LLMs in an OSINT investigation

A source-led Log4Shell case file

This is not a prompt collection. It is a worked investigation: one question, four public records, a visible evidence trail and a decision that remains with the analyst.

CVE-2021-44228

CASE ANCHOR

THE QUESTION

What can public records establish about urgency, and what still requires local asset and telemetry checks?

THE METHOD

Preserve the record. Separate fact from inference. Ask the model to organise, not invent. Re-open every source before action.

The model helps with the middle of the work.

The question starts with a person. The final action ends with one.

Follow the claim, not the chatbot

A useful LLM session has a small job inside a larger investigation. The map below shows where evidence enters, where interpretation happens and where the model must stop.

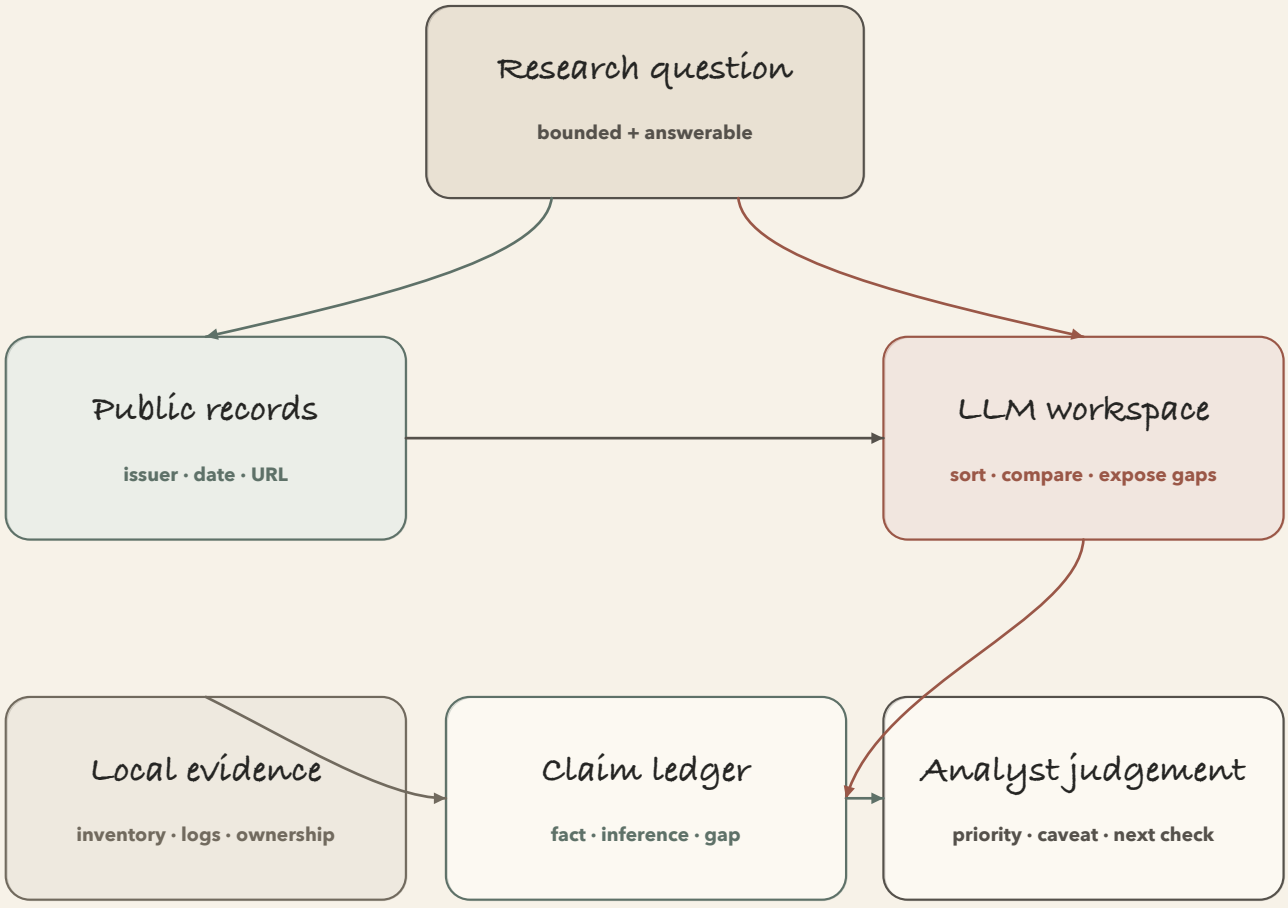

Stop condition

If the public record cannot answer whether our system is affected, the output is not 'safe' or 'compromised'. It is a named gap: check the asset and its telemetry.

Log4Shell, reconstructed from records

The sequence matters. A single CVE page describes severity; later records change what a defender should do. This timeline keeps publication, correction and observed exploitation separate.

10 DEC 2021 — NVD publishes CVE-2021-44228 with a critical 10.0 CVSS 3.x base score. [S3]

17 DEC 2021 — Apache releases 2.17.0 for Java 8 after the first response evolved. [S2]

23 DEC 2021 — The joint CISA advisory tells organisations to identify assets, patch and hunt. [S1]

23 JUN 2022 — CISA reports continued exploitation against VMware Horizon and UAG systems. [S4]

READING THE RECORD

Public evidence supports urgency and a response path. It does not reveal whether a particular organisation runs an affected component or whether exploitation succeeded there. Those are local questions.

