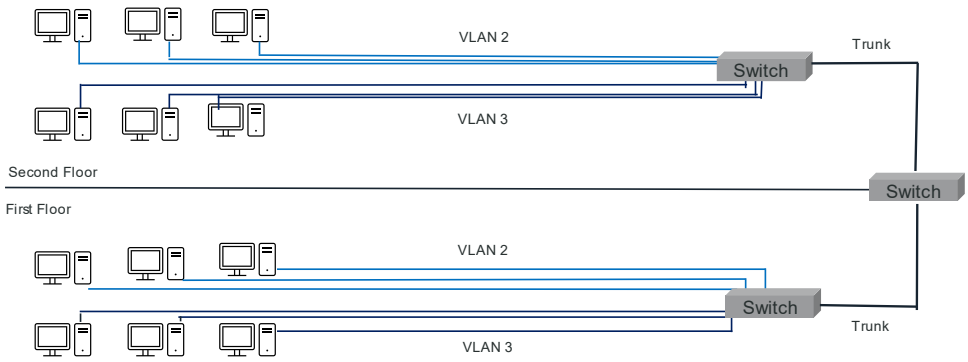

Trunking means that VLANs can be configured across more than one switch device without having to configure the VLANs on each device manually. This means that hosts connected to different switches (and perhaps in completely different locations) can be part of the same local network. The VLANs can be used to define organizational boundaries without having to put hosts in the same physical location.

Pruning refers to removing transmissions related to designated VLANs from a trunk to preserve bandwidth. If a VLAN is not associated with a given trunk link, pruning it from the trunk reduces the amount of broadcast traffic passing over the link. There may also be security reasons for removing a VLAN from a trunk link. Pruning can either be done via VTP or by configuring the trunk manually.

VLAN Assignment Issues:

- Each VLAN is a discrete broadcast domain. Ensure that services such as name resolution and IP autoconfiguration are properly available to all VLANS.
- Ensure those troubleshooting the network have an accurate and up to date map of physical and logical network infrastructure.
- Configuration errors (ie trunking, layer 3 routing config and incorrect host placement) can cause hosts to isolate.
- VLAN to VLAN communications must be configured for devices on two separate VLANs to communicate.

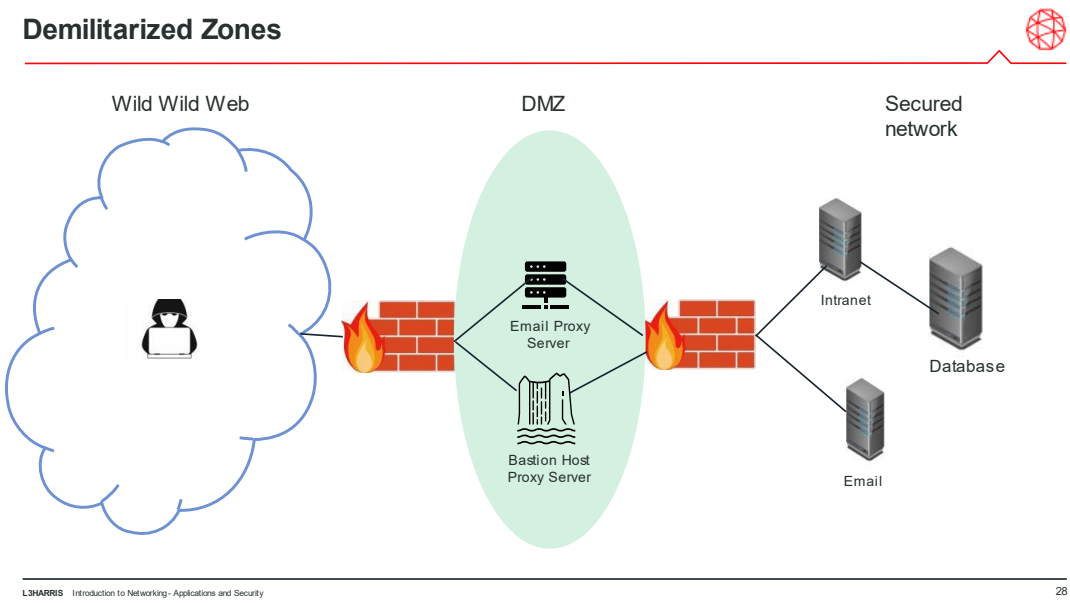

Demilitarized Zone (DMZ) is a perimeter network that protects and adds an extra layer of security to an organization's internal local-area network from untrusted traffic.

