Firewalls and IDS

Dr. Natarajan Meghanathan
Professor of Computer Science
Jackson State University
E-mail: natarajan.meghanathan@jsums.edu

Firewalls

- A firewall is a software running on a dedicated host computer on which no other application is run.
  - To prevent someone from changing the firewall rules by exploiting the vulnerabilities of the other applications that may be run on the host.
- A firewall is as good as it is configured with, depending on the needs of the admin.
- Allowable or non-allowable traffic are typically identified with source/destination IP/network addresses and ports.
- **Filtering:** Egress filtering (filter outgoing traffic); Ingress filtering (filter incoming traffic)
- **Typical Firewall designs:**
  - **Default-deny approach (white-list):** Have a list of allowable traffic and block the rest.
  - **Default-allow approach (black-list):** Have a list of non-allowable traffic and allow the rest.
  - A good design needs to have a hybrid of these two approaches

Packet Filters

- A packet filter firewall is a stateless firewall that looks at only the packet headers to decide whether or not to drop a packet.
  - Stateless: Does not keep track of the decisions taken on any packet.
    - The decision taken on a packet is independent of the decision taken on the preceding packets.

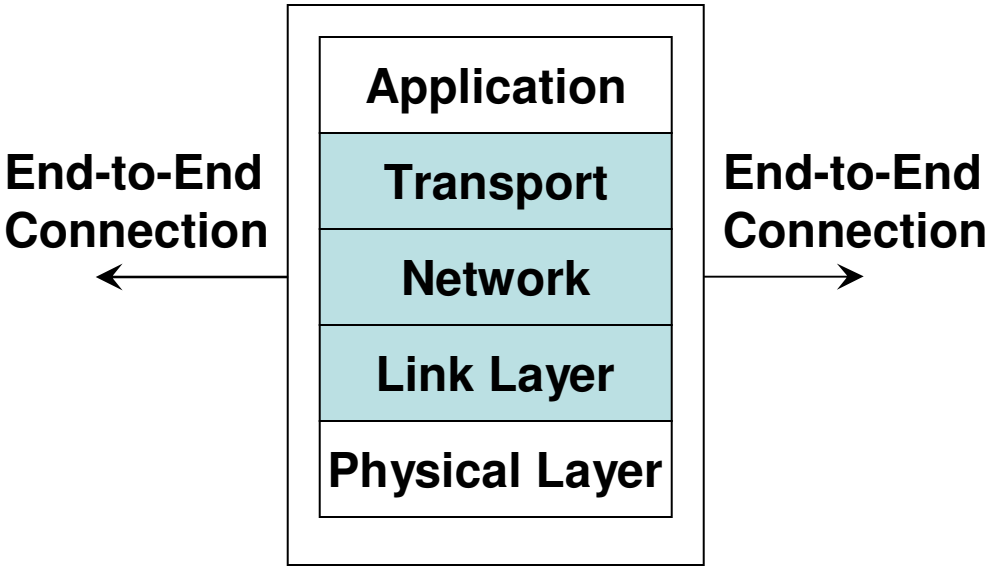

Layers supported by Packet Filter and Stateful Firewalls

The code for packet filters will become lengthy as we want to block traffic belonging to specific networks, IP addresses and transport layer protocols.

