

Securing Networks

AP Cybersecurity

Network Vulnerabilities and Attacks

A **network** connects devices so they can share data - and every connection is a possible way in. You must know the classic network attacks and the tricks behind them.

- **ARP poisoning** - the **address resolution protocol (ARP)** pairs IP addresses with hardware **MAC addresses**. An adversary sends fake ARP messages so traffic meant for the target flows to the adversary instead. This is an **on-path attack** (also called man-in-the-middle): the adversary secretly sits between two parties, reading and even altering their messages.
- **MAC flooding** - flooding a **switch** with fake MAC addresses forces it into broadcast mode, so the adversary can capture all traffic. This is a form of **eavesdropping**.
- **DNS poisoning** - planting a fake record on a **domain name system (DNS)** server redirects users to a malicious site to steal credentials (**credential harvesting**).
- **Smurf attack** - flooding a network with **ICMP** requests aimed at the broadcast address, so every device replies to the victim. It is a **denial of service (DoS)** attack; when many machines attack at once it becomes a **distributed denial of service (DDoS)**.

Adversaries exploit weak networks to flood, map, or spoof devices. A physical data port with no **port security** lets an attacker plug in; an open port lets them install a **rogue access point** that bypasses the firewall entirely. We rate network risk by impact and by how much skill the exploit needs.

To find weaknesses **before** an adversary does, organisations run an **automated vulnerability scanner**: a tool that checks networks, devices, and applications against a database of **known** vulnerabilities, then produces a report listing each one found, how **severe** it is, and a recommended **mitigation**. Fixing the highest-severity items first is a core part of managing network risk.

Protecting Networks: Managerial Controls and Wireless Security


Good network security starts with written **policies** that set a minimum standard: a **router security policy** and **switch security policy** ban local accounts and require port security; a **VPN policy** sets authentication rules and forbids **split tunneling**; and a **wireless security policy** requires strong encryption and authenticated access.

For wireless networks specifically, organisations **disable beacon frames** so the network is harder to find, **control signal strength** so it does not leak outside the building, **enable strong encryption** - **WPA3** Wi-Fi is the current strongest, while old **WEP** and the original **WPA** are broken - and use **MAC filtering** to allow only known devices.

Protecting Networks: Segmentation

Network segmentation divides one network into smaller, isolated pieces (**subnets**). If one subnet is breached, the damage is **contained** and cannot spread.

A key pattern is the **screened subnet** (also called a **DMZ**). It sits between the public internet and the private internal network, holding an organisation's public-facing servers in a lower-security zone - separated from the sensitive internal systems.



AP Cybersecurity · U3 · screened subnet (DMZ) between two firewalls

A screened subnet (DMZ) puts public servers between two firewalls, away from the private network

Segments can also be built with **subnetting** (by IP address) or **VLANs** (logically separating devices on the same switch). Each segment can then get its own security policy - higher-security and lower-security zones.




Server racks: network segmentation isolates systems so one breach does not open everything

Protecting Networks: Firewalls

A **firewall** allows or denies traffic entering or leaving a network. There are several kinds:

- **Stateless** - filters on packet headers alone (IP, port, protocol).
- **Stateful** - also tracks the **state** of each connection for finer control.
- **Next-generation (NGFW)** - adds advanced features like intrusion prevention and deep packet inspection.

A firewall follows an **access control list (ACL)** - an ordered set of rules. Rules are checked **in order**, and the **first match wins**, so the order of rules changes which traffic gets through. Each rule specifies a direction, a thing to filter by (IP, port, service), and an action (permit or deny).




AP Cybersecurity · U3 · firewall ACL: first matching rule wins

A firewall checks its ACL top to bottom; the first matching rule decides

Worked example. A firewall has Rule 3: DENY TCP 443 from 192.168.*, and lower down Rule 7: ALLOW TCP 443 from ALL. A user at 192.168.45.37 cannot reach port 443 - even though Rule 7 would allow them - because Rule 3 matches first, and the *first match wins*. The fix is to move the ALLOW rule **above** the DENY. This is why rule order, not just rule content, decides what traffic gets through.

Firewalls belong at every point where data crosses between zones - at each network segment and at every gateway to the public internet.



