CyberPrompts

AI prompt sheet for defensive and authorised cybersecurity work

12 categories and 50 ready-to-use prompt templates. Each one has a use case, a fill-in template and a pro tip.

> **Use these only on systems, data and people you are authorised to assess.** Unauthorised access to a computer system can create civil and criminal liability in India (see the companion sheet, Cyber Law India). Every prompt here is built for defence, analysis, learning and reporting.

CyberWithMadhav | Madhav Pandya

By Madhav Pandya, CyberWithMadhav. Edition: November 2026. AI output can be wrong: verify before you act.

How to use this sheet

The scaffold behind every prompt

Part	What to include
Role	Who the model should act as, for example a Tier-1 SOC analyst or a secure-code reviewer.
Context	Environment, asset criticality, jurisdiction, audience, and your authorisation.
Task	One clear job per prompt. Split big jobs into steps.
Inputs	The redacted data to analyse, pasted between clear markers.
Constraints	What it must not do: no exploitation steps, no invented facts, no personal data.
Output format	Table, checklist, or JSON, plus length and reading level.
Check	Ask it to label Observed versus Inferred and to say where it is unsure.

Before you paste anything

- **Authorisation first.** Written scope for scanning and testing; consent or lawful basis for personal data.
- **Redact.** Remove credentials, tokens, customer data, and personal data you do not need. Check your organisation's policy and the AI provider's terms before sharing any client material.
- **Verify.** Models can invent CVEs, log fields, IOCs, and legal section numbers. Confirm against primary sources: NVD, vendor advisories, MITRE ATT&CK;, the bare Act.
- **Keep a human decision-maker.** AI drafts and explains; you approve containment, reports and legal wording.
- **Log your prompts** with the engagement notes so results are reproducible and reviewable.

Reading the templates

Marking	Meaning
`[SCOPE]`	Replace with your own details. Square brackets in purple are always yours to fill.
`[A/B/C]`	Choose one option.
`[LIST]`	Paste several items, one per line.

> **Tip for better answers:** ask for evidence quotes, a confidence level on each claim, and what would change the conclusion. Then iterate: paste the model's table back and ask it to challenge its own weakest point.

CyberWithMadhav | Madhav Pandya

Contents

Category	Prompts
OSINT (authorised and ethical)	4
Recon and scanning (owned or authorised targets)	5
SOC and log triage	5
Incident response	4
Threat intelligence summaries	4
Malware analysis (static and report reading)	4
Phishing analysis	4
Cloud security review	4
Secure code review	4
GRC and policy writing	4
CTF learning	4
Report writing	4

CyberWithMadhav | Madhav Pandya

OSINT (authorised and ethical)

Passive, lawful open-source research. Define scope first, grade every source, verify before you report, and minimise personal data.

Scope and legality check

Run this before any OSINT task starts.

```
Act as an OSINT lead bound by [JURISDICTION] law and our rules of engagement. Engagement: [CLIENT_OR_PURPOSE]. Written authorisation: [YES/NO, REFERENCE]. Subject type: [ORGANISATION/DOMAIN/CONSENTING_PERSON]. List: (1) permissible passive sources, (2) activities that would amount to intrusion or a privacy violation, (3) data I must not collect or must minimise, (4) a go/no-go checklist. Ask me clarifying questions before you advise.
```

Pro tip Save the answer in the engagement file. If the model hesitates on a source, treat that as a prompt to ask your lawyer.

Passive footprint summary

Turn raw passive findings into a ranked external risk view.

```
From these passive-only findings about [ORG]: domains and subdomains [LIST], certificate transparency entries [EXCERPT], public repositories [LIST], job-post snippets [SNIPPETS]. Build a footprint summary: technologies, likely exposed services, third-party dependencies, and the 5 highest risks. Label every item Observed or Inferred and name the input that supports it.
```

Pro tip The Observed/Inferred label is what stops the model inventing assets.

Source reliability grading

Grade claims before they reach a report.

```
Here are claims about [SUBJECT] with their sources: [CLAIMS_AND_URLS]. For each claim give source reliability (A to F) and information credibility (1 to 6) using the Admiralty system, flag contradictions, and suggest one independent way to corroborate it. Output a table.
```

Pro tip Ask for the single weakest claim in the set and why. It is usually the one you were about to rely on.