A DMZ enables external clients to access data on private systems, such as web servers, without compromising the security of the internal network. If communication is required between hosts on either side of a DMZ, a host within the DMZ acts as a proxy.

Bastion Host: Hosts in the DMZ are not fully trusted by the internal network because of the possibility that they could be compromised from the Internet and would not be configured with any services that run on the local network, such as user authentication.

Screened Subnet: One important use of subnets is to implement a DMZ. In a screened subnet, two firewalls are placed at either end of the DMZ. One restricts traffic on the external interface; the other restricts traffic on the internal interface.

Three-Legged Firewall: A DMZ can also be established using a single router/firewall appliance. A three-legged firewall (or triple-homed firewall) is one with three network ports, each directing traffic to a separate subnet.

Screened Host: Smaller networks may not have the budget or technical expertise to implement a DMZ. In this case, Internet access can still be implemented using a dual-homed proxy/ gateway server acting as a screened host.

SOHO DMZ/DMZ Host: Sometimes the term DMZ (or DMZ host) is used by SOHO Internet router vendors to mean an Internet-facing host or zone not protected by the firewall. This might be simpler to configure and solve some access problems, but it makes the whole network vulnerable to intrusion and DoS. A true DMZ is established by a separate network interface and subnet so that traffic between hosts in the DMZ and the LAN must be routed (and subject to firewall rules). Most SOHO Internet routers do not have the necessary ports or routing functionality to create a true DMZ.

Packet Filtering Firewall- Stateless

Rule	Direction	Src Address	Dest address	Protocol	Dest Port	Action
1	In	105.15.33.10	10.1.2.09	TCP	80	Allow
2	Out	10.1.2.09	105.15.33.10	TCP	80	Allow
3	Either	Any	Any	Any	Any	Deny

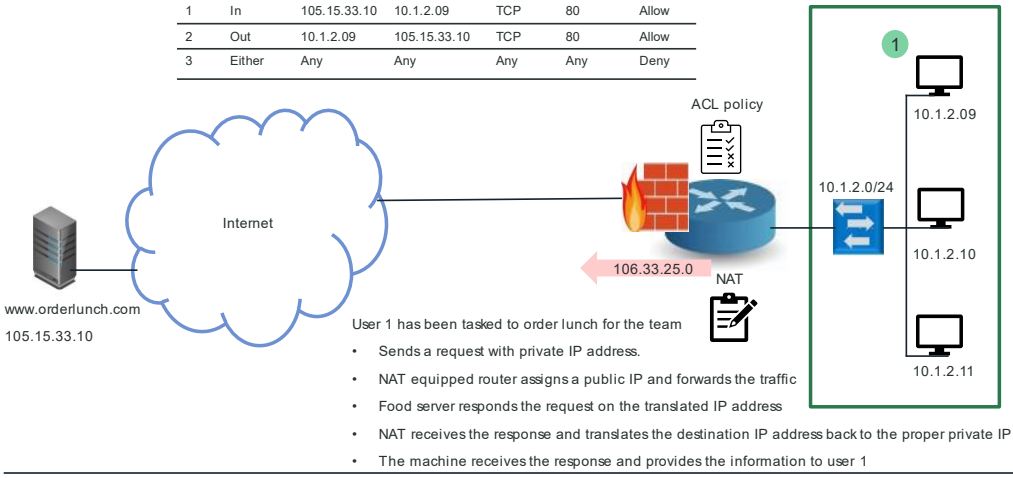

Firewalls are the devices principally used to implement security zones by processing traffic according to rules; traffic that does not conform to a rule that allows it access is blocked. There are many types of firewalls and many ways of implementing a firewall to fit your networks configuration.

- Firewalls that protect a whole network or firewalls that protect a single host only.
- Border firewalls filter traffic between the trusted local network and untrusted external networks, such as the Internet.
- Control only inbound “ingress” or both inbound and outbound “egress” traffic.
- A further distinction can be made about what parts of a packet a particular firewall technology can inspect and operate on.

Packet Filtering Firewall: Configured by specifying rules which are called an access control list (ACL). Each rule defines a specific type of data packet and the appropriate action to take when a packet matches the rule. A packet filtering firewall can inspect the headers of IP packets. This means that rules can be based on the information found in those headers. Ex:

National Address Translation (NAT): A service translating between private (local) addressing scheme used by hosts on the LAN or a DMZ and a public (global) addressing scheme used by an internet-facing device and is configured on a border device, such as a router, proxy server or firewall.

