



Writing for Security Work


Professional track | Incident reports, findings, executive summaries, tickets, runbooks and playbooks, post-incident reviews: templates, plain-language rules, evidence and timestamps, severity and certainty wording, audience adaptation

#TechnicalWriting #IncidentReporting #Runbooks #Postmortem #Communication

Why writing is a security control

Intuition: the document is the handover

Most security work reaches other people only as text: the night analyst's ticket is all the day shift knows, the incident report is all the regulator sees, the finding is all the developer reads before fixing. A vague sentence costs hours (someone re-investigates), a wrong timestamp breaks a timeline, an overstated severity burns credibility for the next real alert. Good security writing is **short, verifiable, dated and aimed at one reader**, and it separates **what was observed** from **what is concluded** and **what must be done**.



BLUF — Bottom Line Up Front: military writing rule, put the conclusion and the ask in the first two lines, details after.


Executive summary — Half a page that lets a decision-maker act without reading the rest: situation, impact, what is being done, decision needed.

Runbook — Step-by-step procedure for one known technical task (rotate a key, isolate a host), exact commands and expected output.

Playbook — Scenario-level plan (ransomware, phishing, data leak): decisions, roles, communications and which runbooks to call.

Post-incident review (PIR) — Structured, blameless write-up after an incident: timeline, causes, what worked, actions with owners and dates; also called postmortem or lessons learned.

TLP (Traffic Light Protocol) — FIRST sharing label, version 2.0 (2022): TLP:RED (named recipients only), TLP:AMBER (organisation and clients need-to-know; AMBER+STRICT = organisation only), TLP:GREEN (community), TLP:CLEAR (public).




Plain-language rules that survive a crisis

Ten rules (checked = apply to every document)

✓ baseline / mandatory □ recommended – verify in your context

- ✓ One idea per sentence, 15-20 words on average; one topic per paragraph
- ✓ Active voice with a named actor: 'the attacker ran', 'we isolated', not 'was isolated' passive voice hides who did it, the first question an investigator asks
- ✓ Numbers, not adjectives: '3 400 records', not 'a large number'
- ✓ Every time in UTC, ISO 8601: 2026-03-14T02:17:09Z
- ✓ Separate fact, assessment and recommendation, and label them
- ✓ Expand every acronym on first use; readers change jobs, tools and languages
- ✓ Name hosts, accounts and tickets exactly as the systems do (SRV-FIN-02, svc_backup)
- **MUST** State confidence with calibrated words (see certainty scale)
- **SHOULD** Prefer lists and tables for steps and comparisons
- **MUST** No blame, no sarcasm, no jokes: tickets and chats become evidence in disputes and courts



The same update written twice: the right side can be acted on without a phone call

At ~~around 3am~~ 03:12Z ¹ the ~~server~~ host WEB-07 ² was ~~hacked~~ accessed with the stolen account j.doe ³. ~~It was blocked~~ The SOC analyst disabled j.doe at 03:20Z ⁴ and ~~everything is fine now~~ no further logons were seen until 06:00Z ⁵.

COMMENTS

- ¹ UTC time, not a guess
- ² Exact host name
- ³ Say what happened, not a loaded word
- ⁴ Named actor, active voice, time
- ⁵ Observed fact instead of reassurance

Editing a shift handover: each numbered edit applies one plain-language rule

✗ Writing your investigation diary instead of the result

New analysts narrate ('first I opened Splunk, then I searched...'). The reader wants the result first: what is true, how sure you are, what to do. Keep the chronological diary as working notes or an appendix; it is still valuable evidence of what you checked.

🕒 Evidence and timestamps: make every claim checkable

Mechanism. A claim in a report is only as strong as the path back to its source. For each fact, the reader should be able to find **where** it came from (log source, host, query), **when** (UTC timestamp and the clock it came from) and **who** collected it. Tools disagree on time: Windows event logs store UTC and display local time, many Linux syslog lines carry local time without a zone, cloud audit logs (AWS CloudTrail, Microsoft Entra ID) use UTC, and a drifting clock without NTP (Network Time Protocol) can be minutes off.

Rule	Write	Why
Normalise to UTC	2026-03-14T02:17:09Z (Z = UTC)	offices, clouds and logs in different zones
Keep the original too	'02:17:09Z (03:17:09 CET on host)'	lets anyone re-check the conversion
Name the clock	'EDR console time', 'firewall log time'	sources drift; say which one you trust
Hash exported evidence	SHA-256 of each file in the report	proves the file was not changed later
Quote, do not paraphrase	exact log line in a code block	paraphrase loses fields and invites dispute
Record the query	search string, time range, tool version	result can be reproduced or challenged
Redact, not delete	mask passwords and personal data	report is shared wider than the evidence