Four records, four different jobs

A source list becomes useful only when each record has a role. The model receives the saved text, not an invitation to browse from memory.

ID	RECORD	ROLE IN ANALYSIS	DATE NOTE
S1	CISA AA21-356A	response guidance	23 Dec 2021
S2	Apache release notes	version history	living page
S3	NIST NVD CVE record	severity + scope	modified 2026
S4	CISA exploitation alert	later observation	23 Jun 2022

CHAIN OF CUSTODY, LIGHTWEIGHT BUT REAL

For every page, keep the original URL, publisher, access time and a saved copy or archive reference. Record later changes as new observations; do not silently replace the old note.

Method note: Bellingcat/GLAN emphasises traceable, repeatable research records; the Berkeley Protocol frames collection, preservation and analysis.

From record to defensible statement

The analyst does not ask the model for an answer. The analyst asks it to transform bounded material into a ledger that can be checked line by line.

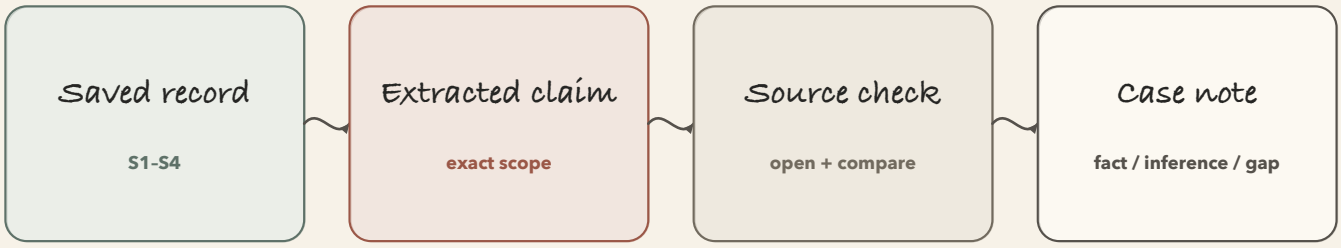

WORKED CLAIM / [S1]

CISA recommends identifying affected assets, updating them and initiating hunt and incident-response procedures.

WHAT IT SUPPORTS

A response sequence: identify, mitigate, update and hunt. This is a public fact about the guidance.

WHAT IT CANNOT SUPPORT

That our organisation is affected or compromised. The advisory has no access to our inventory or logs.

A bounded Log4Shell triage request

Scenario: a defensive analyst has four saved public records and needs a first-pass evidence ledger before checking the organisation's inventory and logs.

PROMPT GIVEN TO THE MODEL

ROLE
You are assisting a defensive threat-intelligence analyst.

MATERIAL
Use only the saved text labelled S1, S2, S3 and S4.

QUESTION
What supports urgency around CVE-2021-44228, what response path is supported, and what remains unknown for our environment?

OUTPUT
Return four rows: statement | type | source | confidence | local check. Type must be fact, inference or unknown.

RULES
Do not add facts. Do not claim exposure or compromise. If unsupported, write 'not established'. Do not perform or recommend an automatic action.

WHY THIS PROMPT WORKS

It names the evidence boundary, asks one case question, fixes the output schema and defines when the model must say it does not know.

Next page: the model output and the analyst's review

The output is useful only after every source marker is opened and checked.

Model draft, analyst-owned finding

Below is a realistic compact answer to the prompt. The right-hand marks show what the analyst accepts and what still needs local evidence.

NVD records a critical 10.0 CVSS 3.x base score.

ACCEPT

FACT · S3 · HIGH

CISA records active exploitation and tells organisations to identify, patch and hunt.

ACCEPT

FACT · S1 · HIGH

Affected assets should receive earlier review.

NARROW

INFERENCE · S1/S2 · MED

Whether our environment runs an affected component or was exploited is not established.

CHECK

UNKNOWN · LOCAL · OPEN

ANALYST CORRECTION

Replace the third row with: 'Prioritise only after matching affected versions to the local inventory.' Then open S1-S4, verify each phrase and record the access date.

Result: useful draft, not an autonomous decision

Leave a trail another analyst can follow

A credible investigation ends with its limits and its records in view. These are the four primary case sources used in this worked example.

[S1] CISA – AA21-356A

https://www.cisa.gov/news-events/cybersecurity-advisories/aa21-356a

[S2] Apache Log4j – release notes

https://logging.apache.org/log4j/2.x/release-notes.html

[S3] NIST NVD – CVE-2021-44228

https://nvd.nist.gov/vuln/detail/CVE-2021-44228

[S4] CISA – continued Log4Shell exploitation

https://www.cisa.gov/news-events/alerts/2022/06/23/malicious-cyber-actors-continue-exploit-log4shell-vmware-horizon-systems

FIELD NOTE

Keep the question narrow. Save the source before asking the model. Let the model organise claims and expose gaps. Re-open the record. Write the limit beside the finding. Connect public intelligence to local evidence only through an accountable human review.

The tool can shorten the path. It cannot own the finding.