- **Dynamic NAT**: NAT service builds a table of public to private address mappings. Each new session creates a new public-private address binding in the table. When the session is ended or times out, the binding is released for use by another host.

Port Address Translation (PAT): Allocates each new connection a high-level TCP or UDP port by creating a port map in its state table, substituting the private IP for the public Ip and forwards the request to the internet.

Stateful Inspection Firewalls

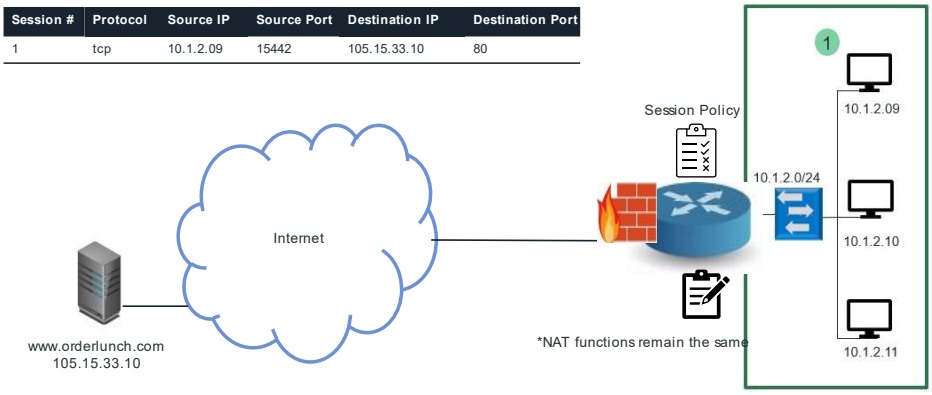

Circuit Level Gateway: Maintains stateful information about the session established between two hosts in the dynamically updated state table.

- A stateful firewall operates at session layer 5 of the OSI Model.
- When a packet arrives, the firewall checks it to confirm whether it belongs to an existing connection.
  - If it does not, it applies the ordinary packet filtering rules to determine whether to allow it
  - Once the connection has been allowed, the firewall allows traffic to pass unmonitored, to conserve processing effort.
  - A circuit-level firewall examines the TCP three-way handshake and can detect attempts to open connections maliciously.
  - It also monitors packet sequence numbers and can prevent session hijacking attacks. It can respond to such attacks by blocking source IP addresses and throttling sessions.

Next Generation Firewall (NGFW)

While a traditional firewall typically provides stateful inspection of incoming and outgoing network traffic, a next generation firewall includes additional features like application awareness and control, integrated intrusion prevention, and cloud-delivered threat intelligence.

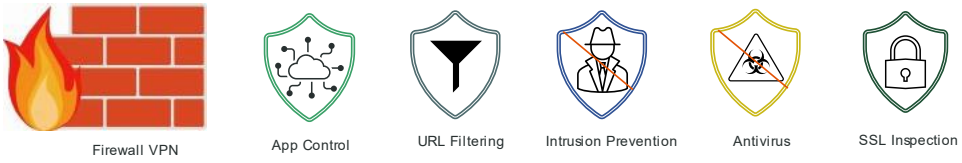

NGFW/Layer 7 Firewalls: Can inspect and parse the contents of packets at the Application layer (7) of the OSI model, all data in every packet. **You may hear this referred to as the** Application Layer gateway, stateful multilayer inspection, or deep packet inspection.

Application-based Firewall Software designed to run on a server to protect a particular type of application, rather than just general access to the host.

Network based firewalls

- Control traffic flows based on application

Intrusion prevention systems

- Identify the application
- Apply application-specific vulnerability signatures to the traffic

Content Filtering: URL Filtering

- Filtering can be applied to a mix of permitted/restricted URLs, keyword matching, web object matching (looking at usage of plug-ins), time of day use, total usage, and so on.

Web application firewall (WAF): Analyzes the header and the HTML code present in HTTP packets to try to identify code that matches a pattern in its threat database.

Forward Proxy Servers: Rather than inspecting traffic as it passes through, the proxy deconstructs each packet, performs analysis, then rebuilds the packet and forwards it on, providing it conforms to the rules.

