Cybersecurity Basics: A Beginner's Guide to GRC, SOC & VAPT

Nov 16, 2026 · @Vaishnavi Sharma

Start here

This guide gives beginners and career switchers the core knowledge expected for an entry-level cybersecurity role. It covers the three areas most job postings fall into: GRC, SOC and VAPT. No prior security experience is assumed.

The three career areas

Area	What it means	Typical entry roles	Question it answers	Suits people who enjoy
GRC (Governance, Risk & Compliance)	Security rules, risk decisions and audits	GRC Analyst, Compliance Associate, IT Auditor, Risk Analyst	Are we doing security the right way, and can we prove it?	Writing, process, working with business teams
SOC (Security Operations Centre)	24x7 monitoring and response to attacks	SOC Analyst L1, Security Monitoring Analyst	Is someone attacking us right now?	Investigating, pattern-spotting, fast decisions
VAPT (Vulnerability Assessment & Penetration Testing)	Finding weaknesses before attackers do	VAPT Analyst, Junior Penetration Tester	Where can an attacker get in?	Puzzles, breaking things, deep technical work

Most beginners start in SOC or GRC, because those roles hire more freshers. VAPT roles usually expect some hands-on proof such as lab write-ups or bug bounty findings.

Coming from another field? Your experience counts

Previous background	Skills that transfer	Natural first area
Customer or technical support, BPO	Ticketing, triage, SLAs, escalation, documentation, handling sensitive customer data	SOC
IT helpdesk or system administration	Windows, Active Directory, networking, troubleshooting	SOC, then VAPT
Engineering (any discipline)	Systems thinking, root-cause analysis, process discipline	SOC or VAPT
Audit, finance, legal, quality	Controls, evidence, regulations, documentation	GRC
Software development	Code, web applications, APIs	VAPT (application security)
Teaching or training	Explaining complex ideas, structured content	GRC (security awareness), then any

In interviews, connect your old work to security in concrete terms. For example: "I verified customer identity before account changes, which is authentication; I escalated suspicious account activity, which is incident triage."

How to use this guide

- **Study order:** Foundations → Networking & Systems → SOC → GRC → VAPT → Cloud. Sections 1 and 2 underpin everything else, so don't skip them.
- **For each section:** read it, look up every term you can't explain in one sentence, then do the matching lab from section 7.
- **Learn in public:** after each section, write a short post about one thing you learned. Recruiters notice consistent learners.
- **Depth:** aim to explain every heading in this guide out loud in under two minutes. That's roughly the level an entry-level interview tests.

1. Core foundations

Every security decision protects one of three properties, called the CIA triad. It is the single most-asked concept in entry-level interviews.

The CIA triad

Principle	Meaning	Example attack	Example control
Confidentiality	Only authorised people can see data	Stolen customer database	Encryption, access control
Integrity	Data isn't changed without permission	Attacker edits a bank transfer amount	Hashing, digital signatures, change control
Availability	Systems work when needed	DDoS attack takes a website offline	Backups, redundancy, DDoS protection

Some frameworks extend this with **non-repudiation** (someone can't deny an action they took, e.g. through digital signatures and logs) and **authenticity** (data comes from who it claims to).

AAA: how access works

- **Authentication:** proving who you are. There are three factor types: something you know (password, PIN), something you have (phone, hardware key), something you are (fingerprint, face). **MFA** combines two or more different types; two passwords is not MFA.
- **Authorization:** what you're allowed to do once logged in, e.g. a support agent can view an account but not delete it.
- **Accounting (auditing):** recording what you did in logs. Security monitoring runs on these records.

Asset, vulnerability, threat, risk

- **Asset:** anything worth protecting: data, servers, laptops, people, reputation.
- **Vulnerability:** a weakness, such as unpatched software, a weak password, or an untrained employee.

- **Threat:** something that could exploit a weakness: a criminal group, ransomware, a careless insider, a flood.
- **Exploit:** the actual method or code used to take advantage of a vulnerability.
- **Risk:** the chance of harm combined with its impact. Risk = likelihood × impact.

Example: an unpatched VPN server (vulnerability) + ransomware gangs actively scanning for it (threat) = high risk to company data (asset).

Security principles

- **Least privilege:** each person and system gets only the access their job needs, and nothing more.
- **Defense in depth:** several layers of protection (firewall, MFA, EDR, backups, training) so one failure isn't fatal.
- **Zero trust:** "never trust, always verify". Every request is checked, even from inside the network.
- **Separation of duties:** no single person controls a whole sensitive process (e.g. one person requests a payment, another approves it).
- **Secure by default:** systems ship with safe settings switched on.
- **Need to know:** access to specific data only when a task requires it.

Cryptography basics

Concept	What it does	Examples	Used for
Symmetric encryption	One shared key encrypts and decrypts; fast	AES-256	Encrypting disks, files, bulk data
Asymmetric encryption	Public key encrypts, private key decrypts; slower	RSA, ECC	Key exchange, digital signatures
Hashing	One-way fingerprint of data; can't be reversed	SHA-256 (MD5 and SHA-1 are outdated)	Password storage, file integrity, malware IDs
Digital signature	Proves who sent something and that it wasn't changed	RSA or ECDSA signatures	Software updates, documents, email

Concept	What it does	Examples	Used for
TLS	Encrypts data in transit using all of the above	HTTPS padlock	Web, email, APIs

Key ideas: **encryption at rest** protects stored data; **encryption in transit** protects data moving across a network. Passwords should be stored as **salted hashes** (a random value added before hashing) using slow algorithms like bcrypt or Argon2, never as encrypted or plain text.