OSINT brief with privacy minimisation

Convert notes into an executive brief that does not over-collect.

```
Turn these raw notes [NOTES] into a brief for [AUDIENCE]: scope, passive-only method, key findings with confidence levels, limitations, and a data-handling appendix. Redact any personal data not needed for [PURPOSE] and list what you removed.
```

Pro tip Redact by default and add data back only when you can state why it is needed.

Recon and scanning (owned or authorised targets)

Interpreting scanner output and planning safe scans. These prompts explain results and remediation, not exploitation.

Nmap output triage

Make sense of a service scan and decide what to fix first.

```
I hold written authorisation to scan [SCOPE]. Interpret this nmap output: [NMAP_OUTPUT]. Group by host and list exposed services, version-related risks (say 'verify' unless you are confident), and a remediation order for the defender. Do not suggest exploitation steps.
```

Pro tip Cross-check every CVE the model names against NVD or the vendor advisory before it goes in a report.

Nuclei findings de-duplication

Cut noisy template results down to real risk.

```
Authorisation: I hold written authorisation to scan [SCOPE]. Without it, scanning can engage IT Act s.43(a) (civil) and s.66 (criminal), and the target may have to report it as targeted scanning or probing (CERT-In Annexure I(i)). Here are nuclei results: [NUCLEI_OUTPUT]. De-duplicate, separate genuine risk from informational findings, map each to its asset and template ID, estimate exploitability in our context [ENVIRONMENT_NOTES], and propose fixes with a suggested owner. Finish with the top 5 actions.
```

Pro tip Paste an asset criticality list as well. Prioritisation improves sharply with business context.

Scan drift comparison

Spot new, removed and changed services between two scans.

```
Compare scan A ([DATE_A]) with scan B ([DATE_B]) for [SCOPE]: [SCAN_A] and [SCAN_B]. Report new, removed and changed ports and services, possible shadow IT, and which changes need approval versus investigation.
```

Pro tip Sort and normalise both outputs (same flags, same order) before pasting.

Scan plan reviewer

Catch production risk and authorisation gaps before you scan.

```
Review my scan plan for [CLIENT]: targets [LIST], tools [LIST], window [TIME_WINDOW], rate limits [LIMITS], exclusions [LIST], authorisation reference [REF]. Identify risks to production (service disruption, account lockouts), gaps in authorisation, and logging I should keep. Suggest safer settings.
```

Pro tip Have a second person sign off on scope and window before the first packet is sent.

Web headers and TLS review

Turn header and TLS scan output into a fix list.

```
Authorisation: I hold written authorisation to test [DOMAIN]. Without it, testing can engage IT Act s.43(a) (civil) and s.66 (criminal), and the target may have to report it as targeted scanning or probing (CERT-In Annexure I(i)). Interpret these HTTP response headers and TLS results for [DOMAIN]: [PASTE]. List missing or weak controls (HSTS, CSP, cookie flags, protocol versions, ciphers) in priority order, and give a sample secure configuration for [SERVER_OR_CDN].
```

Pro tip Test the suggested configuration in staging. Models sometimes produce syntax for the wrong server version.

SOC and log triage

Faster, more consistent alert handling. Always pass asset context and the log schema you actually have.

Alert triage assistant

First-pass assessment of a single alert.

```
Act as a Tier-1 SOC analyst. Alert: [ALERT_NAME_AND_RULE]. Raw event: [LOG_SNIPPET]. Asset: [ASSET_AND_CRITICALITY]. User: [USER_ROLE]. Give: a likely true or false positive assessment with reasoning, 5 pivot queries for [SPLUNK/ELASTIC/SENTINEL], the enrichment I need, and an escalate or close recommendation with rationale.
```

Pro tip Always state asset criticality. The same alert deserves different handling on a domain controller and a lab VM.

Log explainer

Understand unfamiliar logs in time order.

```
Explain these logs [REDACTED_LOGS]: what happened in sequence, which parts look normal and which look suspicious, and which MITRE ATT&CK techniques might apply (mark each as 'possible'). Quote the log line that supports each claim.
```

Pro tip Strip secrets and personal data first. Requiring a quoted line per claim exposes guesses.