Reverse Proxy Servers: provides for protocol-specific inbound traffic. Typical applications for reverse proxy servers include publishing a web server, publishing IM or conferencing applications, and enabling POP/IMAP mail retrieval.

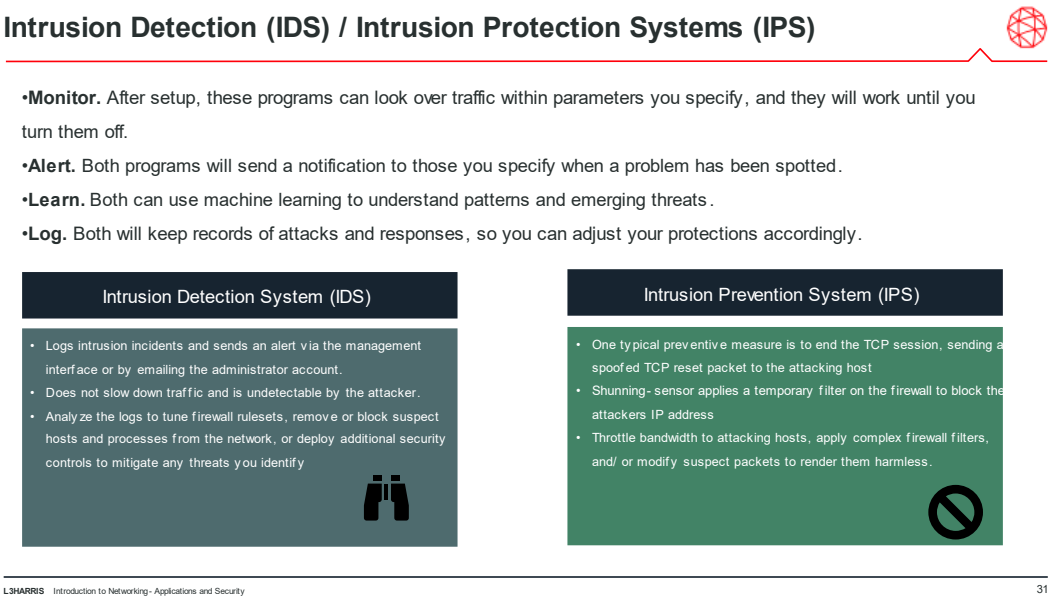

Intrusion Detection System (IDS) is a means of using software tools to provide real-time analysis of either network traffic or system and app. A network IDS (NIDS) is basically a packet sniffer (referred to as a sensor) with an analysis engine to identify malicious traffic and a console to allow configuration of the system.

Intrusion Prevention System (IPS): Provides an active response to any network threats that it matches. Positioned at the border between two network zones- “inline” meaning all traffic passes through.

Host Based IDS and IPS: Captures information from a single host, such as a server, router, or firewall. Installing HIDS/HIPS is simply a case of choosing which hosts to protect, then installing and configuring the software. There will also normally be a reporting and management server to control the agent software on the hosts.

Unified Threat Management: Refers to a system that centralizes various security controls—firewall, anti-malware, network intrusion prevention, spam filtering, content inspection, etc.—into a single appliance.

+UTM systems help to simplify the security process by being tied to only one vendor and requiring only a single, streamlined application to function

-When defense is unified under a single system, this creates the potential for a single point of failure that could affect an entire network. Potential latency issues with heavy network activity.