Common attacks

Attack	What happens	Main defences
Phishing	Fake email or message tricks users into clicking, opening a file or sharing credentials	Email filtering, user training, MFA
Spear phishing / BEC	Targeted phishing; Business Email Compromise impersonates an executive or vendor to redirect payments	Payment verification calls, DMARC, training
Smishing / vishing	Phishing by SMS or voice call	Training, caller verification procedures
Malware	Malicious software: viruses, worms, trojans, spyware, keyloggers	EDR/antivirus, patching, least privilege
Ransomware	Encrypts files and demands payment; often steals data first (double extortion)	Offline backups, EDR, patching, MFA
Brute force / password spraying / credential stuffing	Guessing passwords, trying one common password on many accounts, or reusing leaked ones	MFA, lockout policies, breached-password checks
Social engineering	Manipulating people, e.g. a fake "IT support" call asking for an OTP	Verification procedures, training

Attack	What happens	Main defences
DoS / DDoS	Floods a service until it fails	DDoS protection, rate limiting, CDNs
SQL injection / XSS	Malicious input attacks a web application	Input validation, parameterised queries, WAF
Man-in-the-middle	Attacker intercepts traffic between two parties	TLS everywhere, VPNs, avoiding open Wi-Fi
Supply chain attack	Attacker compromises a vendor or software update to reach its customers	Vendor risk management, software integrity checks
Insider threat	An employee or contractor misuses access, deliberately or by accident	Least privilege, monitoring, offboarding processes

2. Networking & systems fundamentals

You can't recognise an attack in traffic or logs until you know what normal looks like, which makes this the most important technical section for every security role.

The OSI and TCP/IP models

These models describe how data travels across a network, one layer at a time.

OSI layer	Name	What it does	Examples	Attacks at this layer
7	Application	What applications speak	HTTP, DNS, SMTP, FTP	SQL injection, phishing
6	Presentation	Formatting and encryption	TLS/SSL, JPEG	SSL stripping
5	Session	Keeps conversations open	Login sessions, RPC	Session hijacking
4	Transport	Reliable or fast delivery	TCP, UDP	SYN flood

OSI layer	Name	What it does	Examples	Attacks at this layer
3	Network	Routing between networks	IP, ICMP, routers	IP spoofing, ping floods
2	Data link	Delivery on the local network	MAC addresses, switches, ARP	ARP spoofing
1	Physical	Cables and signals	Ethernet, Wi-Fi, fibre	Cable tapping, jamming

The TCP/IP model groups these into 4 layers: Application (OSI 5–7), Transport (4), Internet (3), Network Access (1–2).

Key concepts

- **IP addresses:** IPv4 looks like 192.168.1.10; IPv6 like 2001:db8::1. Private ranges (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16) are used inside networks; public addresses face the internet. **NAT** translates between them.
- **Subnetting and CIDR:** 192.168.1.0/24 means the first 24 bits are the network, leaving 256 addresses (254 usable).
- **MAC address:** a hardware address on the local network. **ARP** maps IP addresses to MAC addresses.
- **TCP vs UDP:** TCP is reliable and connection-based, starting with the three-way handshake (SYN → SYN-ACK → ACK). UDP is faster with no delivery guarantee, used for DNS, video and voice.
- **DNS:** translates names (example.com) to IP addresses. Record types to know: A (IPv4), AAAA (IPv6), MX (mail server), CNAME (alias), TXT (used for SPF and DMARC). Attackers abuse DNS for phishing domains, data exfiltration and command-and-control.
- **DHCP:** automatically assigns IP addresses to devices joining a network.
- **HTTP and HTTPS:** HTTPS is HTTP inside TLS encryption. Methods: GET (fetch), POST (send). Status codes: 200 OK, 301/302 redirect, 401 unauthenticated, 403 forbidden, 404 not found, 500 server error.
- **Network devices:** router (connects networks), switch (connects devices in a network), firewall (allows or blocks traffic by rules), proxy (sits between users and the web), load balancer (spreads traffic across servers).

- **Segmentation:** splitting a network into zones (VLANs, a DMZ for internet-facing servers) so an attacker in one zone can't reach everything.
- **VPN:** an encrypted tunnel across the internet, used for remote access.

Ports to memorise

Port	Protocol	Notes for security
20/21	FTP	Unencrypted file transfer; replace with SFTP
22	SSH / SFTP	Encrypted remote login; frequent brute-force target
23	Telnet	Unencrypted remote login; should be disabled
25	SMTP	Email sending; spam and phishing path
53	DNS	Name lookups; watch for unusual volumes (tunnelling)
80 / 443	HTTP / HTTPS	Web traffic
88	Kerberos	Windows domain authentication
110 / 143	POP3 / IMAP	Email retrieval
389 / 636	LDAP / LDAPS	Directory queries (Active Directory)
445	SMB	Windows file sharing; abused by ransomware such as WannaCry
1433 / 3306	MS SQL / MySQL	Databases; should never face the internet
3389	RDP	Windows remote desktop; a top ransomware entry point

Linux essentials

Most servers, cloud workloads and security tools run on Linux. Get comfortable in the terminal:

Task	Commands
Navigate	`pwd`, `ls -la`, `cd`, `find / -name "*.log"`

Task	Commands
Read files	`cat`, `less`, `head`, `tail -f` (follow a log live)
Search text	`grep "Failed password" /var/log/auth.log`, `grep -i`, `wc -l`
Permissions	`chmod`, `chown`, reading `-rwxr-x---`
Users	`whoami`, `id`, `sudo`, `/etc/passwd`, `/etc/shadow` (password hashes)
Processes	`ps aux`, `top`, `kill`
Network	`ip a`, `ss -tulpn` (listening ports), `ping`, `curl`, `dig`
Scheduled tasks	`crontab -l` (a common persistence spot for attackers)

Key log files: `/var/log/auth.log` or `/var/log/secure` (logins and sudo), `/var/log/syslog` (system events), web server logs under `/var/log/apache2/` or `/var/log/nginx/`.