Detection query builder

Draft a detection and its test cases.

```
Write a [KQL/SPL/SIGMA] detection for [BEHAVIOUR] using log source [SOURCE] and these fields only: [FIELD_LIST]. Include false-positive considerations, tuning options, and test cases (one benign, one malicious).
```

Pro tip Give the real schema. Models invent field names unless you list them.

Authentication anomaly review

Find brute force, spraying and suspicious successes.

```
Analyse this authentication log extract. I have pseudonymised usernames and IP addresses (USER_01, IP_A and so on) and keep the mapping offline: [LOG_DATA]. Report top source aliases, targeted accounts, success-after-failure sequences, signs of password spraying and impossible travel, and next containment steps for [ENVIRONMENT]. Expected user locations: [GEOS]. Timezone: [TZ].
```

Pro tip Pseudonymise consistently so patterns survive, and give a coarse location label per IP alias (for example, IP_A is an Indian mobile ISP) so travel analysis still works.

Shift handover summary

Clean handovers between analysts.

```
Summarise this shift's tickets [TICKET_LIST] into a handover: open incidents, risk level, pending actions with owners and deadlines, and what the next analyst should check first.
```

Pro tip Ask for a 'known unknowns' line so the next shift knows what was not checked.

Incident response

Structure under pressure. Keep facts separate from inferences and keep a decision log.

First-four-hours checklist

Immediate priorities for a suspected incident.

```
We suspect [RANSOMWARE/BEC/DATA_EXFILTRATION/OTHER] in [ENVIRONMENT]. Known facts: [FACTS]. Unknowns: [UNKNOWNS]. Produce a first-four-hours checklist: containment options with business impact, evidence to preserve (volatile first), who to notify (legal, management, regulators), and decision points.
```

Pro tip In India, check early whether CERT-In's 6-hour reporting direction applies. Confirm with counsel. The clock runs from when you notice the incident or are told of it.

Incident timeline builder

One normalised timeline from many sources.

```
Build a consolidated incident timeline from these sources: [LOGS_AND_NOTES_WITH_TIMESTAMPS]. Normalise to UTC and IST, flag gaps and conflicts between sources, and label each entry Fact or Inference.
```

Pro tip Always say which timezone each source uses. Mixed timezones cause most timeline errors.

Containment option analysis

Compare containment choices before acting.

```
Compare containment options [OPTIONS] for [SYSTEMS] on speed, evidence loss, business impact and reversibility. Constraints: [CONSTRAINTS]. Recommend a sequence and say what could go wrong at each step.
```

Pro tip Ask for the option that preserves the most evidence, then decide whether the trade-off is acceptable.

Stakeholder updates

Plain updates without speculation.

```
Draft a [EXECUTIVE/TECHNICAL/CUSTOMER] update for incident [ID]. Status: [STATUS]. Impact: [IMPACT]. Actions taken: [ACTIONS]. Next update at [TIME]. Use plain language, avoid speculation, and list open unknowns explicitly.
```

Pro tip Have legal or the incident commander approve any external wording before it goes out.

Threat intelligence summaries

Condense reports into briefings people will act on. Keep indicators separate from analysis.

Report-to-brief summary

Summarise a long threat report for a specific audience.

```
Summarise this threat report [REPORT_EXCERPTS] for [CISO/SOC/IT_TEAM]: actor or campaign, TTPs (give ATT&CK IDs only if the source states them), an IOC table kept separate from analysis, relevance to our stack [STACK], and recommended actions. Use your own words and cite the section you drew from.
```

Pro tip Tell the model to say 'not stated in the report' rather than fill gaps.

CyberWithMadhav Madhav Pandya

IOC extraction and defanging

Pull indicators out of text safely.

```
Extract every indicator of compromise from [TEXT] into a table with type, value, context and confidence. Defang URLs, domains and IPs. Do not add any indicator that is not present in the text.
```

Pro tip Spot-check ten rows against the source. Extraction errors are usually small and easy to miss.

Relevance assessment

Decide whether a threat matters to you.

```
Given our stack [STACK], sector [SECTOR] and location [GEOGRAPHY], assess this threat [THREAT_SUMMARY]: likelihood, impact, gaps in existing controls, and 3 prioritised mitigations. State your assumptions.
```

Pro tip List your existing controls so the model does not recommend what you already have.