Need efficient filtering algorithms

Packet Filters

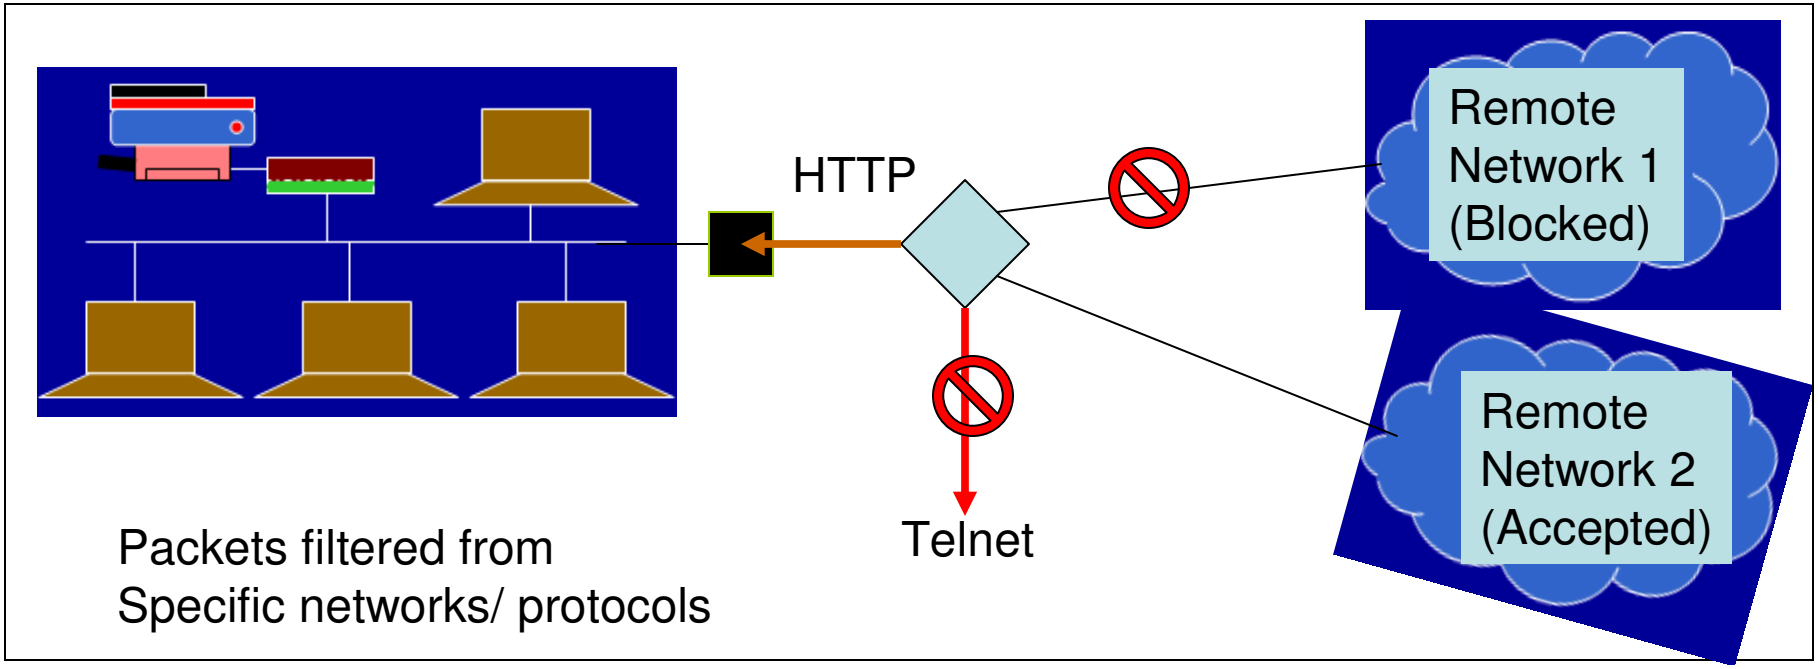

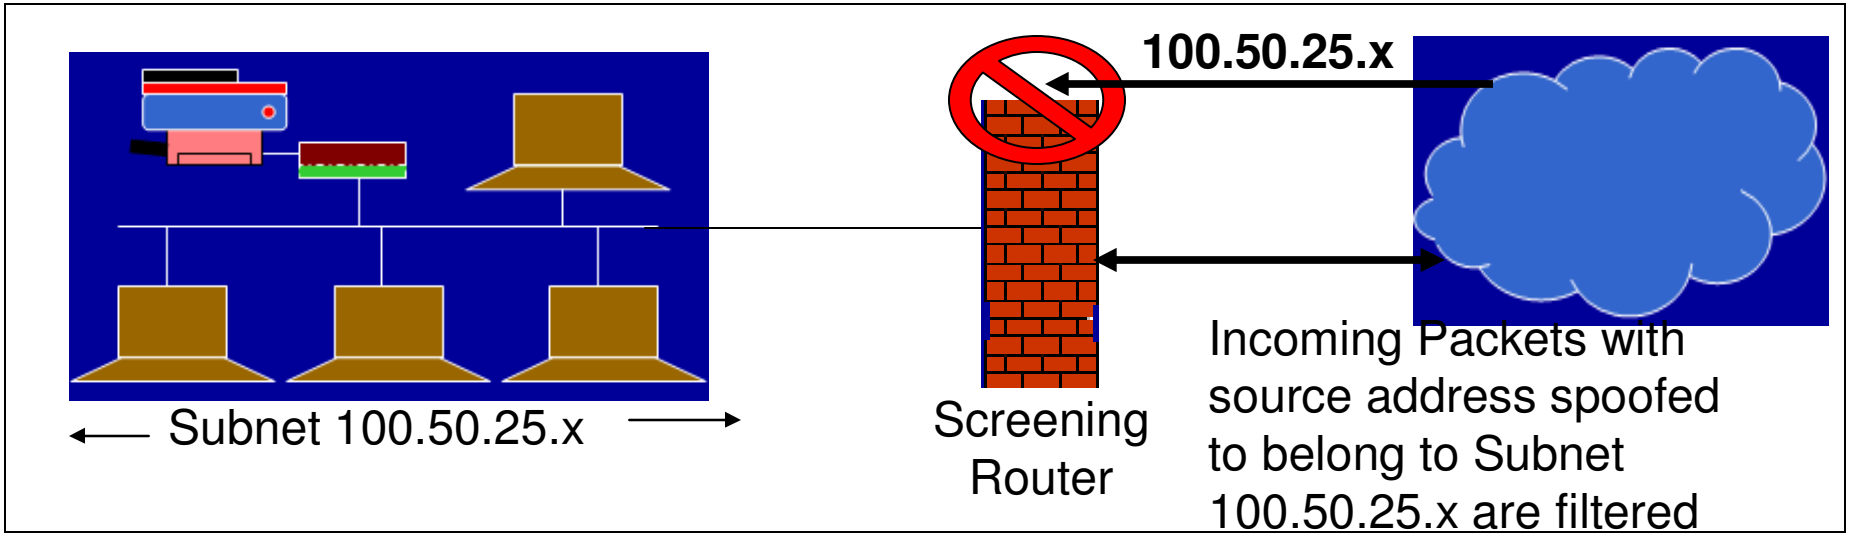

Attacks Detected by Packet Filters

- **IP Spoofing Attacks:** Have the packet filter configured not to let in packets having a source address that corresponds to the internal network.
  - For example, the attacker has spoofed the source IP address to be the IP address of a machine belonging to the network being protected by the firewall.
- **Source routing attacks:** where source specifies the route that a packet should take to bypass security measures, should discard all source routed packets
- **Tiny fragment attacks:** intruder uses the IP fragmentation option to create extremely small fragments and force the TCP header information into fewer separate fragments to circumvent filtering rules needing full header info; can enforce minimum fragment size to include full header.

Stateful Inspection Firewalls

- Stateful firewalls (also called circuit firewalls) examine the contents of each packet with regards to their placement within the packet series belonging to a specific session/connection.
- Stateful firewalls maintain records of all connections passing through the firewall and is able to determine whether a packet is the start of a new connection or part of an existing connection.

- **Attacks and actions prevented with Stateful Inspection Firewalls**
- **Session Hijacking Attack**: Stateful firewalls can remember the sequence numbers expected on both sides as part of a TCP session and can block attempts to hijack the session, when an intruder sends several TCP segments with different sequence numbers (trial-and-error).
- **SYN Flood Attack**: Stateful firewalls can remember the number of connection requests that have been let through for an IP address/TCP port and block connection requests beyond a threshold.
  - Also, do not let more than a certain number of simultaneous TCP connections to originate per (source) IP address.
- **Bandwidth Exhaustion**: Do not let more than a specific amount of data to be transferred per day from the inside network to any outside IP address.

Application Proxy Firewall

- Packet filters and stateful inspection firewalls look only at the headers of the packets, not at the data inside the packets.
- An application layer firewall (proxy; also called as bastion) simulates the proper effects of a packet on a receiving application so that the application receives only requests to act properly.
- A proxy firewall/gateway is a two-headed device: It looks to the inside as if it is the outside (destination) connection; while to the outside, it responds as it is from the inside.