Windows and Active Directory

Most company laptops run Windows, and most companies manage them with Active Directory (AD), so this is where many investigations happen.

- **Active Directory:** the central directory of users, computers, groups and policies in a company network. A **domain controller** runs it. Attackers who reach **Domain Admin** effectively control the company.
- **Group Policy (GPO):** pushes settings to every machine, such as password rules.
- **Authentication:** Kerberos (the modern default) and NTLM (older, often abused in pass-the-hash attacks).
- **PowerShell:** a powerful scripting shell and a favourite attacker tool. Encoded or obfuscated PowerShell commands are a classic alert.
- **Sysmon:** a free Microsoft tool that adds detailed process, network and file logging.

Windows Event IDs to memorise:

Event ID	Meaning	Why it matters
4624	Successful logon	Check logon type: 2 = keyboard, 3 = network, 10 = RDP
4625	Failed logon	Many in a row = brute force
4648	Logon using explicit credentials	Possible lateral movement
4672	Special privileges assigned	An admin logged in
4688	New process created	What ran, and what launched it
4720	User account created	Attackers create accounts for persistence
4728 / 4732	User added to a security group	Watch for additions to admin groups
4740	Account locked out	Brute force or a forgotten password
7045	New service installed	Common malware persistence
1102	Audit log cleared	Attacker covering tracks; always investigate

Logs: the raw material of security

A log is a timestamped record of an event. Common sources: firewalls, proxies, DNS servers, Windows Event Logs, Linux syslog, web servers, email gateways, EDR tools, VPNs, and cloud platforms (AWS CloudTrail, Azure Activity Logs, Entra ID sign-in logs).

For every log line, practise answering five questions: **who** did **what**, **when**, from **where**, and **is that normal** for them? Timestamps are often in UTC, so always confirm the time zone before building a timeline.

3. GRC: Governance, Risk & Compliance

GRC is the business side of security: setting the rules, deciding which risks to fix first, and proving to auditors, regulators and customers that controls actually work. It needs less coding and more writing, analysis and communication.

The three parts

- **Governance:** who is responsible for security and what the rules are. It covers the security strategy, roles (Board, CISO, risk owners), policies, and regular reporting to leadership.
- **Risk management:** finding risks, scoring them, and deciding what to do about each one.
- **Compliance:** meeting the laws, standards and contract terms that apply, and keeping evidence that you do.

The policy hierarchy

Document	What it is	Example
Policy	High-level rule; mandatory; approved by management	"All remote access must use MFA."
Standard	Specific mandatory requirement supporting a policy	"Passwords must be at least 14 characters."
Procedure	Step-by-step instructions	"How to reset a user's MFA device"
Guideline	Recommended, not mandatory, advice	"Tips for spotting phishing"

Common policies: Information Security Policy, Acceptable Use, Access Control, Password, Data Classification, Incident Response, Backup, Vendor Management, Remote Work.

Risk management in practice

1. **Identify** assets and the threats and vulnerabilities affecting them.
2. **Assess** each risk. Qualitative scoring rates likelihood and impact from 1 to 5 each (risk score = likelihood × impact, so 1–25). Quantitative scoring uses money: Annualised Loss Expectancy (ALE) = Single Loss Expectancy (SLE) × Annual Rate of Occurrence (ARO).
3. **Record** each risk in a **risk register**: description, owner, score, treatment, due date, status.
4. **Treat** each risk in one of four ways:
    - **Mitigate:** add controls to reduce it (e.g. enable MFA).
    - **Transfer:** shift the financial impact (cyber insurance, outsourcing).

    - **Accept:** consciously live with it, with sign-off from the risk owner.
    - **Avoid:** stop the risky activity altogether.
5. **Monitor** regularly. **Inherent risk** is the level before controls; **residual risk** is what remains after them. **Risk appetite** is how much risk leadership is willing to accept.

Types of controls

By purpose	Example
Preventive	Firewall rules, MFA, door locks
Detective	SIEM alerts, CCTV, log reviews
Corrective	Restoring backups, patching after an incident
Deterrent	Warning banners, visible cameras
Compensating	Extra monitoring where a primary control isn't possible

By type	Example
Administrative (managerial)	Policies, training, background checks
Technical (logical)	Encryption, antivirus, access control lists
Physical	Locks, badges, guards

Data classification

Data is labelled by sensitivity so it gets the right protection, typically as **Public**, **Internal**, **Confidential** and **Restricted**. Personal data (names, phone numbers, ID numbers) and sensitive personal data (health, financial) usually sit in the top two levels.