Time	Level	Source	Message
02:16:58Z	INFO	VPN	login j.doe from 185.220.101.4 (MFA push approved)
02:17:09Z	WARN	EDR SRV-FIN-02	rclone.exe started by svc_backup ▶ claim 1: exfiltration tool executed
02:17:40Z	ERROR	Proxy	12.4 GB upload to mega.nz blocked
02:31:02Z	AUDIT SUCCESS	EDR	host isolated by analyst m.leroy

Quote evidence as records: time (UTC), source, level, exact message; mark the line your claim rests on

🕒 See the time-zone trap on your own machine

Run `date`; `date -u`; `date -u +%Y-%m-%dT%H:%M:%SZ` (macOS or Linux) or `Get-Date`; `(Get-Date).ToUniversalTime().ToString('s')+'Z'` (PowerShell). Expected: two different hours unless you live in UTC, and one ISO 8601 string ending in Z. Now open a log you have (`/var/log/system.log`, Event Viewer) and decide whether its times are local or UTC: if you cannot tell, neither can your reader. Finally hash a file you would export: `shasum -a 256 file` (macOS) or `sha256sum file` (Linux) and note the value with the UTC time.

⚠️ Severity and certainty wording

Two separate scales appear in almost every document: **severity** (how bad the impact is or could be) and **confidence** (how sure you are). Mixing them produces sentences like 'a critical possible breach' that nobody can act on. State both, and define the severity words once in the document or point to the organisation's scale.

Word	Use it when (example definition, adapt locally)	Typical response time
critical	confirmed compromise of crown-jewel systems, active data theft or encryption, safety impact	immediately, 24/7
high	likely compromise or exploitable flaw on an exposed system with sensitive data	same day
medium	flaw needing extra conditions, or limited scope incident	days
low	hardening gap, no direct path to impact	next planned cycle

Certainty term	Probability band	Example sentence
almost no chance	1-5 %	'We see almost no chance that backups were reached.'
very unlikely	5-20 %	
unlikely	20-45 %	
roughly even chance	45-55 %	
likely	55-80 %	'The account was likely phished (MFA fatigue pattern).'
very likely	80-95 %	
almost certainly	95-99 %	'Data was almost certainly exfiltrated: 12 GB left.'

Where the certainty scale comes from

The bands above are the US intelligence community standard ICD 203 (Intelligence Community Directive 203, Analytic Standards, 2015). Say 'confirmed' only for what you observed directly; everything else gets a band word. Separately, give a confidence level (low, moderate, high) in the evidence behind the judgement.

Severity vs priority

Severity describes impact; **priority** is the order of work set by the business (a medium on the payment site may outrank a high on a lab server). CVSS (Common Vulnerability Scoring System) gives a technical severity score 0-10, not a business risk: write both when they differ and explain why.

Document templates: incident report, finding, executive summary

Incident report INC-2026-014

Classification: TLP:AMBER Version: 1.2

Author: _____ Status: contained

1. Summary (BLUF) *

Phished VPN account reached 2 finance servers; 12 GB upload blocked; no data loss assessed.

2. Scope and impact

Systems: _____

Data and volumes: _____

Regulatory triggers: _____

3. Timeline (UTC) *

4. Root cause and attack path

5. Actions taken and open

Owner: _____

Due: _____

Ticket: _____

6. Appendices

TLP:AMBER - INTERNAL DISTRIBUTION

1 5 lines: what, when, impact, status, decision needed

2 GDPR 72 h, NIS2 24 h early warning if met

3 time, source, event, confidence

4 label confirmed vs assessed

5 one owner and a date per action

6 IOCs (indicators of compromise), queries, hashes

Incident report skeleton (internal, after containment); margin notes say what each part must hold

HIGH 7.1

F-03

Any customer can read other customers' invoices (IDOR)

Asset api.shop.example/v2/invoices/{id}
Vector Authenticated customer, changes numeric id
Impact names, addresses, amounts of all 180 000 customers; GDPR breach
Fix check invoice.owner == session.user on the server; use random ids as defence in depth
Refs OWASP API1:2023 Broken Object Level Authorization

GET /v2/invoices/10442 as user A -> 200, invoice of user B

A pentest or audit finding: title states the flaw and the consequence; evidence proves it; fix is actionable

[TLP:AMBER] Incident 2026-014: finance servers, data theft attempt blocked - decision by 10:00

2026-03-14 08:05 UTC

From: CISO <ciso@corp.example>
To: exec-committee@corp.example

Bottom line: an attacker used a phished VPN account to reach two finance servers at 02:17 UTC and tried to upload 12 GB to a cloud storage site. The upload was blocked; we assess it very unlikely that data left. Payroll runs normally.

Done: servers isolated, account disabled, forensic copy taken.

Decision needed by 10:00: approve a forced password reset for all 900 VPN users (1 h of help-desk load).

Next update: 12:00 UTC or earlier if the assessment changes.