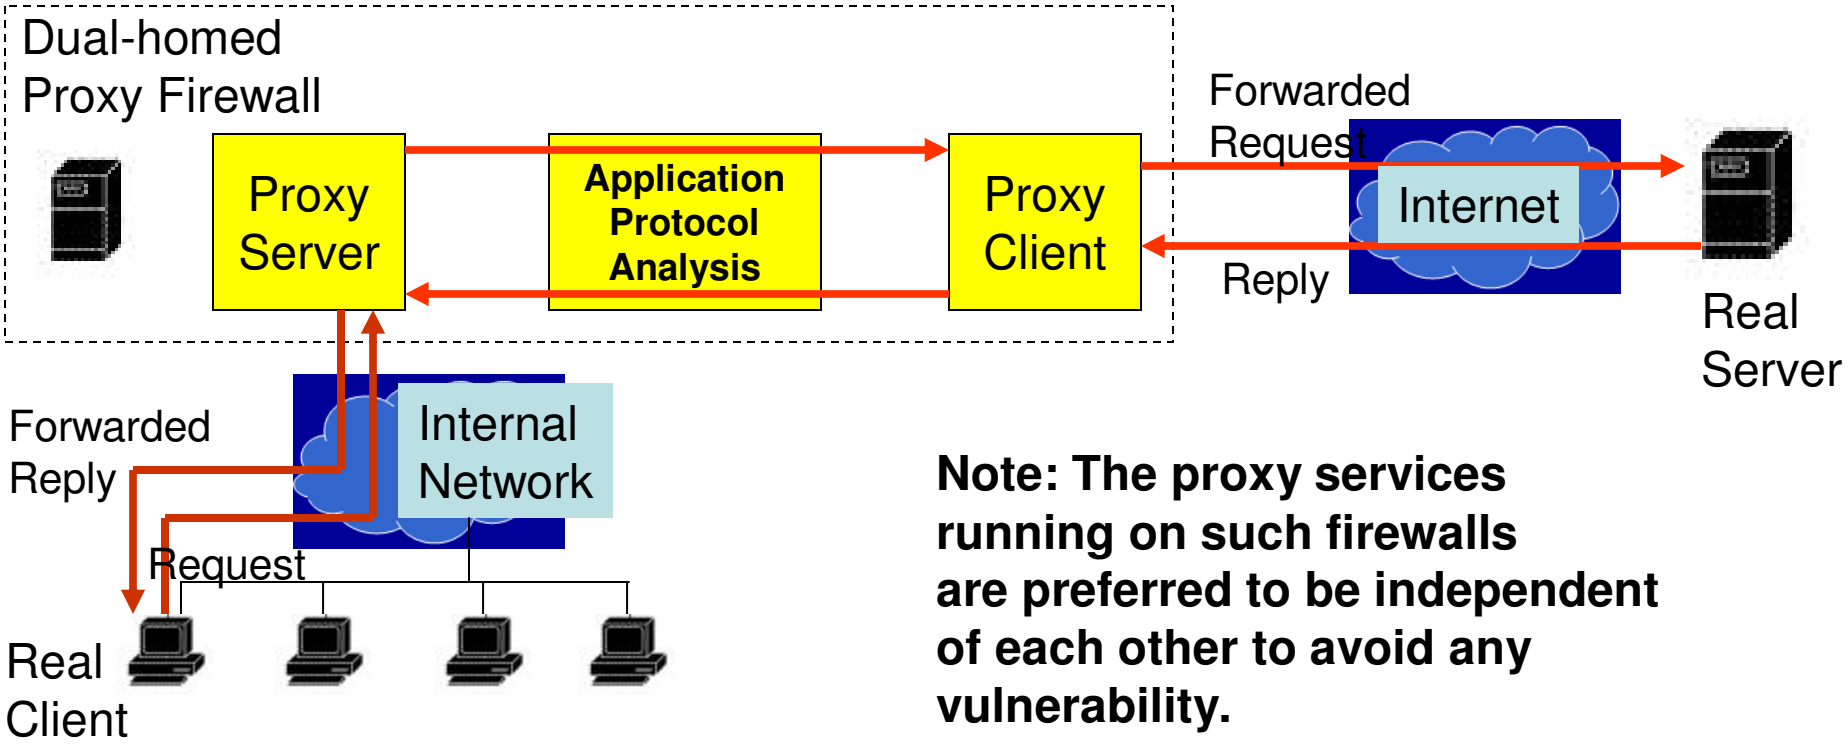

Note: The proxy services running on such firewalls are preferred to be independent of each other to avoid any vulnerability.

Application Proxy Firewall

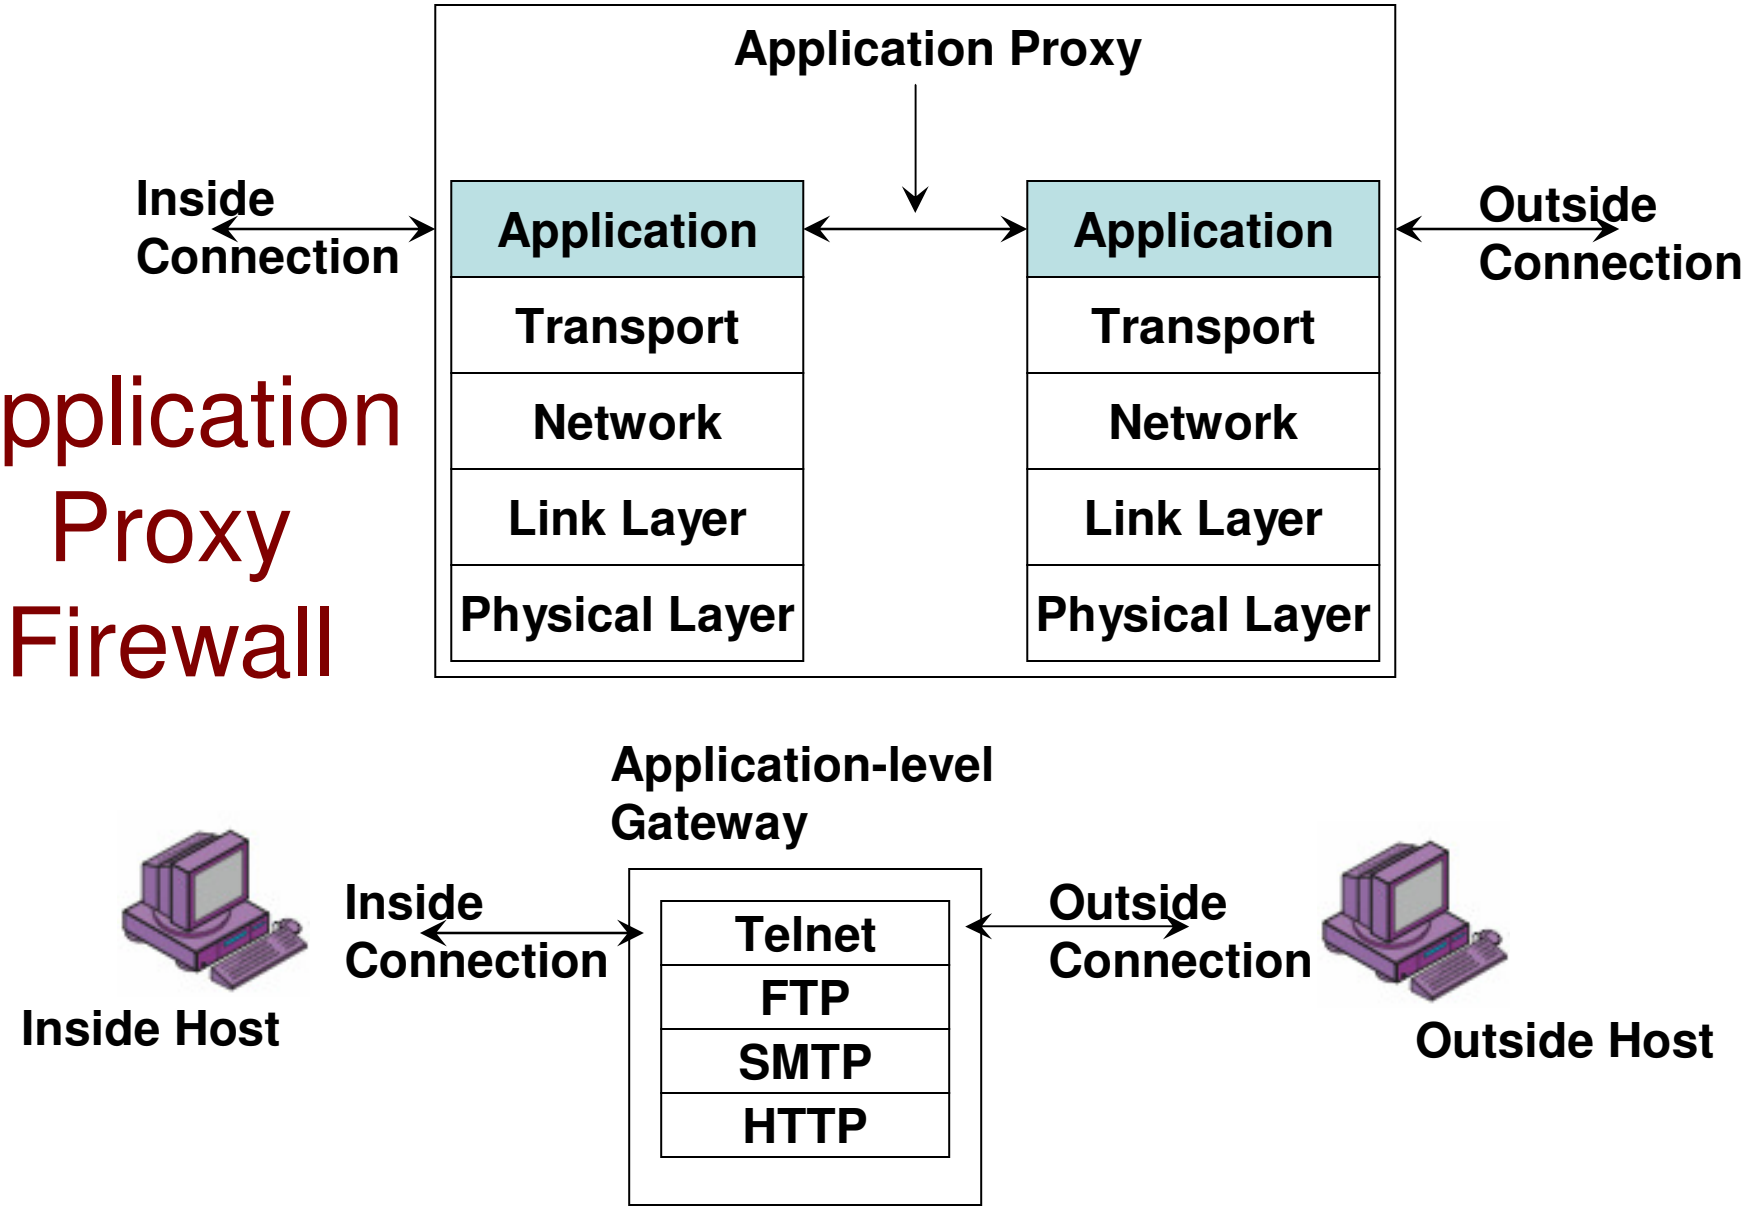

Source (adapted from): Figure 22.1(d) from William Stallings – Cryptography and Network Security, 5th Edition