Frameworks, standards and laws

Name	What it is	Who uses it
ISO/IEC 27001:2022	International standard for an Information Security Management System (ISMS). Annex A lists 93 controls in 4 themes: organisational, people, physical, technological	Organisations seeking certification; very common in IT services and BPO
ISO/IEC 27002	Detailed guidance on implementing the 27001 controls	Alongside 27001
NIST Cybersecurity Framework 2.0	US framework with 6 functions: Govern, Identify, Protect, Detect, Respond, Recover	Widely used worldwide as a reference model
NIST SP 800-53	Large catalogue of security controls	US government and its suppliers
CIS Controls v8	18 prioritised, practical security controls	Organisations wanting a practical starting point
SOC 2	US audit report (an attestation, not a certification) covering security, availability, processing integrity, confidentiality and privacy	SaaS and service companies with US clients
PCI DSS v4.0	Security standard for payment card data	Anyone storing, processing or transmitting card data
GDPR	EU data protection law; fines up to €20 million or 4% of global annual turnover	Anyone processing personal data of people in the EU
DPDP Act 2023 (India)	India's Digital Personal Data Protection Act, with its rules being phased in	Anyone processing digital personal data of people in India
CERT-In Directions 2022 (India)	Report specified cyber incidents to CERT-In within 6 hours	Organisations operating in India

Name	What it is	Who uses it
HIPAA	US health data law	Healthcare providers and their vendors

Local laws and timelines change, so check the official source for the countries you work with.

Audits and evidence

- **Internal audit:** the organisation checks itself. **External audit:** an independent auditor checks, e.g. for ISO 27001 certification or a SOC 2 report.
- **Evidence:** proof a control works, such as access review records, screenshots of settings, training completion reports, change tickets and logs.
- **Gap analysis:** comparing current practice against a framework to list what's missing, then planning fixes.
- **Findings:** audit results are usually rated (e.g. major or minor non-conformity) and tracked to closure in a remediation plan.
- **Third-party (vendor) risk management:** checking that suppliers protect your data, through security questionnaires, certifications, contract clauses and periodic reviews.
- **Security awareness:** regular training and phishing simulations, often a compliance requirement.

Business continuity and disaster recovery

- **BCP (Business Continuity Plan):** how the business keeps running during a disruption.
- **DRP (Disaster Recovery Plan):** how IT systems are restored.
- **BIA (Business Impact Analysis):** identifies critical processes and the cost of downtime.
- **RTO (Recovery Time Objective):** how quickly a system must be back up.
- **RPO (Recovery Point Objective):** how much data loss, measured in time, is acceptable. An RPO of 4 hours means backups at least every 4 hours.
- **3-2-1 backup rule:** 3 copies of data, on 2 different media, with 1 offsite (and ideally one offline or immutable, to survive ransomware).

What a GRC analyst does day to day

Maintains the risk register, collects audit evidence, reviews and updates policies, sends and reviews vendor security questionnaires, tracks remediation of findings, runs access reviews, and prepares reports for management. Excel, clear writing and stakeholder follow-up are core skills.

4. SOC: Security Operations Centre

A SOC is the team that monitors an organisation's systems around the clock, investigates alerts, and responds to attacks. Day to day it looks like a support desk: shifts, a ticket queue, triage, SLAs, documentation and escalation, except every ticket is a possible attack.

SOC structure

Tier	Role	Main work
L1	SOC Analyst (entry level)	Monitors the alert queue, triages alerts, closes false positives, escalates real incidents
L2	Incident Responder	Investigates escalated incidents in depth, contains and cleans up threats
L3	Threat Hunter / Detection Engineer	Hunts for hidden threats, writes and tunes detection rules
Lead	SOC Manager	Runs shifts, metrics, reporting and improvement

SOCs are either **in-house** (one company protecting itself) or run by an **MSSP / MDR provider** that monitors many client companies. MSSPs hire many freshers and are a common first job.

Core tools

Tool type	What it does	Examples
SIEM	Collects logs from everywhere, correlates them, raises alerts	Splunk, Microsoft Sentinel, IBM QRadar, Elastic, Google SecOps

Tool type	What it does	Examples
EDR / XDR	Monitors endpoints (laptops, servers) for malicious behaviour; can isolate machines	Microsoft Defender for Endpoint, CrowdStrike Falcon, SentinelOne
SOAR	Automates repetitive response steps, such as blocking an IP or enriching an alert	Splunk SOAR, Sentinel playbooks, Cortex XSOAR
IDS / IPS	Detects or blocks suspicious network traffic	Snort, Suricata, Zeek
Email security	Filters phishing and malware in email	Microsoft Defender for Office 365, Proofpoint, Mimecast
Threat intelligence	Checks whether an IP, domain, URL or file is known to be malicious	VirusTotal, AbuseIPDB, URLScan.io, Cisco Talos, AlienVault OTX
Ticketing / case management	Tracks every alert and incident	ServiceNow, Jira, TheHive
Sandbox	Safely runs suspicious files to watch what they do	ANY.RUN, Hybrid Analysis, Joe Sandbox

Alert outcomes

Outcome	Meaning	Action
True positive	Real malicious activity, correctly alerted	Escalate and respond
Benign true positive	The rule worked, but the activity was authorised (e.g. an IT admin's scan)	Close with notes; consider tuning
False positive	Alert fired, but nothing malicious happened	Close with notes; suggest tuning
False negative	Real attack, no alert	The dangerous one; found by hunting or after damage

How to triage any alert

1. **Read the alert:** which rule fired, on which user or host, at what time, and with what severity?
2. **Scope it:** is this one machine or many? Is the user a VIP or the host a critical server?
3. **Gather context:** is this behaviour normal for this user or system? Check their recent activity.
4. **Enrich indicators:** look up IPs, domains, URLs and file hashes in threat intelligence tools.
5. **Check surrounding logs:** what happened just before and after the alert?
6. **Decide:** true or false positive, and the real severity.
7. **Document and act:** write clear ticket notes (what you checked, what you found, your verdict, why) and escalate or close within SLA.