Executive summary sent at T+6 h: BLUF, impact in business terms, confidence, ask, next update

Subject lines that route themselves

Put the label, the incident id, the one-line status and the deadline in the subject: [TLP:AMBER] INC-2026-014 contained - decision by 10:00. People triage by subject line on a phone.

Tickets, runbooks, playbooks and post-incident reviews

ticket SOC-48213 (triage note) text

1 STATUS: escalated to IR, next step = memory capture WEB-07

2 SUMMARY: j.doe logged on to WEB-07 at 03:12Z from TOR exit

3 FACTS: Entra ID sign-in log id 8c1e..; MFA push approved 03:11Z

4 user on leave since 2026-03-10 (HR system)

5 RULED OUT: VPN split-tunnel false positive (IP not ours)

6 ASKS: IR to confirm scope; IAM to revoke j.doe sessions

7 EVIDENCE: SOC-48213/attachments, sha256 listed in note 2

A ticket another shift can continue: status and next step first, facts with sources, what was ruled out

Tadata | tadata.fr | contact@tadata.codes

Runbook	Playbook
One task: 'rotate the krbtgt password twice'	One scenario: 'ransomware on a file server'
Exact commands, expected output, rollback	Decisions, roles, communications, escalation
Written for the operator at the keyboard	Written for the incident lead and managers
Tested by running it in staging	Tested by tabletop exercise
Calls nothing; often automated in SOAR	Calls several runbooks at each phase

Runbook step format (one step = one verifiable action)

STEP 4 Disable the compromised account

WHY stops new logons while sessions are revoked

RUN Disable-ADAccount -Identity j.doe

EXPECT Get-ADUser j.doe | select Enabled -> False

IF NOT escalate to IAM on-call (phone in header)

ROLLBACK Enable-ADAccount -Identity j.doe

- 1

Facts timeline agreed UTC timeline before any discussion of causes
- 2

What went well keep it: detections, decisions, tools that helped
- 3

Contributing factors systems and conditions, not people: 'the alert was routed to an unmonitored queue'
- 4

Actions each with one owner, a due date and a tracking ticket; fewer, finished actions beat a long wish list
- 5

Publish share widely inside, TLP-labelled; revisit open actions at 30 and 90 days

Post-incident review (PIR) agenda, within 2 weeks while memory is fresh

🔥

GitLab.com database outage, 31 January 2017

An engineer deleted the primary database directory while fighting a replication problem; the backup methods then turned out not to be working, and about six hours of data were lost. Git-Lab live-streamed the recovery and published a detailed public postmortem listing each failed safeguard and the fixes. It is still used as a model of a blameless, factual review: it names the conditions that allowed the mistake, not a culprit.

🔥

Equifax, 2017: the alert that did not reach the owner

The US GAO (Government Accountability Office) report GAO-18-559 (2018) found that Equifax's internal alert about the Apache Struts flaw (CVE-2017-5638) went to a distribution list that was out of date, so the owners of the affected dispute portal did not act on it. Addressing and routing are part of writing: a perfect message to the wrong list is no message.

👤

Audience adaptation and self-test

Audience	What they need	Format	When	Writer
Executives	impact, cost, decision, next update	BLUF e-mail, 1 page	fixed cadence	CISO
Engineers	hosts, IOCs, exact steps	ticket, runbook	continuous	IR lead
Legal / DPO	data, people, dates, regulators	memo, privileged if asked	at each finding	IR lead
Staff	what to do, what not to do	short intranet note	once, then updates	Comms
Regulator	facts required by the law	official form	legal deadlines	DPO / CISO

Same incident, five audiences: change the content, never the facts

⚠️

Write as if the document will be read in court

Incident notes, chats and reports can be disclosed in litigation or to regulators. Do not speculate about guilt, do not write 'we were negligent', do not delete drafts once a legal hold is in place. When counsel directs an investigation to keep it under legal privilege, follow their instructions on labels and distribution (rules differ by country; privilege is much narrower in France than in the US).

📋

'How would you write up an incident for the board?'

Strong answer: 'One page, bottom line first: what happened in business terms, whether customers or money are affected, what we have done, how confident we are, and the one decision I need from them. No jargon, numbers instead of adjectives, and a fixed time for the next update. Technical detail goes in an appendix for those who ask.'

Statement	True	False
1. 'The server was compromised' is a complete finding title → Name the flaw, the asset and the consequence	☐	☒
2. Timestamps should be written in UTC with the original time kept	☒	☐
3. 'Likely' in ICD 203 means about 55-80 %	☒	☐
4. A playbook contains the exact commands for one task → that is a runbook; a playbook covers a scenario	☐	☒
5. A PIR assigns each action one owner and a due date	☒	☐
6. TLP:CLEAR material may be published	☒	☐

Self-test: true or false?

Q1

What three things must a reader be able to trace for each fact?

A

Its source (log, host, query), its time (UTC and which clock) and who collected it.