Weekly threat digest

A one-page digest from several advisories.

```
Combine these advisories [ADVISORIES] into a one-page weekly digest: top 3 priorities, a patch and mitigation list with suggested deadlines, and an ignore list with reasons.
```

Pro tip Link every item to its advisory so readers can check the primary source.

Malware analysis (static and report reading)

Analysis in an isolated lab and reading existing reports. These prompts never ask for malware to be written, improved or run.

Sandbox report interpreter

Read a sandbox or AV report quickly.

```
Interpret this sandbox report [REPORT]. Summarise behaviour (file, registry, network, process), separate observed activity from heuristic detections, list indicators, and state the limits of this analysis. Do not provide instructions to build or improve malware.
```

Pro tip Ask which conclusions would change if the sample detected the sandbox.

Static strings and imports review

Form hypotheses from static artefacts.

```
Here are strings, imports and section data from a suspicious [PE/ELF/OFFICE] file, analysed in an isolated lab: [STATIC_DATA]. Suggest hypotheses about its function ranked by confidence, and which further static checks would confirm or refute each.
```

Pro tip Treat the output as hypotheses to test, not conclusions.

Report to detection ideas

Turn a published analysis into defensive detections.

```
From this vendor malware analysis [EXCERPT], derive defensive detections: logic for YARA or Sigma rules described in prose, network indicators, and EDR telemetry to hunt. Flag anything likely to cause false positives.
```

Pro tip Test any rule against a clean corpus before deploying.

CyberWithMadhav | Madhav Pandya

Obfuscated script explanation

Understand a script after safe static decoding.

```
Explain what this script [SCRIPT] likely does after the safe static de-obfuscation steps I applied: [STEPS_TAKEN]. I have defanged every URL, domain and IP address in it. List decoded strings, suspected network or payload behaviour, and safe next steps. Do not execute or recreate it.
```

Pro tip Decode in an offline lab first and defang URLs, domains and IPs before pasting. Never paste live malware into a public AI tool.

Phishing analysis

Analyse messages, respond to clicks, and train people. Nothing here helps craft phishing.

Email header and authentication analysis

Verify where a message really came from.

```
Analyse these email headers [HEADERS]: SPF, DKIM and DMARC results, relay path anomalies, mismatches between From, Return-Path and Reply-To, and a verdict with confidence. Quote the header lines that support each point.
```

Pro tip Remove recipient addresses you do not need before pasting.

Phish content triage

Explain the manipulation in plain language.

```
Review this suspicious message [REDACTED_BODY]. All URLs and domains in it are defanged (hxxps and dots in square brackets). Identify the social-engineering tactics, urgency and authority cues, URL and attachment risks, and write a short user-facing explanation in [ENGLISH/HINDI/MARATHI].
```

Pro tip Defang URLs before pasting (hxxps://example[.]com) so nobody clicks one by accident. Reuse the plain-language explanation as the reply to the person who reported it.

Clicked-link response plan

Contain a likely credential compromise.

```
A user in [ROLE] opened [DEFANGED_URL] and may have entered credentials. Provide containment steps (password reset, session revocation, mailbox rule and forwarding checks, EDR hunting), questions to ask the user, and notification triggers for [ORG].
```

Pro tip Paste the URL defanged (hxxps and [.] in the domain). Check for inbox rules and OAuth app consents, which attackers use to persist after a password reset.

Awareness micro-lesson

Teach from a real, de-identified example.

```
Write a 10-minute awareness lesson based on this de-identified phish, with any URLs defanged: [TEXT]: red flags, what to do, how to report at [ORG], and 5 quiz questions with answers.
```

Pro tip Use it only in an authorised awareness programme with management approval.

Cloud security review

Configuration and logs, reviewed with a least-privilege mindset.

IAM policy reviewer

Find over-permissive access and propose a tighter version.

```
Review this IAM or role policy [POLICY_JSON]. Identify over-permissive statements, privilege-escalation paths and missing conditions, then propose a least-privilege version and explain each change. Environment: [AWS/AZURE/GCP].
```

Pro tip Supply the real resource names. Wildcards are where most findings hide.