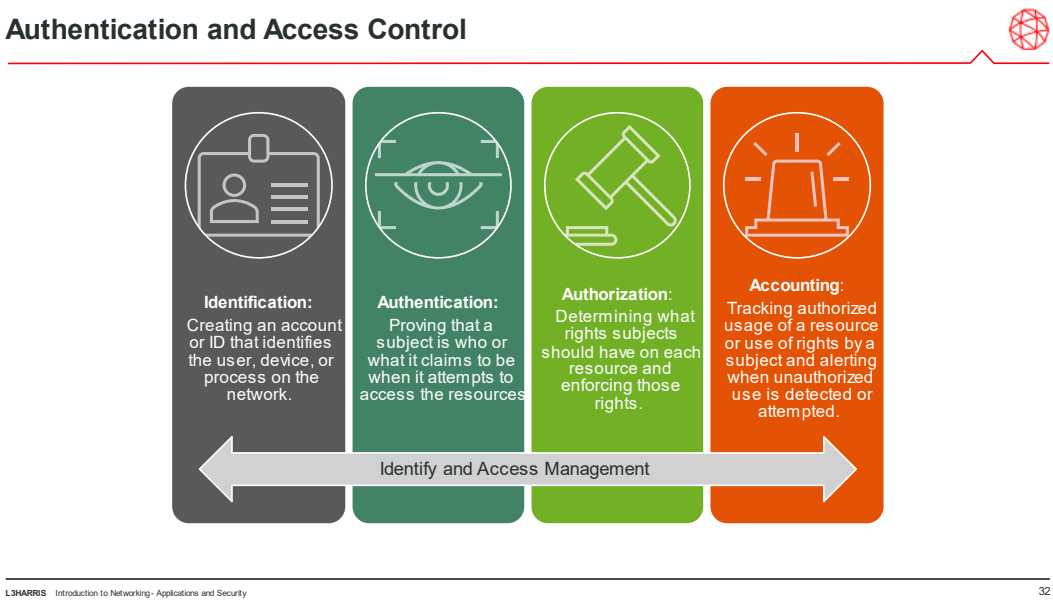

Access Control List (ACL): Firewall and protection systems are structured by a set of technical controls that govern how subjects (ex: users, devices, software processes) may interact with objects (ex. networks, databases, servers, files). This rule-based management is configured on the principle of least access.

- The rules in a firewall's ACL are processed top-to bottom. If traffic matches one of the rules, then it is allowed to pass; consequently, rules that are most specific and that must override other rules (if there is a conflict) are placed at the top.

Identify and access management (IAM) system mediates use of objects by subjects by enabling a system administer to define the attributes that make up an entities identify. These attributes subsequently enable access management systems to make informed decisions about whether to grant or deny an entity access, and if granted, decide what the entity has authorization to do.

Trusted users: Users such as employees and approved partners or contractors

Untrusted users: Users such as site visitors, customers, and suppliers

Role Based Access Control (RBAC):

- Least privileges: Employees network rights and permissions should be specific to their job role and do not allow access outside of the scope of what is needed to perform their job.
- Each employee's access should be limited to the least amount of access required to perform their job.
- Applications should only run with the specific privileges required to run that application and allow amount of data provided to be based on job function.
- Separation of duties states that no one person should have too much power or responsibility.

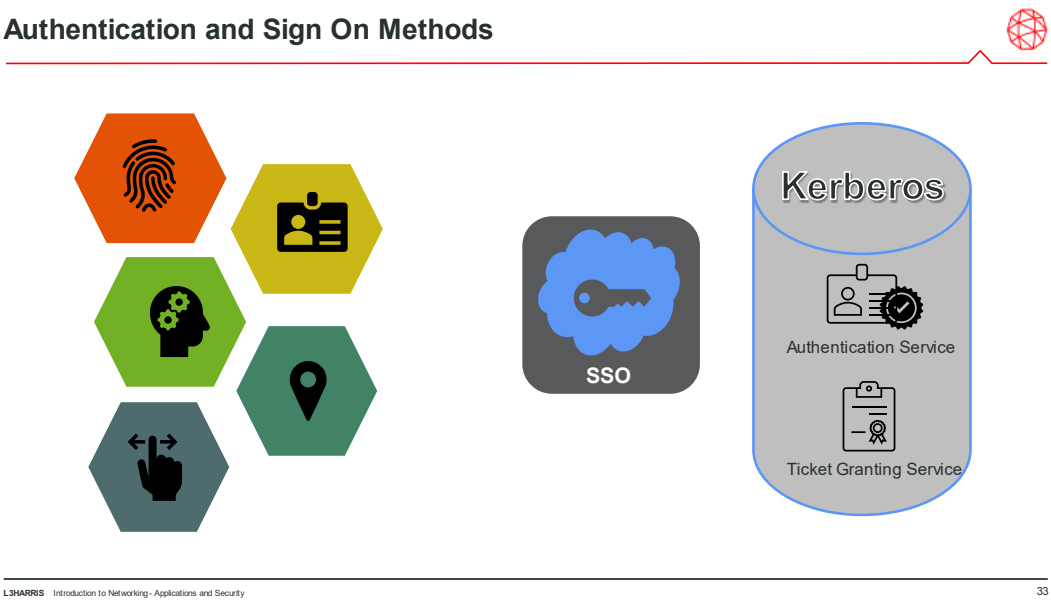

Authentication Factors

Something you are (ex. Fingerprint, facial recognition)

Something you have (ex. Smart card, key fob, hardware token generating a one-time password)

Something you know (ex. username and password, often used for account reset mechanisms)

Somewhere you are (ex. location base services, GPS, IP Address)

Something you do (ex. refers to behavioral biometric recognition)

Multi-factor authentication (MFA) Two- Factor authentication example- Something you have (ie. Cell phone to receive confirmation text) and something you know (ie. Pin or password)

Single sign on (SSO): A user must authenticate to a system only once to gain access to all its resources. This tends to be implemented on enterprise networks.

