Securing Networks

AP Cybersecurity ▪ Unit 3

AP Cybersecurity

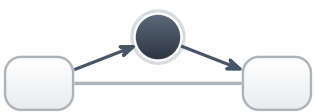

What we will cover

1. Network attacks
2. Policies & wireless security
3. Segmentation & the DMZ
4. Firewalls
5. Access control lists
6. Detecting attacks

1 Network attacks

The classic network attacks

ARP poisoning 地址解析投毒
fake ARP messages divert traffic – an **on-path attack** 中间人攻击

DNS poisoning 域名投毒
a fake record redirects users – **credential harvesting** 凭据收集

MAC flooding 物理地址泛洪
floods a **switch** 交换机 into broadcast mode – **eavesdropping** 窃听

Smurf attack
ICMP to the broadcast address – a **DoS** 拒绝服务; many at once = **DDoS** 分布式拒绝服务

Pair each attack with its tell-tale sign: ARP poisoning = one IP with two MACs; MAC flooding = a surge of new MACs; DNS poisoning = an unexplained drop in web traffic.

AI, thresholds, and alert fatigue

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

Finding weaknesses, and reading the traffic

an automated vulnerability scanner 自动漏洞扫描器	run to find weaknesses before an adversary does
what it checks	open ports, unpatched software, default credentials, weak configuration
the limit	it finds known weaknesses only, and it generates work rather than fixing anything
packet capture	analysts examine captured traffic hunting for network-based indicators of compromise
what those look like	traffic to a known bad address, unusual volumes, protocols on the wrong port, beaconing at a regular interval
an ACL 访问控制列表	an access control list on a router or firewall, stating which traffic is permitted

> A scan tells you what *could* be exploited; packet capture tells you what *is being* exploited. Both are needed, and they answer different questions.

AI, thresholds, and alert fatigue, in detail

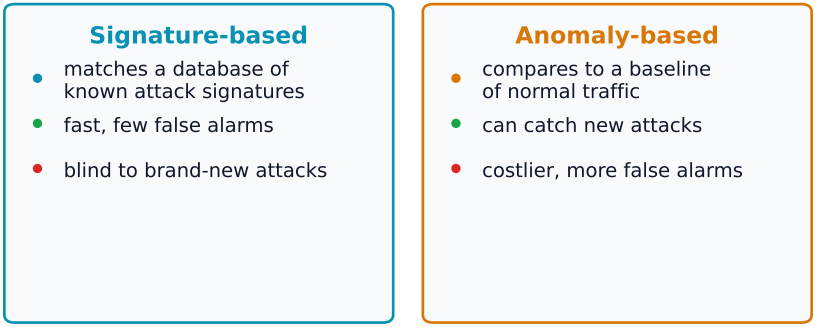

AP Cybersecurity · U3 · signature-based vs anomaly-based detection

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

Why weak networks fall

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

Network switches

2 Policies & wireless

Policies come first

Written policies set a minimum standard:

- router & switch policies – ban local accounts, require port security
- a VPN policy – authentication rules, and no **split tunneling** 分离隧道
- a wireless policy – strong encryption, authenticated access

disable beacon frames – harder to find

control signal strength
– no leak outside

WPA3 无线加密协议
– the current strongest

MAC filtering – known devices only

> WPA3 is strong; WEP and the original WPA are broken – never accept them in an answer.

3 Segmentation

Segmentation and the screened subnet

Network segmentation 网络分段 divides one network into isolated **subnets** 子网. If one is breached, the damage is contained.

> Build segments by subnetting (IP) or with **VLANs** 虚拟局域网 (logically, on one switch). Each gets its own policy.

Segmentation and the screened subnet, in detail

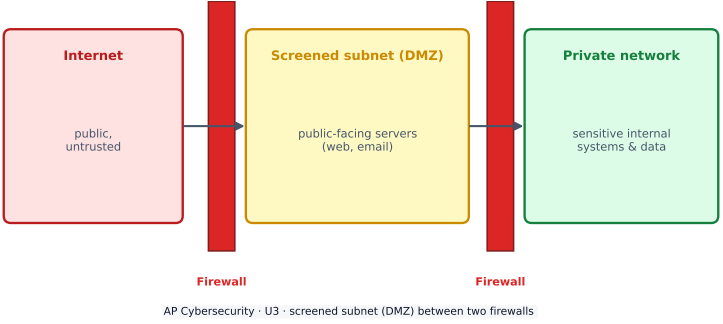

public servers between two firewalls

> If one is breached, the damage is contained.

4 Firewalls

Three kinds of firewall

Stateless 无状态
filters on headers alone – IP, port, protocol

Stateful 有状态
also tracks each connection's state – finer control

Next-generation
adds intrusion prevention, deep packet inspection

> Firewalls belong at every point where data crosses between zones – each segment, and every gateway to the internet.

