

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** 地址解析投毒 - the **address resolution protocol (ARP)** 地址解析协议 pairs IP addresses with hardware **MAC addresses** 物理地址. An adversary sends fake ARP messages so traffic meant for the target flows to the adversary instead. This is an **on-path attack** 中间人攻击 (also called man-in-the-middle): the adversary secretly sits between two parties, reading and even altering their messages.
- **MAC flooding** 物理地址泛洪 - flooding a **switch** 交换机 with fake MAC addresses forces it into broadcast mode, so the adversary can capture all traffic. This is a form of **eavesdropping** 窃听.
- **DNS poisoning** 域名投毒 - planting a fake record on a **domain name system (DNS)** 域名系统 server redirects users to a malicious site to steal credentials (**credential harvesting** 凭据收集).
- **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)** 拒绝服务 attack; when many machines attack at once it becomes a **distributed denial of service (DDoS)** 分布式拒绝服务.

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** 非法接入点 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** 自动漏洞扫描器: 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** 缓解措施. 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** 分离隧道; 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** 无线加密协议 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 网络分段 divides one network into smaller, isolated pieces (**subnets** 子网). If one subnet is breached, the damage is **contained** and cannot spread.

A key pattern is the **screened subnet** 屏蔽子网 (also called a **DMZ** 隔离区). 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.



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** 虚拟局域网 (logically separating devices on the same switch). Each segment can then get its own security policy - higher-security and lower-security zones.

Protecting Networks: Firewalls

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

- **Stateless** 无状态 - filters on packet headers alone (IP, port, protocol).
- **Stateful** 有状态 - 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)** 访问控制列表 - 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).



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

Image: Dsimic, CC BY-SA 4.0 (commons.wikimedia.org)

Detecting Network Attacks

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

- a **network intrusion detection system (NIDS)** 网络入侵检测系统 analyses traffic and **raises an alert**, but does not block;
- a **network intrusion prevention system (NIPS)** 网络入侵防御系统 can also **stop** an attack by closing ports or blocking addresses;
- a **security information and event management (SIEM)** 安全信息与事件管理系统 gathers data from many sources to spot patterns.

There are two detection methods. **Signature-based** 基于特征 detection compares traffic to a database of known attack **signatures** - fast and low on false alarms, but blind to brand-new attacks. **Anomaly-based** 基于异常 detection compares traffic to a normal **baseline** 基线 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** 网络入侵指标 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.

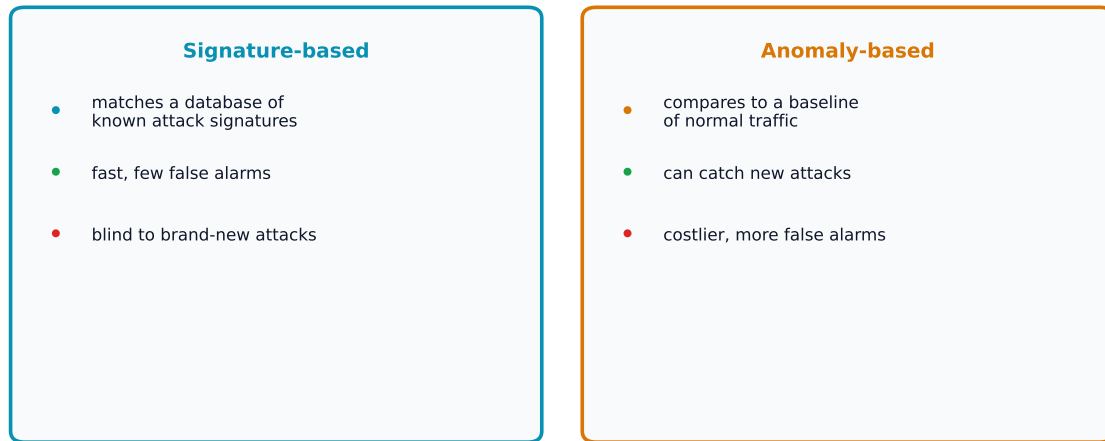
AI, thresholds, and alert fatigue

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** 概率的: rather than a yes/no, each event gets a **percentage** likelihood of being malicious.

The organisation then sets a **threshold** 阈值 - 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.

Too many false alerts cause **alert fatigue** 警报疲劳: 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 has almost none, while anomaly-based and hybrid detection trade a higher false-positive rate for the ability to catch novel attacks.



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** = fast, few false positives, misses new attacks (more false negatives); **anomaly-based** = catches new attacks, costs more, more false positives. Memorise this trade-off.
- 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.