- **Kerberos:** provides SSO. Once authenticated, a user is trusted by the system and does not need to re-authenticate to access different resources.
  - **Clients** request services from a **server,** which both rely on an intermediary—a **Key Distribution Center (KDC)**—to vouch for their identity. KDC consists of Authentication Service and Ticket Granting Service.

Local Authentication: Primary method for logging into systems when onsite, generally requiring username and password credentials.

LAN Manager/NTLM: is a challenge/response authentication protocol using an encrypted hash of the user's password. This means that the user's password is not sent to the server in plaintext and cannot (in theory) be obtained by an attacker.

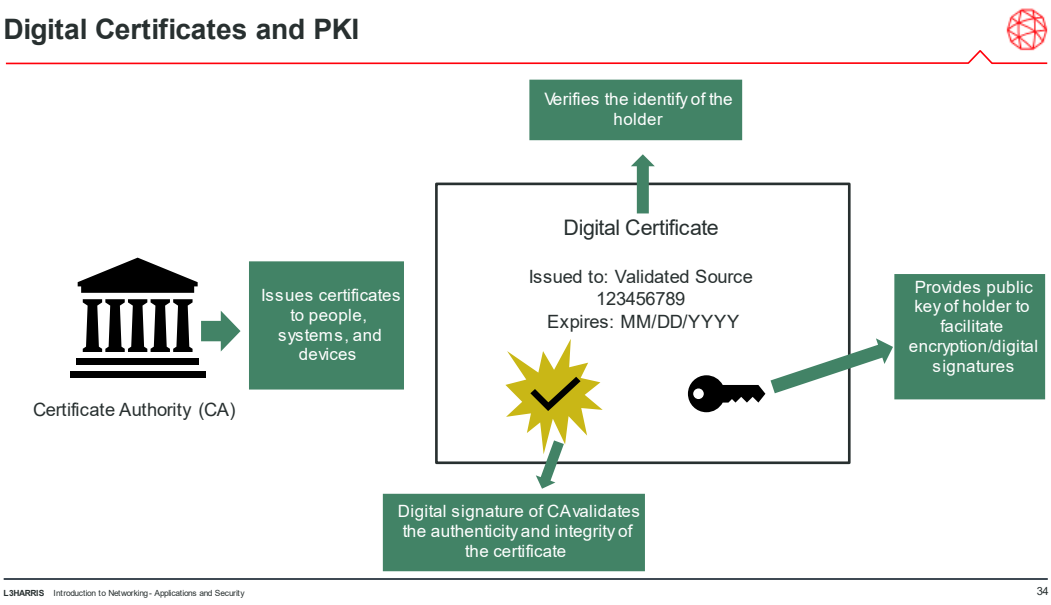

Digital certificates: are used in protocols such as Kerberos and Transport Layer Security (TLS).

- A certificate can be installed on a web server or email server to validate its identity and establish a secure transmission channel.

Digital certificates depend on the concept of public key cryptography. Public key cryptography, also referred to as **asymmetric encryption**, solves the problem of distributing encryption keys when you want to communicate securely with others, authenticate a message that you send to others, or authenticate yourself to an access control system. With asymmetric encryption, you generate a key pair. The private key in the pair remains a secret that only you know. The public key can be transmitted to other subjects. The private key cannot be derived from the public key.

Public key infrastructure (PKI) aims to prove that the owners of public keys are who they say they are.