Filtering, and the protocols attackers abuse

Access control list	An **access control list (ACL)** on a router or firewall permits or denies traffic by address, port and protocol – the coarsest and cheapest filter.
ARP	The **address resolution protocol (ARP)** maps an IP address to a MAC address on the local network, and it trusts any reply – which is what makes ARP spoofing possible.
DNS	The **domain name system (DNS)** maps names to addresses; poisoning it sends a visitor to the attacker's server with the right name in the bar.
DoS and DDoS	A **denial of service (DoS)** exhausts a resource; a **distributed denial of service (DDoS)** does it from thousands of machines at once, so blocking one source achieves nothing.

> Both protocols were designed for a trusted network. Every attack on them exploits that assumption, not a bug.

The access control list – first match wins

A firewall follows an **ACL** 访问控制列表 – an ordered set of rules. Each names a direction, a filter (IP, port, service) and an action.

> Rules are checked in order, and the first match wins – so rule order, not just rule content, decides what gets through.

The access control list – first match wins, in detail

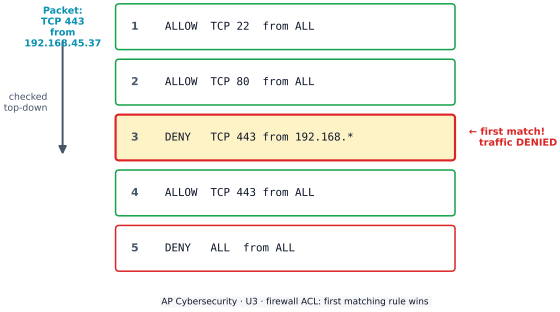

checked top to bottom

A firewall follows an **ACL** 访问控制列表 – an ordered set of rules, checked top to bottom, where the first match wins.

Worked example – why the ALLOW rule fails

Rule 3: `DENY TCP 443 from 192.168.*`
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. The fix: move the ALLOW rule above the DENY.

> Read ACLs top-to-bottom and stop at the first match. A Deny above an Allow blocks the traffic, however many Allows sit below it.

5 Detection

When prevention fails – detection takes over

Automated tools read the **log files** 日志文件 that record network activity:

NIDS 网络入侵检测系统
– alerts, does not block

NIPS 网络入侵防御系统
– can also stop an attack

SIEM 安全信息与事件管理 –
gathers many sources to spot patterns

The D in NIDS is detect; the P in NIPS is prevent.

Two detection methods – memorise the trade-off

Signature-based 基于特征

matches known attack signatures

✓ fast, few false positives

✗ blind to brand-new attacks

Anomaly-based 基于异常

compares to a normal **baseline** 基线

✓ catches novel attacks

✗ costs more, more false positives

> A hybrid approach combines both – and is the safe answer when a question asks for the best of each.

Detecting, preventing, and correlating

NIDS	A **network intrusion detection system (NIDS)** watches traffic and raises an alert. It sees everything and stops nothing.
NIPS	A **network intrusion prevention system (NIPS)** sits inline and can drop the packet – which means a false positive becomes an outage.
SIEM	**Security information and event management (SIEM)** collects logs from everywhere and correlates them, so a pattern invisible on one device becomes visible across many.
Anomaly detection	Signature matching catches the known; **probabilistic** anomaly detection flags the unusual, and produces most of the false positives.

Detection buys time; **mitigation** is what you do with it. Name both halves in an incident answer.

Why a detection system needs a threshold

the volume	a medium network logs millions of events a day – far more than any team can read
so	organisations score events by how likely each is to be an attack
the threshold 阈值	the likelihood at which an alert actually fires
set it low	few real attacks are missed, and the team drowns in false positives
set it high	the queue is manageable, and real attacks slip through
alert fatigue 警报疲劳	too many false alerts and responders get so used to them that a real one is dismissed

> The threshold is a deliberate trade between missed attacks and wasted attention, and there is no setting that avoids both. Naming that trade is the answer.

Protecting Networks: Segmentation

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

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.

Exam tips (continued)

- **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.

6 Recap

Unit 3 – key takeaways

1 The tells

ARP = one IP two MACs; MAC flood = a surge of MACs

2 DMZ

public servers between the internet and the private network

3 ACL order

first match wins – a Deny above an Allow blocks

4 Detection

signature is fast but blind; anomaly catches new, cries wolf

Check yourself

1. A DENY sits above an ALLOW for the same traffic. What happens?
2. Which detection method misses brand-new attacks?
3. Which wireless encryption is currently strongest?
4. NIDS or NIPS: which can block an attack?

Answers **1.** it is denied – first match wins **2.** signature-based **3.** WPA3 **4.** NIPS