Infrastructure-as-code review

Check Terraform or CloudFormation before deployment.

```
Review this [TERRAFORM/CLOUDFORMATION] code [CODE] for public exposure, missing encryption, logging gaps, network rules and hard-coded secrets. Map findings to [CIS_BENCHMARK/NIST] controls and provide corrected snippets.
```

Pro tip Run the fixed code through a linter such as a policy-as-code tool before trusting it.

Architecture threat model

A structured STRIDE pass on a design.

```
Threat-model this cloud architecture [DESCRIPTION] using STRIDE. List assets, trust boundaries, the top 10 threats, controls present versus missing, and a prioritised roadmap. State your assumptions.
```

Pro tip Include data flows and who administers what. Threat models fail on missing flows.

Cloud activity log review

Look for suspicious control-plane activity.

```
Review these [CLOUDTRAIL/ACTIVITY_LOG] events [EVENTS]: unusual API calls, new access keys, region anomalies, console sign-ins without MFA, and changes to logging. Suggest containment steps for [ACCOUNT_TYPE].
```

Pro tip Give the normal regions and admin roles so anomalies are judged against your baseline.

Secure code review

Use AI as a second reviewer. It finds candidates; you confirm them.

Language-aware security review

Review a function or file for common flaws.

```
Review this [LANGUAGE] code [CODE] for security issues: injection, authentication and authorisation flaws, insecure deserialisation, SSRF, secrets, crypto misuse and race conditions. For each issue give the line, risk, CWE, and a fix. Say 'not sure' where you are uncertain.
```

Pro tip Ask for a proof-of-concept test, not an exploit, so you can confirm the issue in your own test suite.

CyberWithMadhav | Madhav Pandya

Authentication and session review

Check a login design against a standard.

```
Here is our login and session design [DESCRIPTION_OR_CODE]. Check it against OWASP ASVS level [LEVEL]: credential storage, session handling, MFA, rate limiting, logout and account recovery. List gaps in priority order.
```

Pro tip Recovery and reset flows are the most common weak point. Ask for them specifically.

Dependency risk triage

Prioritise a dependency or SBOM list.

```
From this dependency list or SBOM [LIST], flag high-risk packages (unmaintained, typosquat-prone, known-vulnerable only if you are certain), and suggest upgrade paths and mitigations. Tell me which to confirm in [OSV/NVD/GITHUB_ADVISORIES].
```

Pro tip Treat vulnerability claims as leads and verify each in an advisory database.

Fix with regression tests

Patch a confirmed issue and prove it stays fixed.

```
Fix this confirmed vulnerability [DESCRIPTION] in [CODE] with a minimal patch. Explain why it works and add regression tests, including negative tests that should fail on the vulnerable version.
```

Pro tip Run the tests against the unpatched code first to confirm they actually fail.

GRC and policy writing

Policies, control mappings and risk entries that an auditor can follow.

Policy drafter

A first draft of a security policy.

```
Draft an [INFORMATION_SECURITY/ACCESS_CONTROL/INCIDENT_RESPONSE] policy for [ORG_NAME, SIZE, SECTOR] aligned to [ISO_27001_2022/NIST_CSF_2.0/DPDP]. Include purpose, scope, roles, numbered requirements, exceptions process and review cycle. Plain language, about [LENGTH] words.
```

Pro tip A policy states what must happen. Ask for a separate procedure document for how.

Control mapping and gap view

Map your controls to a framework.

```
Map our controls [CONTROL_LIST] to [ISO_27001_ANNEX_A/SOC_2/NIST_CSF_2.0]. Show covered controls, gaps, and suggested evidence for each. Mark any mapping you are unsure of.
```

Pro tip Spot-check the mapping with the framework text. Control numbers are easy to get wrong.

Risk register entries

Turn an issue into proper risk statements.

```
Convert this issue [DESCRIPTION] into risk register entries: risk statement, asset, threat, vulnerability, likelihood and impact with rationale, existing controls, treatment options and a proposed owner.
```

Pro tip Ask for two scenarios (likely and worst-case) rather than a single number.