- Anyone issuing public keys should obtain a digital certificate.
- The validity of the certificate is guaranteed by a certificate authority (CA).
- A digital certificate is essentially a wrapper for a subject's (or end entity's) public key. As well as the public key, it contains information about the subject and the certificate's issuer or guarantor.
- The certificate is digitally signed to prove that it was issued to the subject by a particular CA

Authentication, Authorization and Accounting (AAA) Servers

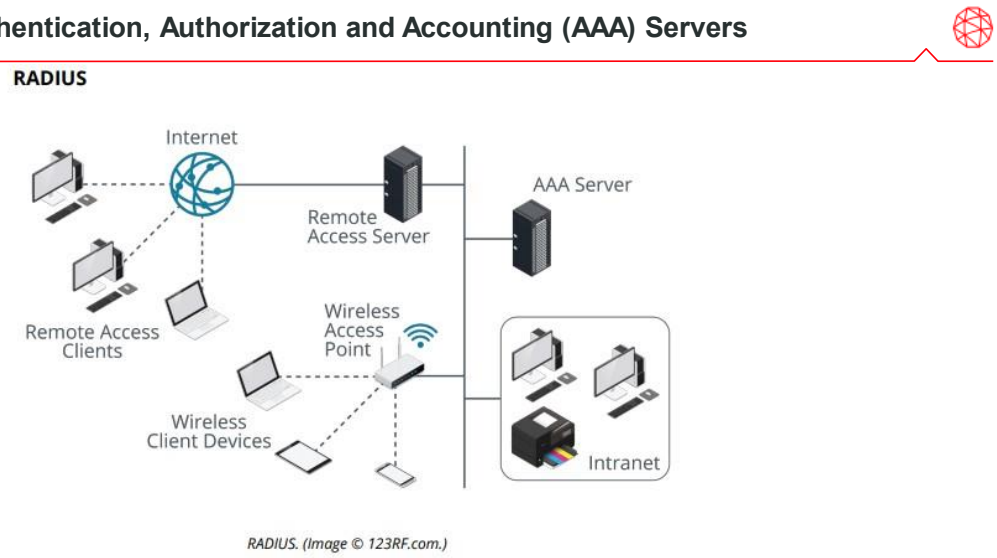

AAA servers were developed to mediate authentication operations between network clients, network access devices, and user authentication and credential management servers. This is accomplished by providing the following functions:

- Validate user credentials for authentication, AAA
- Transmit authorizations for rights and permissions
- Collect accounting information from the user session.

Remote Authentication Dial-in User Service (RADIUS) A standard protocol used to manage remote and wireless authentication infrastructures.

Terminal Access Controller Access Control System (TACACS+) is a similar protocol to RADIUS but designed to be more flexible and reliable. TACACS+ was developed by Cisco but is also supported on many of the other third-party and open-source RADIUS server implementations.

Lightweight Directory Access Protocol: A directory is like a database, where an object is like a record, and things that you know about the object (attributes) are like fields.

Auditing and Logging: The accounting function in AAA is generally performed by logging subject and object activity. All NOSs and many applications and services can be configured to log events. The main decision is which events to record. Logs serve the following general purposes:

- Accounting for all actions that have been performed by users. Change and version control systems depend on knowing when a file has been modified and by whom. Accounting also provides for non-repudiation.
- Detecting intrusions or attempted intrusions. Here records of failure-type events are likely to be more useful, though success-type events can also be revealing if they show unusual access patterns

End Point and Port Security

Port based Network Access Control (PNAC)

- The device (supplicant) requests access to the network
- The switch (authenticator) enables EAPoL protocol and waits for the device to supply authentication data
- The authenticator passes the data to an authentication server which checks the credentials and grants or denies access.
- Access Granted: The switch will configure the port to the appropriate VLAN and enable it for network traffic
- Access Denied: May be denied ay type of access or placed in a guest VLAN with limited access

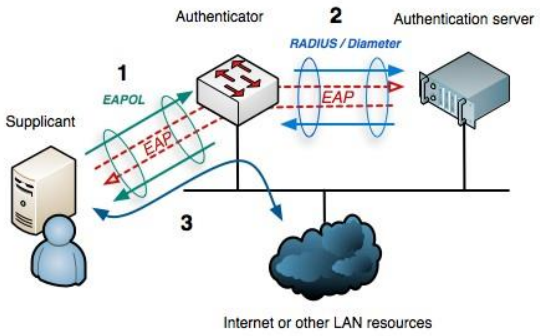

Image

MAC Filtering: Defining which MAC addresses are permitted to connect to a particular port.

- Prevent unauthorized users from connecting to a switch interface
- Each port has its own configuration- Example configure a maximum number of source MAC addresses on an interface

Dynamic ARP Inspection (DAI) prevents a host attached to an untrusted port from flooding the segment with gratuitous ARP replies.