Good ticket notes are a skill on their own: another analyst should understand your conclusion without redoing your work.

Worked example 1: brute force followed by a successful login

Alert: 50 failed logins (Event ID 4625) for user `a.verma` within 5 minutes, then one successful login (4624, logon type 10, i.e. RDP).

- The source IP is outside the country, and AbuseIPDB shows hundreds of abuse reports for it.
- The user normally logs in from one city during office hours, so this is abnormal.
- A success after many failures suggests the password was guessed.
- **Verdict:** true positive, high severity.
- **Actions:** escalate to L2; recommend disabling the account, resetting the password and revoking sessions; block the IP; review what the account did after logging in (new processes, new accounts, file access).

Worked example 2: suspicious PowerShell

Alert: EDR flags `powershell.exe -nop -w hidden -enc JABjAGwA...` launched by `WINWORD.EXE` on a finance laptop.

- Word should never launch PowerShell, so this parent-child relationship is a strong red flag (likely a malicious macro).

- `-enc` means a Base64-encoded command, and `-w hidden` hides the window: both are classic evasion signs.
- Decoding the command (e.g. with CyberChef) shows it downloading a file from an unknown domain.
- **Verdict:** true positive, critical. Isolate the host through EDR, find the email that delivered the document, and check whether other users received it.

Searching logs in a SIEM

Every SIEM has a query language. These two queries both find possible brute force: more than 20 failed logins for one account from one IP within 5 minutes.

Microsoft Sentinel (KQL):

```
SecurityEvent
| where EventID == 4625
| summarize Failures = count() by Account, IpAddress, bin(TimeGenerated, 5m)
| where Failures > 20
| sort by Failures desc
```

Splunk (SPL):

```
index=windows EventCode=4625
| bin _time span=5m
| stats count AS failures BY Account_Name, src_ip, _time
| where failures > 20
| sort - failures
```

You don't need to memorise syntax for an entry-level interview, but you should be able to read a query like this and explain what it does.

Phishing email analysis

Phishing is the most common alert type. Work through it in this order:

1. **Sender:** lookalike domains (paypa1-support.com), free email accounts claiming to be companies, or a display name that doesn't match the address.
2. **Headers:** SPF, DKIM and DMARC results show whether the email really came from the domain it claims. Also compare the Return-Path and Reply-To addresses, and trace the Received lines.

3. **Links:** hover to see the real destination; check it in URLScan.io or VirusTotal. Never click it directly.
4. **Attachments:** check the file hash in VirusTotal; be wary of .zip, .iso, .html, .lnk, .one and macro-enabled Office files. Detonate suspicious files only in a sandbox.
5. **Content:** urgency, threats, unusual payment requests, requests for credentials or OTPs.
6. **Scope and respond:** who else received it? Did anyone click or reply? Remove it from all mailboxes and block the sender, domain and URLs.

Incident response lifecycle (NIST SP 800-61)

1. **Preparation:** plans, contact lists, tools and training ready before anything happens.
2. **Detection & analysis:** spot the incident, confirm it, and work out its scope and severity.
3. **Containment:** stop the spread, e.g. isolate hosts, disable accounts, block IPs. Short-term containment comes first, then longer-term fixes.
4. **Eradication:** remove the cause, e.g. delete malware, remove attacker accounts, close the vulnerability.
5. **Recovery:** restore systems from clean backups, monitor closely for the attacker's return.
6. **Post-incident activity (lessons learned):** review what happened and improve, which feeds back into Preparation.

Preserve evidence throughout: record timestamps, don't power off machines unnecessarily, and keep a **chain of custody** if legal action is possible.

MITRE ATT&CK and the Cyber Kill Chain

MITRE ATT&CK (attack.mitre.org) is a free knowledge base of real-world attacker behaviour. **Tactics** are the attacker's goals; **techniques** are how they achieve them. For example, T1566 is Phishing (Initial Access), T1059.001 is PowerShell (Execution), T1110 is Brute Force (Credential Access). SOC teams map alerts to ATT&CK to understand attacks and find gaps in detection.

The 14 enterprise tactics, in rough attack order: Reconnaissance, Resource Development, Initial Access, Execution, Persistence, Privilege Escalation, Defense Evasion, Credential Access, Discovery, Lateral Movement, Collection, Command and Control, Exfiltration, Impact.

The older **Cyber Kill Chain** (Lockheed Martin) describes 7 stages: Reconnaissance, Weaponisation, Delivery, Exploitation, Installation, Command & Control, Actions on Objectives. Breaking any link in the chain stops the attack.

Threat intelligence basics

- **IOCs (Indicators of Compromise):** specific evidence such as malicious IPs, domains, file hashes. Easy for attackers to change.
- **TTPs (Tactics, Techniques & Procedures):** how an attacker behaves. Much harder to change, so detections built on TTPs last longer.
- The **Pyramid of Pain** ranks indicators from trivial for attackers to change (hashes, IPs) to painful (tools, TTPs).