CyberWithMadhav | Madhav Pandya

DPDP readiness checklist (India)

Prepare for the Digital Personal Data Protection regime.

```
Create a DPDP Act 2023 readiness checklist for [ORG_TYPE] covering notice, consent, retention and erasure, security safeguards, breach response, children's data and vendor contracts. Mark each item to verify against the Act and Rules, and note the phased commencement dates.
```

Pro tip Main duties are scheduled for about May 2027, but dates can change. Verify current status before you rely on it.

CTF learning

For legal CTFs and your own lab only. Built to teach, not to hand over flags.

Hint ladder

Get unstuck without spoilers.

```
I have checked that this event allows AI assistance. I am stuck on a CTF challenge. Category: [CATEGORY]. Description: [TEXT]. What I tried: [ATTEMPTS]. Give 3 progressive hints, from a gentle nudge to a near-spoiler, and stop. Do not reveal the flag unless I ask.
```

Pro tip Check the event's AI rules first. Some events ban AI help, and breaking that rule can get you disqualified. Then try each hint for 15 minutes before asking for the next one.

Concept explainer

Build a solid mental model of a topic.

```
Explain [CONCEPT] with a mental model, one analogy, common mistakes, how defenders detect and prevent it, and 3 practice labs on [PLATFORM]. Assume I know [CURRENT_LEVEL].
```

Pro tip Ask for the defensive view too. It makes you a better tester and a better analyst.

Write-up coach

Turn notes into a clear write-up.

```
Here are my notes for challenge [NAME]: [NOTES]. Turn them into a write-up: context, approach, dead ends, solution, lessons learned, and how the weakness would be mitigated in a real system.
```

Pro tip Publishing write-ups is great portfolio material. Check the CTF's rules on sharing solutions first.

Study plan

A realistic plan toward a goal.

```
Make a [NUMBER]-week plan to reach [GOAL, e.g. eJPT/SECURITY_PLUS/BLUE_TEAM_ROLE] with [HOURS] hours per week. Current level: [LEVEL]. Include weekly goals, labs, review quizzes and checkpoints.
```

Pro tip Ask the model to build in one rest week and one catch-up week.

CyberWithMadhav | Madhav Pandya

Report writing

Clear, defensible reports. Evidence first, adjectives last.

Finding write-up

A consistent write-up for one finding.

```
Write a finding: title, severity [RATING_OR_CVSS_VECTOR], affected assets [LIST], description, evidence [EVIDENCE_REFERENCES], business impact [CONTEXT], remediation steps, and references. Audience: [DEVELOPERS/EXECUTIVES].
```

Pro tip Provide the evidence yourself. Never let the model invent screenshots, versions or CVSS scores.

Executive summary

One page for decision-makers.

```
From these findings [FINDINGS_TABLE], write a one-page executive summary: overall risk posture, top 3 themes, business impact, and a 30/60/90-day roadmap. Avoid jargon.
```

Pro tip End with the three decisions leadership must make.

Retest plan

Verify fixes properly.

```
Create a retest plan for findings [LIST]: steps to retest, expected result when fixed versus not fixed, and the evidence to capture for each.
```

Pro tip Keep the same test method as the original finding so results are comparable.

Plain-language rewrite

Explain a finding to a non-technical owner.

```
Rewrite this technical finding [TEXT] in simple [ENGLISH/HINDI/MARATHI] for the owner of a [BUSINESS_TYPE], keeping it technically accurate. Add a glossary of 5 terms and one next step they can take today.
```

Pro tip Have a fluent speaker review translations before sending them to the client.

Responsible use

These prompts support learning, defence and professional reporting. They are not a licence to test anything you like. Work inside your written scope, protect personal data, and report vulnerabilities responsibly to the owner.

Legal note (India). Accessing a computer or network without the owner's permission can attract civil compensation under section 43 of the Information Technology Act, 2000, and criminal liability under section 66 where the act is dishonest or fraudulent. This is general information, not legal advice. See the companion sheet, Cyber Law India, for the detail and for points to verify.

Victim of cyber fraud? Report it quickly at cybercrime.gov.in or call the national cybercrime helpline 1930. Speed matters for freezing money, so do not wait to gather every detail first.

CyberWithMadhav | Madhav Pandya. Cybersecurity awareness and learning content. Feedback and corrections are welcome.