Application Proxy Firewall

- Each application proxy in the firewall requires two components: a proxy server and a proxy client.
- All communication between internal users and the Internet passes through the proxy server rather than allowing users to directly communicate with servers on the Internet.
- An internal user (client) sends a request to connect to an external service. The request goes through the Application Proxy Firewall that runs a proxy server for that particular service being requested.
- The proxy server evaluates the request and decides to permit or deny the request based on a set of rules that are managed for the individual network service.
- Proxy servers allow only those packets that comply with the services of the application protocol.
- Proxy servers are also useful to collect audit records of session information
- If the proxy server approves the request, it forwards that request to the proxy client.
- The proxy client then contacts the real server on behalf of the real client and proceeds to relay requests from the proxy server to the real server and to relay responses from the real server to the proxy server.
- The proxy server relays requests and responses between the proxy client and the real client.

- Note: The above discussion assumes the client is in the internal network and the server is in the external network. The same discussion applies for the other scenario too:
  - The real client (from the outside network) contacts the proxy server, the proxy server evaluates the request and forwards to the proxy client, the proxy client contacts the real server (running in the internal network).
  - The proxy client forwards the response from the real server to the proxy server, which forwards the response to the real client (in the outside network).

Reverse Proxy Firewall

Examples of using Proxy/Reverse Proxy Firewall

- Scenario 1: A company wants to allow dial-in access by its employees, without exposing its company resources to login attacks from remote non-employees. Suppose the internal network has a mixture of operating system types, none of which support strong authentication through a challenge-response system.
- Solution:
  - The requirement could be handled by a specifically written proxy firewall that requires strong authentication such as a challenge-response, in addition to a valid username and corresponding password.
  - The proxy validates the challenge-response itself, and then pass on only the username and password in a form required by the internal host's operating system.
- Scenario 2: A company wants to set up an online price list so that outsiders can see the products and prices offered. It wants to be sure that (a) no outsider can change the prices or product list and (b) outsiders can access only the price list and not any of the more sensitive files stored inside.
- Solution:
  - The requirement could be handled by a specifically written reverse proxy firewall that monitors the file transfer protocol data to ensure that only the price list file was accessed, and that the file could be only read, not modified.

- Note: A proxy firewall or reverse proxy firewall can also function more as a guard, monitoring the amount and quality of data exchanged.
  – It could keep track of the amount of data exchanged per user from the internal network and deny access if exceeded a pre-defined limit. (proxy firewall)
  – It could also run a virus scanner to scan all the incoming files to a file server and if required outgoing files too. (reverse proxy firewall)

Personal Firewalls

- Motivation: Home users, individual workers, and small businesses use cable modems or DSL connections with unlimited, always-on access.
- These people need a firewall, but a separate firewall computer to protect a single workstation can seem too complex and expensive.

- A workstation could be vulnerable to malicious code or malicious active agents (ActiveX controls or Java applets), leakage of personal data stored in the workstation, and vulnerability scans (like nmap) to identify potential weaknesses.
- A personal firewall is an application program that runs on a workstation to screen traffic on the workstation and block unwanted traffic leaving or entering the workstation to the network to which it is connected.
- A user could configure the personal firewall to accept traffic only from certain sites, and not from specific sites, and to generate logs of activities happened in the past
- A personal firewall could be also configured with a virus scanner which would be then automatically invoked to scan any incoming data to the workstation.

- A static machine is a vulnerable target for the attack community and adding a personal firewall can save it more secure compared to machines that are not behind such a firewall.

What Firewalls Can and Cannot Block

- Firewalls cannot alone secure an environment.
- A firewall protects only the perimeter of its environment against attacks from outsiders who want to execute code or access data on the machines in the protected environment.
- Firewalls cannot protect from internal threats (through disgruntled employees).
- Firewalls cannot protect against malware imported via laptop, PDA, or portable storage device infected outside the network, then attached and used internally.
- Firewalls can be held responsible for any security breach in if they are the only means to control the entire network perimeter.
  - If a host in the inside network has a connection to the outside network through a modem, the whole of the inside network is exposed to the outside network through the modem and the host. A firewall cannot be responsible for any attack
- Firewalls cannot protect data after they have left them.
- A firewall is often a single point of failure for a network.
  - A more layered approach like a screening router, followed by a proxy firewall, followed by a personal firewall may be more helpful.
- Firewalls must be frequently configured and updated to take into account the changes in the internal and external environment and based on the review of the firewall activity reports that may indicate intrusion attempts.
- The machine hosting the firewall code will not have any other software like an editor, compiler, etc. in order to reduce the chances of an attack.

Demilitarized Zone (DMZ) Networks

- A DMZ network (also called perimeter network) is a subnet that contains an organization's services that are exposed to a larger untrusted network (like the Internet).
- In other words, the DMZ comprises of hosts that provide services to users outside the internal LANs, such as e-mail, web, DNS servers.
- Because of the higher chances of these hosts being compromised, they are placed into their own sub-network in order to protect the rest of the network if an intruder were to succeed in attacking them.
- Thus, a DMZ network adds an additional layer of security to an organization's LAN – an external attacker only has access to the hosts in the DMZ and not to any other internal networks.
- Hosts in the DMZ provide services to both the internal and external networks – an external ("front-end") firewall monitors the traffic between the DMZ network and the external Internet; while, an internal ("back-end") firewall monitors the traffic between the DMZ hosts and the internal network clients.