SOC metrics

- **MTTD (Mean Time to Detect)** and **MTTR (Mean Time to Respond)**: lower is better.
- Alert volume per analyst, false positive rate, SLA compliance, escalation accuracy.

What an L1 analyst's shift looks like

Shift handover → review open tickets → work the alert queue by priority → triage, enrich, document → escalate true positives to L2 → handle phishing reports from users → write the shift summary → handover. Most SOCs run 24x7 rotating shifts, including nights.

5. VAPT: Vulnerability Assessment & Penetration Testing

VAPT is finding security weaknesses before attackers do. It is offensive security, and it is legal **only with written authorisation** from the system owner. Testing anything without permission is a crime in most countries, including under India's IT Act, 2000.

VA vs PT

	Vulnerability Assessment (VA)	Penetration Testing (PT)
Goal	Find and list as many weaknesses as possible	Prove which weaknesses can actually be exploited, and how far an attacker gets

	Vulnerability Assessment (VA)	Penetration Testing (PT)
Method	Mostly automated scanning, then validation	Manual, skilled testing supported by tools
Depth	Wide but shallow	Narrow but deep
Output	Prioritised list of vulnerabilities	Exploitation proof, attack paths and business impact
Frequency	Monthly or quarterly	Yearly, or after major changes

By knowledge given to the tester: black box (nothing, like an outside attacker), grey box (partial, such as a normal user login), white box (full access to source code and architecture).

By target: web application, API, mobile app, external network, internal network, wireless, cloud configuration, and social engineering (phishing simulations, physical entry tests).

Related activities: red teaming is a long, stealthy, goal-based simulation of a real attacker that also tests the SOC's detection; **purple teaming** is red and blue teams working together to improve detections; **bug bounty** programmes pay independent researchers for valid vulnerabilities, within published rules.

Penetration testing phases

1. **Planning & scoping:** agree in writing on targets, timings, allowed techniques, and emergency contacts (the **Rules of Engagement**), plus a signed authorisation. Without this, it is illegal hacking.
2. **Reconnaissance:** gather information. **Passive recon** uses public sources (company website, DNS records, certificate transparency logs, LinkedIn, search engines) without touching the target. **Active recon** interacts with the target.
3. **Scanning & enumeration:** find live hosts, open ports, running services and versions, users, shares, and web directories.
4. **Vulnerability analysis:** match findings against known weaknesses (CVE databases, scanner results) and spot misconfigurations.
5. **Exploitation:** safely prove a weakness can be used, without damaging the system.
6. **Post-exploitation:** see how far access goes: privilege escalation, lateral movement, access to sensitive data. Document everything.
7. **Reporting:** the main deliverable (see below).

8. **Remediation & retest:** the client fixes the issues and the tester confirms the fixes work.

Well-known methodologies: **PTES** (Penetration Testing Execution Standard), the **OWASP Web Security Testing Guide** for web apps, and **NIST SP 800-115**.

OWASP Top 10: web application risks

The OWASP Top 10 is the standard awareness list of critical web application risks. The 2021 edition is still the one most interviews ask about:

#	Risk	Simple example	Typical fix
A01	Broken Access Control	Changing `order_id=1001` to `1002` in a URL shows another customer's order (IDOR)	Server-side authorisation checks on every request
A02	Cryptographic Failures	Passwords stored in plain text; sensitive data sent over HTTP	Strong hashing, TLS everywhere
A03	Injection	Typing `' OR 1=1 --` into a login field bypasses it (SQL injection); XSS also falls here	Parameterised queries, input validation, output encoding
A04	Insecure Design	Password reset with no rate limit or verification	Threat modelling during design
A05	Security Misconfiguration	Default admin credentials; detailed error pages; open cloud storage	Hardening baselines, configuration reviews
A06	Vulnerable & Outdated Components	Running a library with a known critical CVE	Patching, dependency scanning
A07	Identification & Authentication Failures	No lockout or MFA; session IDs that never expire	MFA, lockouts, secure session handling
A08	Software & Data Integrity Failures	Installing unsigned updates; insecure deserialisation	Signed updates, integrity checks

#	Risk	Simple example	Typical fix
A09	Security Logging & Monitoring Failures	Attacks go unnoticed because nothing is logged	Log security events, send them to a SIEM
A10	Server-Side Request Forgery (SSRF)	Tricking the server into fetching internal or cloud-metadata URLs	Allow-lists for outbound requests

XSS (Cross-Site Scripting) injects a script into a page that runs in other users' browsers, e.g. to steal session cookies. The three types are stored, reflected and DOM-based. **CSRF (Cross-Site Request Forgery)** tricks a logged-in user's browser into sending an unwanted request.

OWASP has released a newer 2025 edition of the list; check owasp.org for the current version. For APIs, see the separate **OWASP API Security Top 10**.

Vulnerability scoring and management

- **CVE:** a unique public ID for a known vulnerability, e.g. CVE-2021-44228 (Log4Shell, in the Log4j logging library).
- **CWE:** a category of weakness, e.g. CWE-79 (XSS), CWE-89 (SQL injection).
- **CVSS:** a 0–10 severity score. 0.1–3.9 Low, 4.0–6.9 Medium, 7.0–8.9 High, 9.0–10.0 Critical. Version 3.1 is still widely used; version 4.0 was released in 2023.
- **EPSS:** estimates the probability a vulnerability will be exploited in the wild.
- **CISA KEV catalogue:** a list of vulnerabilities known to be actively exploited; these should be patched first.
- **Prioritise by context, not score alone:** a Critical bug on an internet-facing server beats a Critical bug on an isolated test machine.
- **Vulnerability management lifecycle:** discover assets → scan → prioritise → remediate → verify → report, repeated continuously.