Real network hardware: a firewall is a device (or software) sitting where these cables meet the outside world

Detecting Network Attacks

When prevention fails, **detection** takes over. Automated tools read the **log files** that record network activity:

- a **network intrusion detection system (NIDS)** analyses traffic and **raises an alert**, but does not block;
- a **network intrusion prevention system (NIPS)** can also **stop** an attack by closing ports or blocking addresses;
- a **security information and event management (SIEM)** system gathers data from many sources to spot patterns.

There are two detection methods. **Signature-based** detection compares traffic to a database of known attack **signatures** - fast and low on false alarms, but blind to brand-new attacks. **Anomaly-based** detection compares traffic to a normal **baseline** and flags anything unusual - it can catch novel attacks but needs more resources and raises more false alarms. A **hybrid** approach combines both.

Examining captured traffic (**packet-capture** files), analysts hunt for **network-based indicators of compromise** in the source and destination IP addresses, ports, and protocols. Four common ones: connections to **known-malicious IP addresses**, **unauthorized network scans** (an outsider probing your ports), unusual **spikes or slowdowns** in traffic, and **mismatched port-application traffic** (for example, non-web traffic flowing over port 80). These complete the host-, file-, and behaviour-based indicators a single device logs.

■ AI, thresholds, and alert fatigue

A medium network logs **millions** of events a day - far more than any team can read - so organisations train **AI** models to sort likely-malicious patterns from normal ones. These models are **probabilistic** : a detector can report a **percentage** score for how suspicious an event appears. The course describes this as a likelihood; a reported percentage should not automatically be treated as a perfectly calibrated probability.

The organisation then sets a **threshold** - the likelihood at which an alert fires - and that choice is a genuine trade-off:

- set the threshold **too high** and real attacks slip through **undetected**;
- set it **too low** and the team is **overwhelmed** with false alerts.

Try it – apply an explicit threshold. In this invented teaching dataset, higher scores mean more suspicious. Alert when score is greater than or equal to the threshold. Investigation has established the outcomes.


Event	Score	Outcome
A	90%	attack
B	70%	normal
C	60%	attack
D	40%	normal
E	30%	attack
F	10%	normal

At thresholds of 70% and 30%, list the alerted events, count normal events alerted, and count attacks missed. Include scores equal to the threshold.

Worked reasoning. At 70%, A and B alert: B is one normal event alerted, while attacks C and E are missed. At 30%, A, B, C, D and E alert: B and D are two normal events alerted, and no attack in this sample is missed. Lowering the threshold catches C and E, but adds normal D to the review queue. Even a high threshold can produce a false alert. These invented scores and investigated outcomes are not measured detector accuracy or a guarantee about future events.

Capacity transfer. Suppose the team can review only two alerts in the period. The 70% setting fits that capacity but misses two known attacks; the 30% setting produces five alerts and exceeds capacity, although it catches all three attacks in this sample. Neither is automatically best. Weigh the harm of misses, the cost of false alerts and investigation capacity. Priorities and other controls matter too; a threshold alone is not a complete defence.

Too many false alerts cause **alert fatigue** : responders get so used to false positives that they start assuming an alert is false *before* investigating it - so a real attack, when it finally comes, is waved away. This is exactly why a **low false-positive rate matters**: signature-based detection often has fewer false alarms, while anomaly-based detection can flag unfamiliar activity but may add false alarms. Actual results depend on the signatures, baseline, thresholds and traffic; neither method guarantees detection.



AP Cybersecurity · U3 · signature-based vs anomaly-based detection

Signature-based detection matches known attacks; anomaly-based detection flags deviations from normal

Exam tips

- For firewall-ACL questions, **read the rules top-to-bottom and stop at the first match** - a Deny rule above an Allow blocks the traffic even though the Allow exists lower down.
- Pair each attack with its **tell-tale sign**: ARP poisoning = one IP with two MAC addresses; MAC flooding = a surge of new MAC addresses; DNS poisoning = an unexplained drop in web traffic.
- Read **packet captures** for network-based IoCs: known-malicious IPs, unauthorized scans, traffic spikes/slowdowns, and mismatched port-application traffic.
- Run **vulnerability scanners** to find *known* weaknesses proactively, and fix the highest-severity findings first.
- **Signature-based** matches known indicators; **anomaly-based** compares a baseline. Discuss coverage, false alerts and resource costs in the scenario; neither guarantees detection.
- A **screened subnet** / **DMZ** holds public-facing servers between the internet and the private network - name it whenever a question separates public services from internal data.
- **WPA3** is the strong wireless encryption; **WEP** and original **WPA** are insecure.