DMZ Networks

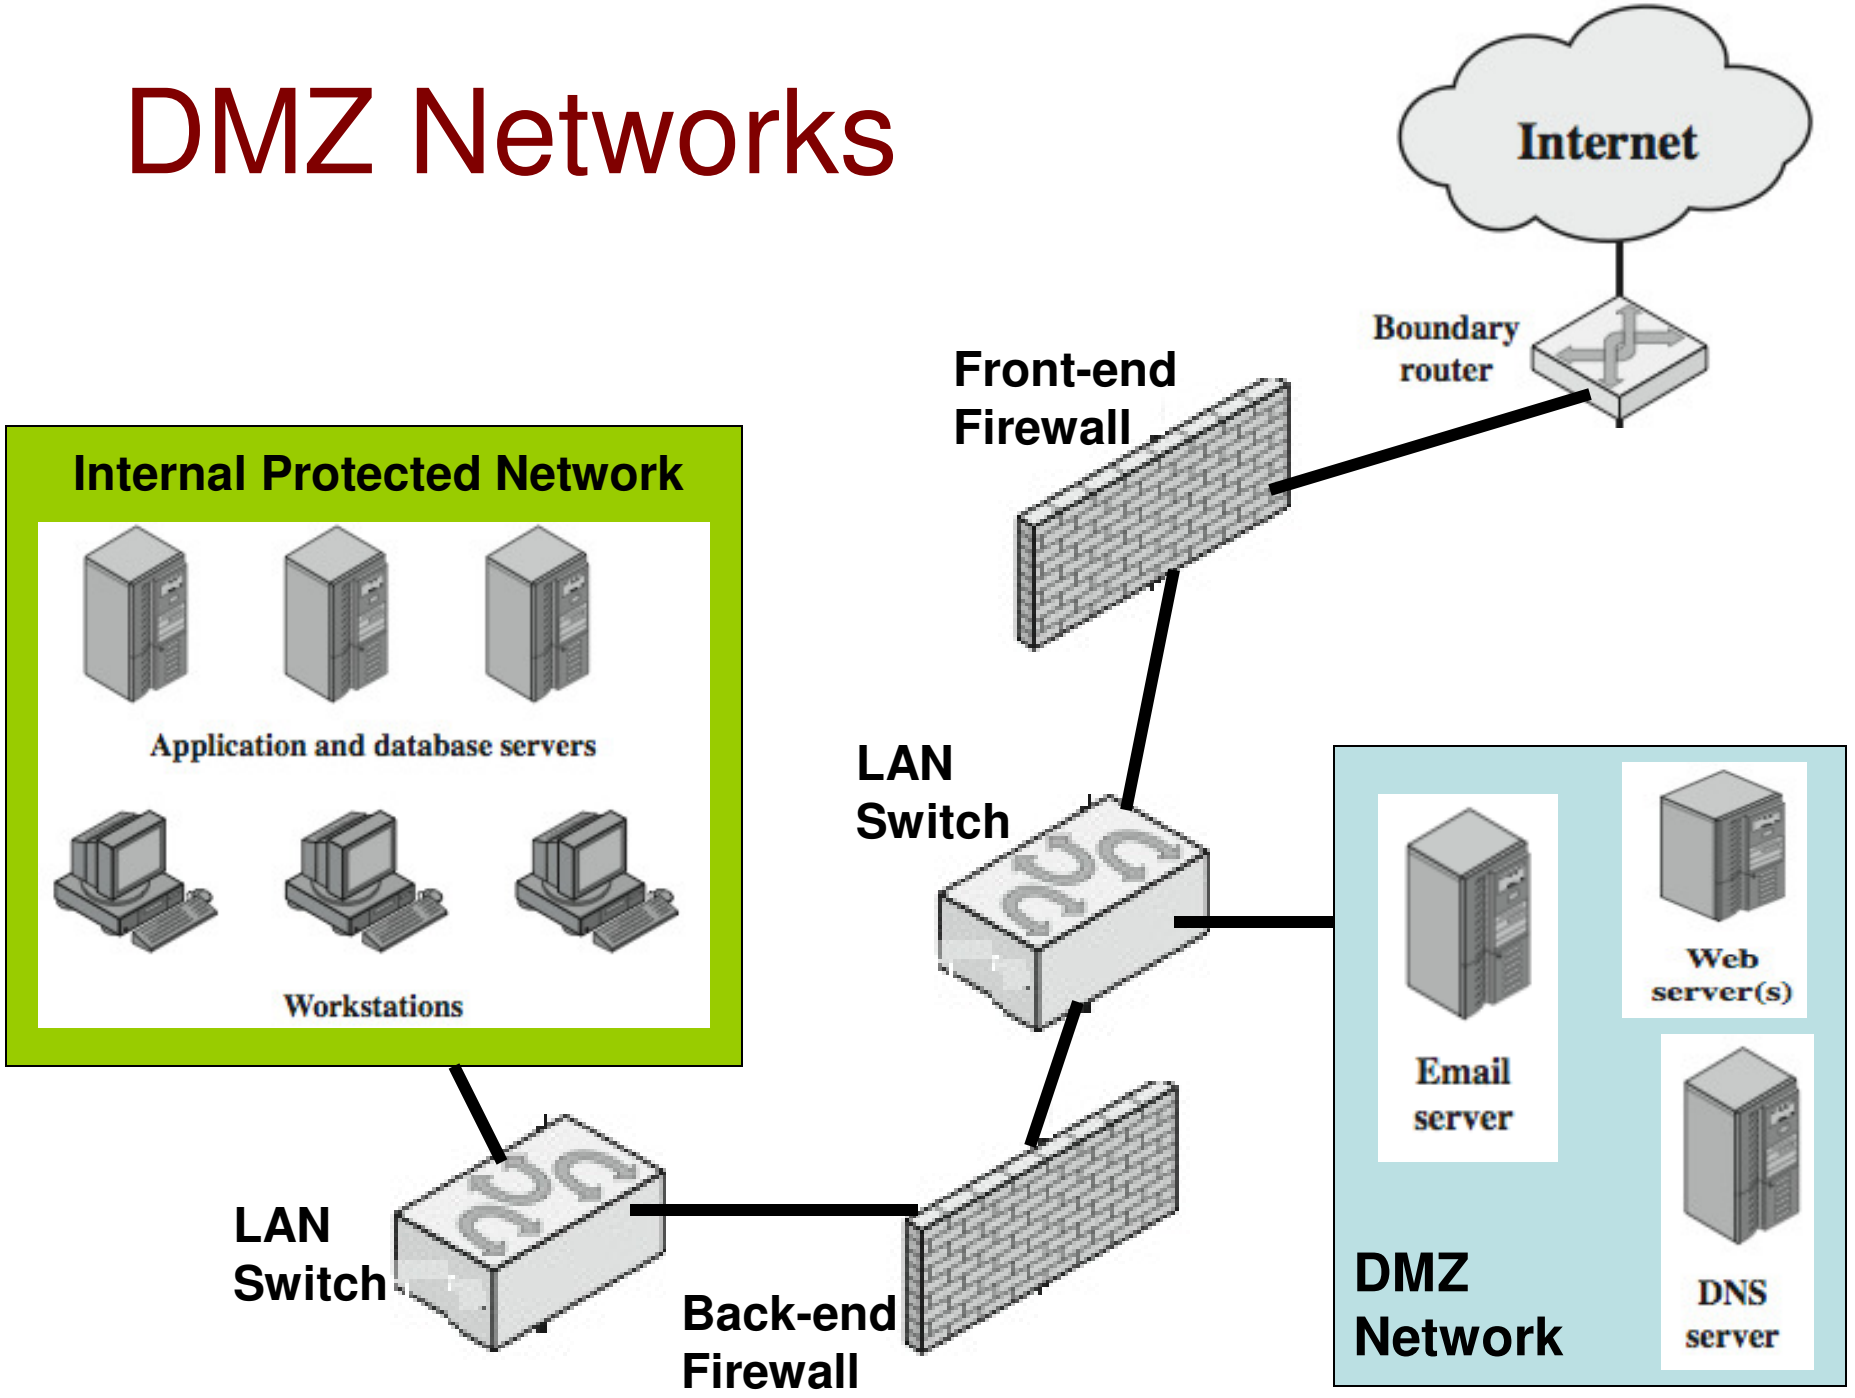

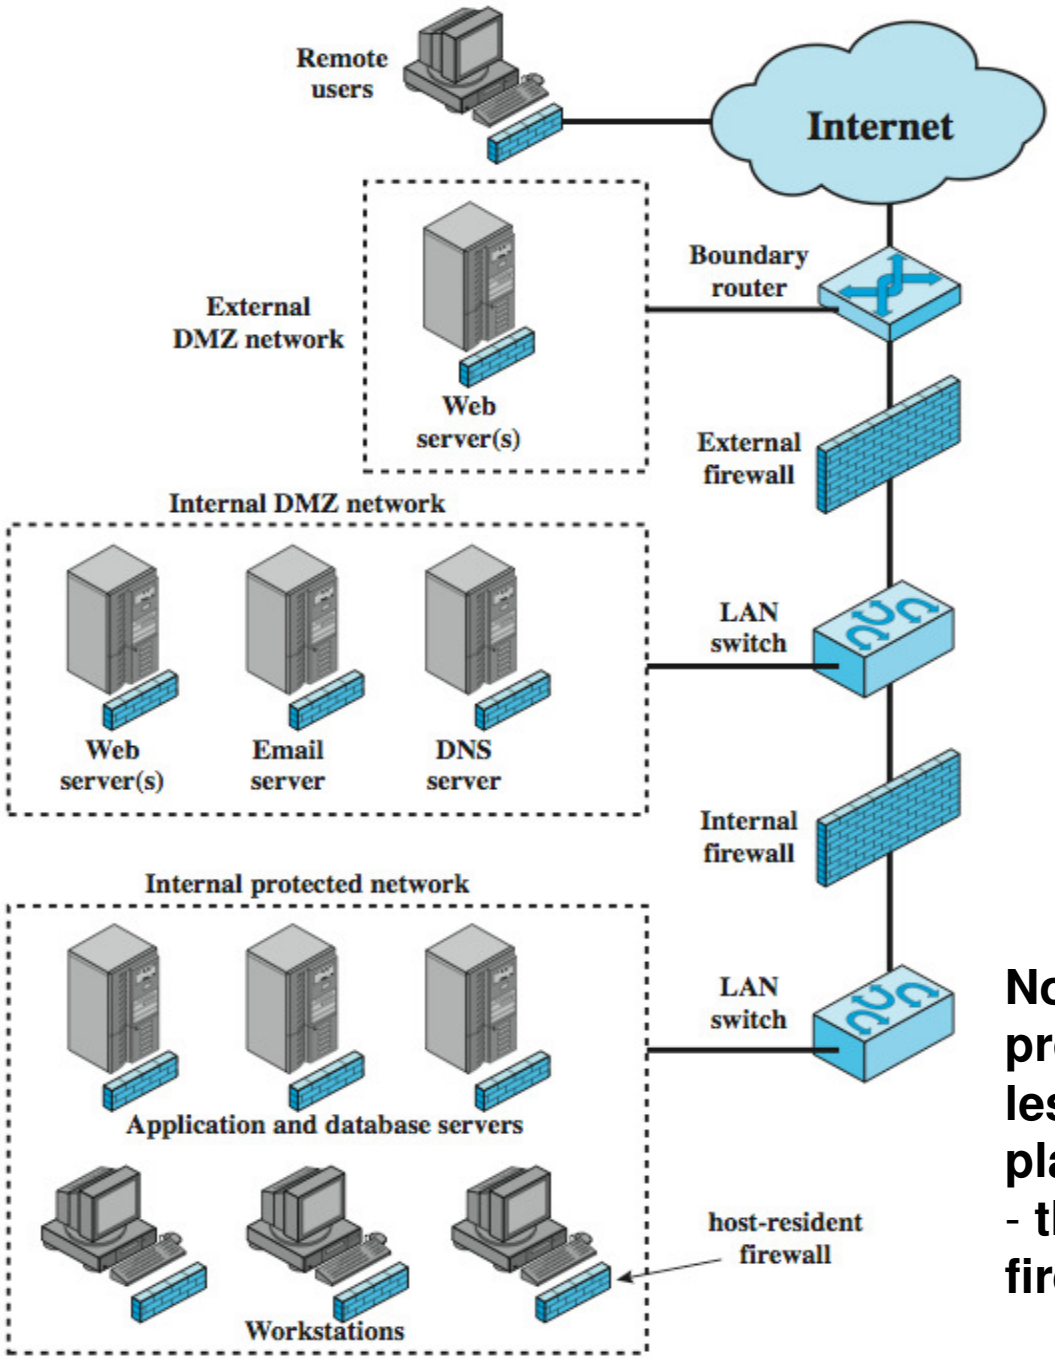

Distributed Firewalls

Note: Servers that need less protection, because they may have less critical information, could be placed in the external DMZ network - that is even outside the external firewall.

Intrusion Detection Systems (IDS)

- An IDS to the networking world is like a burglar alarm to the physical world.
- The main purpose of an IDS is to identify suspicious or malicious activity, note activities that deviate from normal behavior, catalog and classify the activity, and, if possible, respond to the activity.
- Host-based IDS (HIDS): It examines activities on an individual system and not concerned with other systems or the network.
- Network-based IDS (NIDS): It examines activity (traffic) crossing the network it is monitoring and not concerned about individual systems.

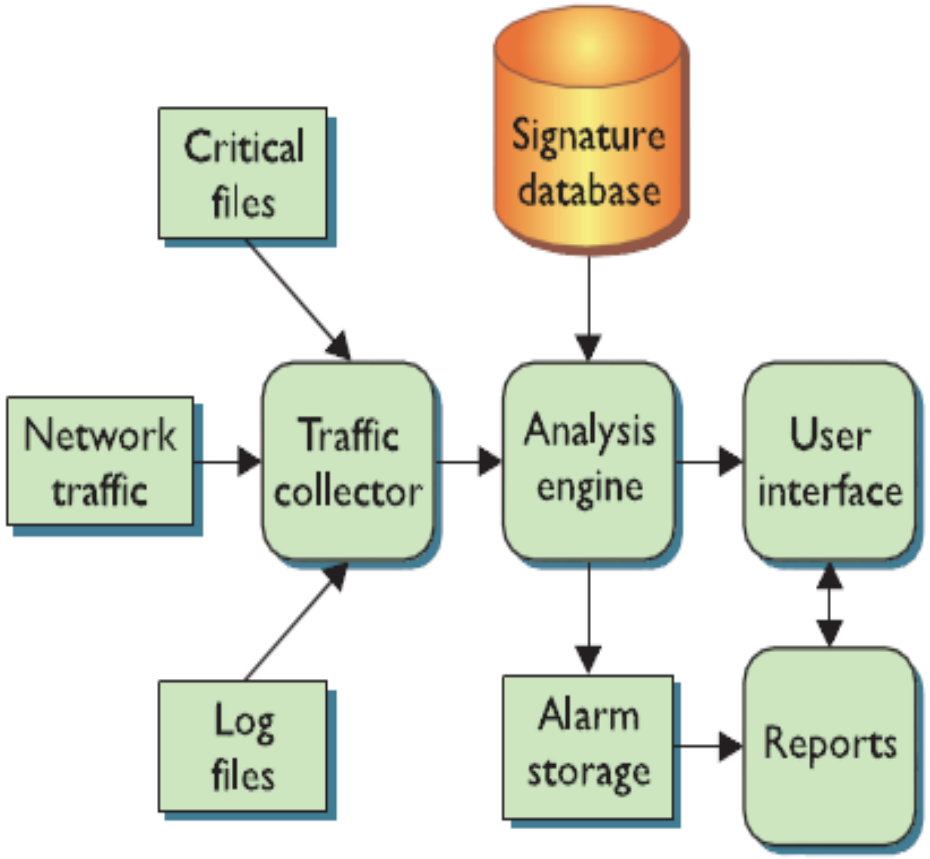

Logical Depiction of IDS Components

Source: Figure 13.2 from Conklin and White – Principles of Computer Security, 2nd Edition

Logical Components of an IDS

- An intrusion detection system (IDS) is a device, typically a separate computer, that monitors activity to identify malicious or suspicious events.
- An IDS typically consists of several special components (often logical and software-based rather than physical) working together on the device in which it is installed.
- Traffic Collector – Collects activity/events for the IDS to examine. For a HIDS, these could be log files, audit logs, or traffic coming to or leaving a specific system. For a NIDS, these could be network traffic captured through a sniffer.
- Analysis Engine – Examines the collected network traffic and compares it to known patterns of suspicious or malicious activity stored in the signature database. It is often referred to as “brain” of the IDS.
- Signature Database: A collection of patterns and definitions of known suspicious or malicious activity.
- User Interface and Reporting: Interfaces with the human element, providing alerts when appropriate and giving the user a means to interact with and operate the IDS.