Common tools

Tool	Used for
Nmap	Host discovery and port/service scanning (e.g. `nmap -sV -p- target`)

Tool	Used for
Nessus / OpenVAS / Qualys	Automated vulnerability scanning
Burp Suite	Intercepting and testing web application traffic (the industry standard)
OWASP ZAP	Free, open-source web application scanner
Gobuster / ffuf	Discovering hidden web directories and files
sqlmap	Automated SQL injection testing
Metasploit	Exploitation framework
Hydra / John the Ripper / Hashcat	Password attacks and hash cracking (in authorised tests only)
Wireshark	Capturing and analysing network traffic (used in SOC work too)
Kali Linux / Parrot OS	Linux distributions with these tools pre-installed

The VAPT report

A strong report contains:

- **Executive summary:** overall risk in plain business language, for management.
- **Scope and methodology:** what was tested, when, and how.
- **Findings:** for each issue, its title, severity (CVSS), affected systems, description, steps to reproduce, evidence screenshots, business impact, and a specific fix.
- **Remediation roadmap:** what to fix first.
- **Appendix:** tool output and raw data.

Clear writing matters as much as technical skill: a finding nobody understands never gets fixed.

Ethics

Stay strictly inside the agreed scope, report everything you find, never keep or share client data, and stop and inform the client immediately if you find evidence of a real attack in progress. Practise only on platforms built for it (section 7).

6. Cloud security

Most organisations now run on AWS, Microsoft Azure or Google Cloud, so every security role touches the cloud. Most cloud breaches come from customer misconfiguration, not from the provider being hacked.

Cloud models and shared responsibility

- **IaaS** (virtual machines, e.g. AWS EC2, Azure VMs): the customer secures the OS, apps, data and access.
- **PaaS** (managed platforms, e.g. Azure App Service, AWS RDS): the provider handles the OS; the customer secures the app, data and access.
- **SaaS** (finished software, e.g. Microsoft 365, Salesforce): the customer still owns their data, user accounts and configuration.

In every model the provider secures the physical data centres and core infrastructure, and the **customer is always responsible for their data and identities**.

Key cloud security areas

- **Identity and access (IAM):** identity is the new perimeter. Use MFA everywhere, least-privilege roles (RBAC), conditional access (e.g. block logins from unexpected countries), and just-in-time admin access. Protect the root or global admin account above all.
- **Network security:** security groups and network security groups (cloud firewall rules), private endpoints, no management ports (22, 3389) open to the internet.
- **Data protection:** encryption at rest and in transit, key management services (AWS KMS, Azure Key Vault), and never storing secrets in code.
- **Logging and monitoring:** turn on audit logs (AWS CloudTrail, Azure Activity Logs, Entra ID sign-in logs, Google Cloud Audit Logs) and send them to a SIEM.
- **Posture management (CSPM):** tools that continuously check for misconfigurations against benchmarks such as the CIS Benchmarks.

Most common cloud misconfigurations

Misconfiguration	Risk	Fix
Public storage bucket or blob container	Data leak to the whole internet	Block public access by default

Misconfiguration	Risk	Fix
Over-permissive IAM roles (e.g. admin for everyone)	One stolen account compromises everything	Least privilege, regular access reviews
No MFA on admin accounts	Account takeover	Enforce MFA
SSH or RDP open to 0.0.0.0/0	Brute force, ransomware	Restrict by IP, use a bastion or VPN
Access keys hard-coded in code or public repositories	Attackers find and use them within minutes	Secret managers, secret scanning
Logging disabled	Attacks go unseen; no evidence	Enable audit logs in every account and region

Native security tools by provider

Purpose	AWS	Microsoft Azure	Google Cloud
Identity	IAM, IAM Identity Center	Microsoft Entra ID	Cloud IAM
Audit logs	CloudTrail	Activity Logs, Entra sign-in logs	Cloud Audit Logs
Threat detection	GuardDuty	Microsoft Defender for Cloud	Security Command Center
Posture and findings	Security Hub	Defender for Cloud (secure score)	Security Command Center
SIEM	(third-party or Security Lake)	Microsoft Sentinel	Google Security Operations
Secrets and keys	Secrets Manager, KMS	Key Vault	Secret Manager, Cloud KMS

Beginner certifications and a hands-on project

- **Microsoft:** AZ-900 (cloud fundamentals), then SC-900 (Security, Compliance & Identity Fundamentals).
- **AWS:** Cloud Practitioner, then Security Specialty later in your career.

- **Project:** in a free-tier account, set a budget alert first, then enable MFA, create a least-privilege user, lock down a network rule, turn on audit logs and the native threat detection service, and write up each step with screenshots.

7. Roadmap, labs and interview prep

12-week plan (about 2 hours a day)

This plan targets an entry-level SOC role. For GRC, swap weeks 7–8 for section 3 plus an ISO 27001 course; for VAPT, extend weeks 11–12 into a full month of web labs.