- Relies on DHCP snooping for intel
- ARP inspection maintains a trusted database of IP: ARP mappings

Port Based Network Access Control (PNAC) provides an authentication mechanism to devices wishing to attach to a LAN or VLAN

- IEEE 802.1X Standard

Network Access Control (NAC)

- Set minimum- security policy devices must meet to be granted network service
  - Health Policy: (ie, malware infection, firmware, and OS patch level, etc)
- Post Admission Control via NAC Policy server
- Redemption: refers to what happens if the device does not meet the security profile. A non-compliant device may be refused connection completely or put in a quarantined guest network or captive portal

IMPLEMENT DEVICE HARDENING

- Change default device credentials on installation and ensure that accounts are secured with strong passwords.
- Use only secure channels for administration traffic or any other protocol where credentials need to be submitted.
- Configure services according to the device's baseline and disable any services which are not required. Consider setting up alerting mechanisms to detect service configuration changes.
- Ensure that only the necessary IP ports (TCP and UDP ports) to run permitted services are open and that access to a port is controlled by a firewall ACL if appropriate.
- Ensure any physical ports on a device that could be used to attach unauthorized USB or network devices are protected by secure access (in a locked room or cabinet).
- If a device is decommissioned, ensure that encryption keys stored on that device are archived (if appropriate) and then securely deleted from the device storage.
- Change encryption keys used to access servers and appliances when employees with credentials to access these devices leave the company or change job roles.

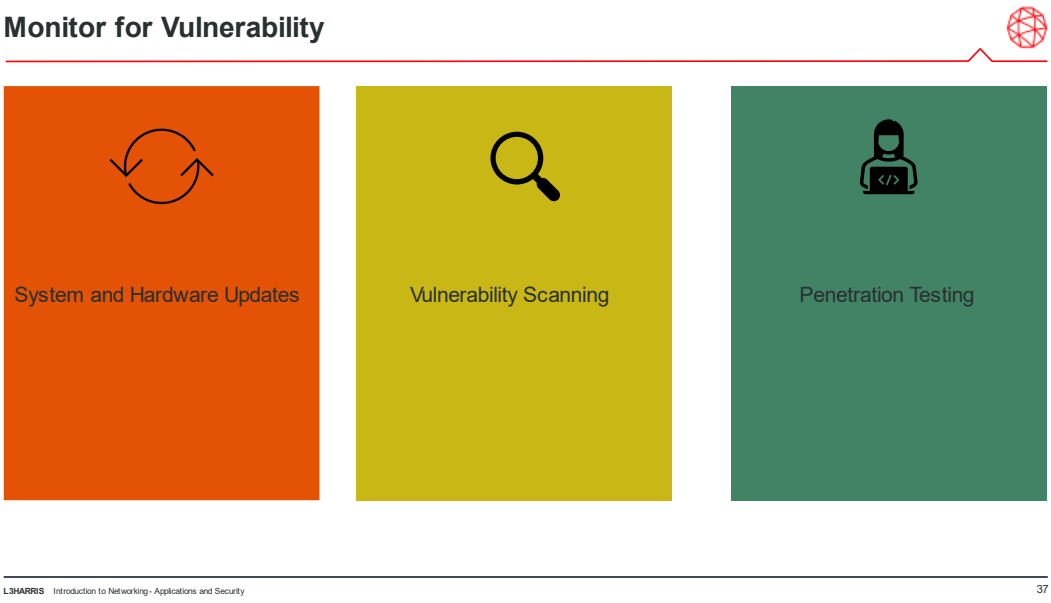

System Updates:

- Many updates address security vulnerabilities
- Patch Management: Procedures put in place to manage the installation of updates for hardware (firmware) and software.
- Make a configuration backup, in case settings must be reapplied after the update

Hardware Updates

- Driver—This is software that provides an interface between the operating system and the device.
- Firmware—This is software instructions stored in flash memory. This type of chip does not require a power supply, so the data does not have to be moved in and out of disk storage

Vulnerability Scanning is the process of auditing a network (or application) for known vulnerabilities. (I.e., unpatched software application, an administrator account with a weak password). Typical results from a vulnerability scan will identify common misconfigurations, the lack of necessary security controls, and missing patches

Penetration Testing: A comprehensive assessment of a network's security systems than a basic vulnerability scan. A penetration test (pen test), also referred to as ethical hacking, essentially involves thinking like an attacker and trying to demonstrate that the security systems can be compromised.