False Positives and False Negatives

- When an IDS matches an activity to a specific pattern and generates an alarm for a non-malicious traffic that is not a threat, it is called a false positive.
- Technically, the IDS is functioning correctly by matching the pattern and has no ability to determine the intent behind the activity; but, from a human standpoint, this is not an information the analyst needed to see, as it does not constitute a threat and does not require intervention.
- Hostile activity that does not match an IDS signature and goes undetected is called a false negative.
- Note that an IDS is limited by its signature set – it can match only activity for which it has stored patterns.

Signature and Anomaly-based IDS

- Based on the approach adopted to detect suspicious or malicious traffic, IDS could be categorized into Signature-based and Anomaly-based IDS.
- *Signature-based IDS*: Relies heavily on a pre-defined set of attack and traffic patterns called signatures.
- A signature-based IDS (like an anti-virus software) can only match against known patterns – if a new attack comes in that the signature-based IDS has never seen before, it would not be able to identify it as suspicious or malicious – a primary weakness of signature-based IDS.
- *Anomaly-based IDS*: Monitors activities and attempts to classify them as either "normal" or "anomalous" (suspicious and unknown) based on self-created rule sets.
- An anomaly-based IDS uses heuristic techniques to categorize and classify traffic while developing and refining their internal rule sets.
- An advantage with anomaly-based IDS is that it can potentially detect new attacks or variant of old attacks.
- A drawback of anomaly-based IDS is that it could generate a potentially high number of false positives while the system is learning what "normal" is. Hence, such IDS should be programmed to dynamically adapt to changes.

Signature-based IDS

- Example for detecting an attack using signature-based IDS:
  - Detecting the TCP SYN flood attack using a port scan
    - An IDS would probably find nothing unusual in the first SYN (say to port 80) and then another SYN (from the same source address) to port 25.
    - But as more and more ports (especially closed ports) receive SYN packets, the pattern will reflect a possible port scan that happened already

- A problem with the signature-based IDS is the signature itself: An attacker will try to modify a basic attack in such a way that it will not match the known pattern of the attack.
  - For example, an attacker may convert lowercase to uppercase characters, characters by the ASCII equivalents and etc.
- An IDS has to learn more signature patterns to catch an attack with different patterns.
- Statistical analyses are nowadays used to detect attacks with patterns that match with the stored signatures within a certain probability of error.

Anomaly/ Heuristic-based IDS

- Like the signature-based IDS, heuristic-based IDS is limited by the amount of information the system has seen (to classify actions into the right category) and how well the current actions fit into one of the categories.
- Activities could be classified into three categories: Good/ Benign, Suspicious and Unknown.
  - Over time, specific kinds of actions can move from one category to another depending on whether the IDS learnt in due course that certain actions are acceptable or not
- Model-based IDS: Develop standard models for certain activities. If the activities violated the model, raise an alarm.
  - Example: A normal behavior of an employee to start his day in the work environment would be to read emails, write many documents using a word processor, and occasionally backup files, etc. If an employee accesses system sensitive management utilities immediately after login, it could raise an alarm.
- State-based IDS: Monitor the system as it goes through different state changes. If the rate of state change is faster than a threshold or the system has entered into previously unseen state, or the system has veered into an unsafe mode, it could raise an alarm.
- Misuse-based IDS: Identify activities that could be easily misused.
  - Example: An attempt to access a password file is suspect, except for few utilities like login, password change, create user.

Network-based IDS (NIDS)

- NIDS are placed next to the firewall on the network perimeter and analyze the traffic as it passes by for the protocols, source, destination, content, traffic already seen and etc.
- A NIDS typically looks for traffic that typify hostile actions or misuse, such as the following:
  - Denial-of-service attacks, Port scans or sweeps, Malicious content in the data payload of a packet or packets, Vulnerability scanning, Trojans, Viruses, Worms, Tunneling and Brute-force attacks.
- The traffic collector of a NIDS logically attaches itself to a Network Interface Card (NIC) that operates in promiscuous mode (stealth mode) and sniffs the passing traffic.

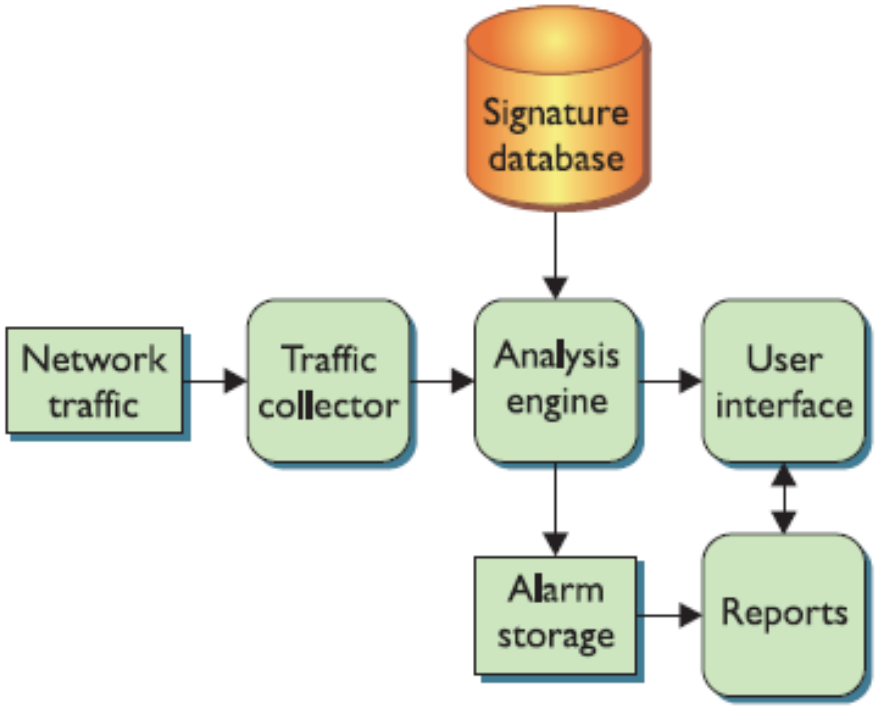

Source: Figure 13.4 from Conklin and White – Principles of Computer Security, 2nd Edition