- [ ] **Weeks 1–2:** Sections 1–2 (foundations, networking). TryHackMe Pre Security path. First LinkedIn post.
- [ ] **Weeks 3–4:** Linux and Windows. OverTheWire Bandit levels 0–15; TryHackMe Linux and Windows Fundamentals. Start ISC2 CC training.
- [ ] **Weeks 5–6:** Pass ISC2 CC. Start TryHackMe SOC Level 1.
- [ ] **Weeks 7–8:** Section 4 (SOC): SIEM practice, phishing analysis, LetsDefend alert queue. Publish 2 investigation write-ups.
- [ ] **Weeks 9–10:** Section 6 (cloud): SC-900 or AWS Cloud Practitioner, plus the cloud lab project.
- [ ] **Weeks 11–12:** Sections 3 and 5 (GRC, VAPT): PortSwigger labs for the OWASP Top 10. Mock interviews, polish your CV and LinkedIn, start applying.

Certifications for beginners

Certification	Cost	Best for
ISC2 Certified in Cybersecurity (CC)	Free training and exam under ISC2's "One Million Certified" offer (check that it's still running)	Broad foundation; a good first certification
Google Cybersecurity Certificate (Coursera)	Subscription; financial aid available	Structured course including SIEM, Linux and Python basics
Microsoft SC-900	Paid exam; Microsoft regularly runs free voucher events	Cloud and identity security, especially after AZ-900

Certification	Cost	Best for
CompTIA Security+	Paid; the most expensive here	The most widely recognised entry-level certification in job postings
Blue Team Level 1 (BTL1) or TryHackMe SOC certifications	Paid	Hands-on SOC skills
ISO/IEC 27001 Foundation	Paid	GRC roles
eJPT (INE)	Paid	Entry-level penetration testing

Start with the free certifications, and pay for one recognised exam (usually Security+) once you're sure of your direction.

Free and low-cost practice platforms

Platform	Best for
TryHackMe	Guided beginner paths: Pre Security, SOC Level 1, Cyber Security 101, Jr Penetration Tester
LetsDefend	A realistic SOC alert queue; the closest thing to a real L1 shift
Blue Team Labs Online / CyberDefenders	Blue-team investigation challenges: logs, packet captures, memory, phishing
OverTheWire Bandit	Linux command-line skills through a game
PortSwigger Web Security Academy	Free, high-quality web hacking labs for the OWASP Top 10
Hack The Box	Harder offensive labs for after the basics
Splunk Boss of the SOC (BOTS) datasets	Practising SIEM searches on realistic data
Professor Messer (YouTube)	Free Security+ video course
CyberChef	Free "Swiss Army knife" for decoding and analysing data

Build a small portfolio

Employers trust proof over claims. Three to five documented projects are enough:

1. **Home SOC lab:** a Windows VM sending logs to a free SIEM (Splunk Free, Elastic or Microsoft Sentinel's free trial). Simulate a brute-force attack and write the detection.
2. **Phishing analysis write-ups:** analyse real phishing samples (from public repositories or your own spam folder), documenting headers, links and verdicts.
3. **LetsDefend or TryHackMe investigations:** write up 2–3 completed cases in your own words.
4. **Cloud hardening project:** the cloud lab from section 6, with before-and-after screenshots.
5. **GRC sample:** a mock risk register or ISO 27001 gap analysis for a fictional small company.

Publish them on GitHub, a free blog or LinkedIn articles, and link them from your CV.

Glossary

Term	Meaning
APT	Advanced Persistent Threat: a skilled, long-term attacker group, often state-backed
C2 (Command & Control)	The channel attackers use to control infected machines
DLP	Data Loss Prevention: tools that stop sensitive data leaving the organisation
EDR / XDR	Endpoint (or extended) Detection & Response tools
Exfiltration	Stealing data out of a network
Hash	A fixed fingerprint of data (e.g. SHA-256), used to identify files
IOC	Indicator of Compromise: evidence of an attack, such as a bad IP or file hash
Lateral movement	An attacker moving from one compromised system to others
Persistence	Techniques an attacker uses to keep access after reboots or password changes
Privilege escalation	Gaining higher access than granted (user → admin)
Sandbox	An isolated environment for running suspicious files safely
TTP	Tactics, Techniques & Procedures: how an attacker operates
WAF	Web Application Firewall
Zero-day	A vulnerability exploited before a patch is available

20 common interview questions

Practise answering each out loud in 1–2 minutes.

1. Explain the CIA triad with an example of each.
2. What's the difference between a threat, a vulnerability and a risk?
3. What is MFA, and why isn't two passwords MFA?
4. Explain TCP vs UDP and the three-way handshake.
5. What happens when you type a website address into a browser?
6. Name common ports and their services.
7. What's the difference between encryption, hashing and encoding?
8. What is a SIEM and how does it work?
9. How do you tell a true positive from a false positive?
10. Walk through how you'd analyse a suspicious email.
11. You see 200 failed logins and then a success. What do you do?
12. Why is Word launching PowerShell suspicious?
13. What are the phases of incident response?
14. What is MITRE ATT&CK, and how is it used?
15. Explain SQL injection and XSS, and how to prevent them.
16. What's the difference between a vulnerability assessment and a penetration test?
17. What is the shared responsibility model in the cloud?
18. What is ISO 27001? Name a data protection law relevant to your country.
19. What's a recent cyber attack in the news, and what could have prevented it?
20. Why cybersecurity, and how does your background help? (Connect concrete past tasks to security skills, then name your labs and certifications.)

Staying current

Security changes every week. Follow The Hacker News, BleepingComputer, Krebs on Security, CERT-In advisories, and the Darknet Diaries podcast, and spend 15 minutes a day reading. Interviewers often ask about recent incidents.