NIDS Placed behind Firewall

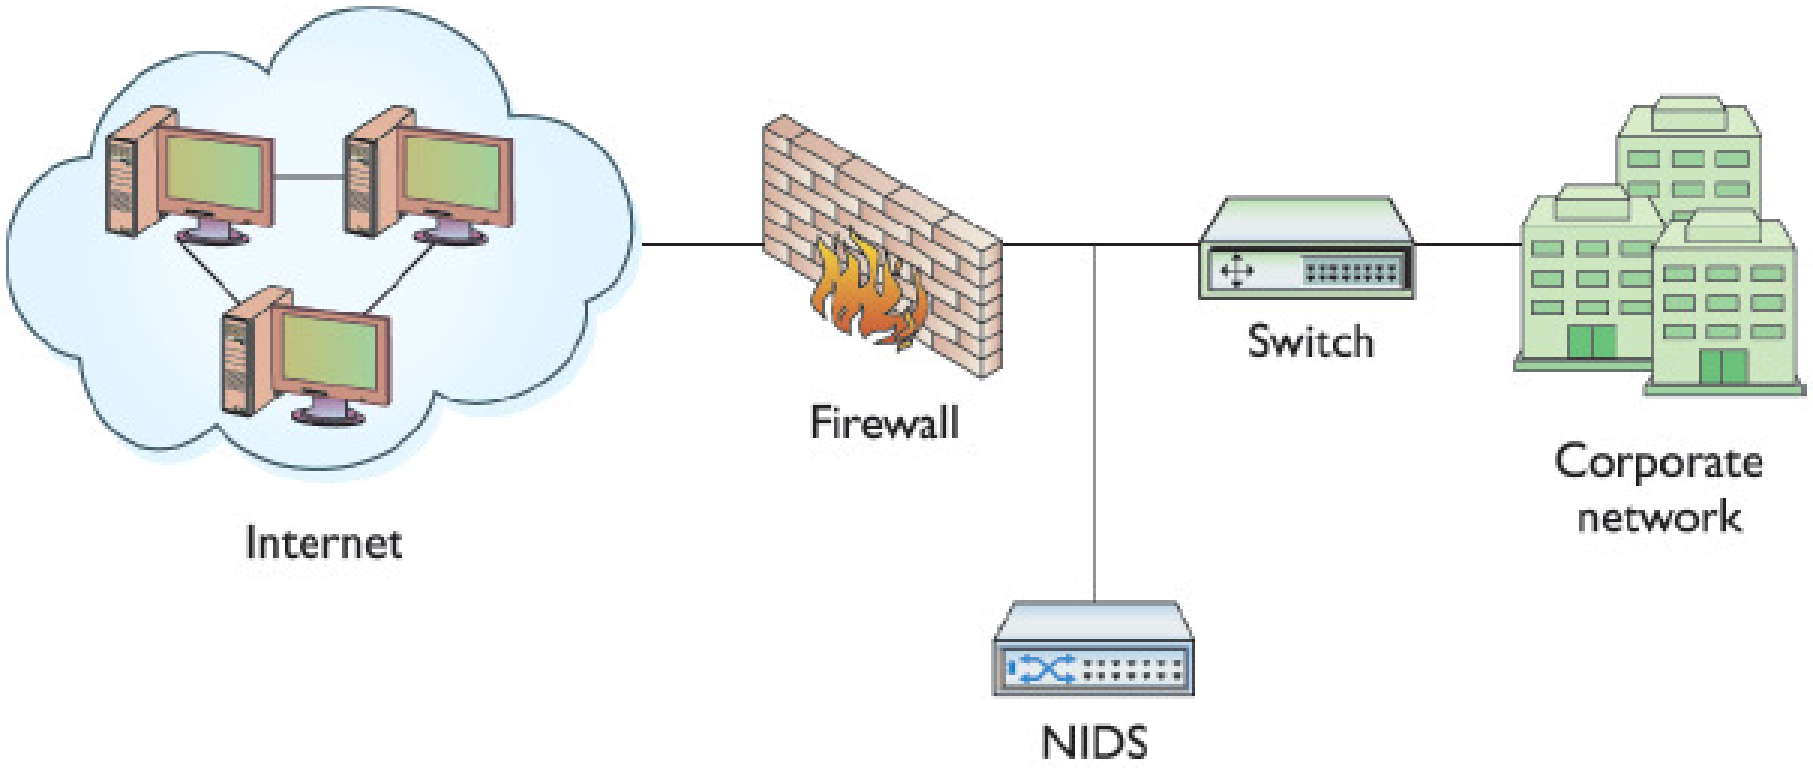

Source: Figure 13.7 from Conklin and White – Principles of Computer Security, 2nd Edition

NIDS: Advantages and Disadvantages

- Advantages of a NIDS
- Less Overhead: With a few well-placed NIDSs, one can monitor the entire network traffic going in and out of the organization. Also, upgrading and maintaining a fewer number of NIDSs is usually much cheaper than upgrading and maintaining hundreds of host-based IDSs.
- Big Picture: The collection of the few NIDSs can have visibility into all the network traffic and can correlate attacks (whether they are widespread or concentrated, unorganized or focused) among multiple systems.
- Disadvantages of a NIDS
- A NIDS is ineffective when traffic is encrypted.
- A NIDS cannot see traffic that does not cross it – If a NIDS is placed only in the perimeter, chances are that it could miss traffic traversing the internal network.
- A NIDS must be able to handle high volumes of traffic (even 1-Gbps is common nowadays) with the availability of networks with larger bandwidth.
- A NIDS does not know about activities on the hosts themselves.

Active vs. Passive NIDS

- Passive NIDS:
- A passive NIDS simply watches the traffic, analyzes it and generates alarms.
- It does not interact with the traffic itself in any way, and it does not modify the defensive posture of the system to react to the traffic.
- Active NIDS:
- An active NIDS contains all the same components and capabilities of the passive NIDS with one critical addition – the active NIDS can react to the traffic it is analyzing.
- The reactions of an active NIDS could range from something simple, such as sending a TCP reset message to interrupt a potential attack and disconnect a session, to something complex, such as dynamically modifying firewall rules to reject all traffic from specific source IP addresses for the next few hours or days.
- Active NIDS are also referred to as Intrusion Prevention Systems (IPSs). When configured with the private keys of the servers in the internal network, IPSs would be able to decrypt the SSH connection establishment messages between a client and server and extract the session keys that would be used during the complete session. This gives an added advantage for the IDS/IPS to handle encrypted traffic.

Host-based IDS (HIDS)

- A host-based IDS (HIDS) examines log files, audit trails (both generated by the local operating system), and network traffic coming into or leaving a specific host.
  - On UNIX systems, the examined logs are those created by syslog, kernel logs and error logs; On Windows systems, the examined logs are the event logs – Application, System and Security.
- Critical files are those that are vital to the system's operation or overall functionality. They may be program (or binary) files, files containing user accounts and passwords, or even scripts to start or stop system processes.
- Any unexpected modifications (for e.g., could be detected using checksum) to the critical files could mean the system has been compromised or modified by an attacker. By monitoring these critical files, the HIDS can warn users of potentially malicious activity.
- Within the log files, the HIDS is looking for certain activities that typify hostile actions or misuse, such as the following:
  - Logins at odd hours, Login authentication failures, Additions of new user accounts, Modification or access of critical system files, Modification or removal of binary files (executables), Privilege escalation

Packet Sniffer (Protocol Analyzer)

- Packet Sniffer: Is a computer software (or even a computer hardware programmed to) intercept and log traffic passing over a LAN.
- A packet sniffer may be used for both beneficial and malicious purposes:
  - Analyze network problems and monitor network usage
  - Gather and report network statistics
  - Filter suspect content from network traffic
  - Spy on other network users and collect sensitive information such as passwords
  - Reverse engineer (study using the structure of the different packet headers) the protocols used over the network
  - Detect network intrusion attempts
  - Gather information for effecting a network intrusion
- In order to capture all the network traffic, the Network Interface Card (NIC) on the IDS hosting the packet sniffer should run in promiscuous mode and analyze every packet crossing the wire.
- Most switches come with SPAN (Switched Port Analyzer) port – a mirrored port that will see all the traffic passing through the switch or through specific virtual LANs. Packet sniffers can be run on the SPAN port of a switch.

Honeypot

- A honeypot is a trap to detect, deflect or in some manner counteract attempts at unauthorized use of information systems. –
- A honeypot is usually a computer, and sometimes data or an unused IP address space that appears to be part of a network but which is actually isolated, unprotected and monitored, and which seems to contain information or a resource that would be of value to attackers.
- Honeypots have no production value and hence should not see any legitimate traffic or activity. Whatever they capture can be surmised as malicious or unauthorized.
- A honeynet is a network of honeypots. A honeynet is used for monitoring a larger and/or more diverse network in which one honeypot may not be sufficient.
- A honeypot/ honeynet is more of a preventative approach of detecting potential attackers existing in the Internet who may target the organization network in the near future.
- Honeypots could be used to fake as open relays to attract spam emails and determine the source e-mail address and destination e-mail addresses used by the spammers.
  - An open relay is an e-mail server that allows anyone on the Internet to send email through it.
  - Once they find an open relay, spammers keep sending the span email to the open relay and expect it to spread the spam.
- Note that no ordinary e-mail will come to a honeypot. All it receives could be categorized as spam.