Intelligence needs a source trail before a summary

Channel counts, narrative reach and actor attribution answer different questions.

RESEARCH BRIEFING · INTELLIGENCE · 11 NOVEMBER 2026

AI-assisted monitoring should preserve the evidence behind an alert before producing a fluent summary. Repeated claims, large channel counts and signs of coordination do not by themselves establish truth, reach or responsibility.

Key points

- Count independent reporting origins separately from amplification.
- Retain source passages, dates, translations and corrections.
- Require named review for consequential attribution.

The scale figures have limits

The EEAS's March 2026 report analysed 540 detected information-manipulation incidents from 2025 involving 10,500 unique channels. It said 65% of the incidents remained unattributed. These figures describe its monitoring sample, not a census of global activity. [1]

A channel can be a website, account or blog. It is not a person reached, an independent source or a separate operator. The unattributed share also cannot be reassigned to a named actor merely because some infrastructure appears connected.

Separate the claim from its distribution

Store the original URL, publication and retrieval times, publisher identity and passage supporting each claim. Keep translations beside the original. Record what corroborates or contradicts the claim and when the assessment changes.

Ten outlets republishing one wire report represent one reporting origin and ten distribution points. Deduplicate before counting growth. Document changes in collection scope: adding a language or platform can create an apparent surge without a corresponding change in activity.

Check provenance and context

C2PA credentials can help establish an asset's recorded provenance and integrity. They do not certify the truth of its subject or caption. A genuine image can carry a false description; missing credentials do not establish fabrication. [2]

Check dates, location, source independence and alternative explanations. Keep confidence in a claim separate from confidence in attribution. Where the evidence supports only a research lead, label it that way.

Audit the alerts

Build a bounded, human-reviewed test set. Measure precision as correct alerts divided by alerts issued, and recall as detected relevant cases divided by all relevant cases in that set. Report the case count, selection method and reviewer disagreement.

Also measure citation support, correction time and analyst review minutes. Count inaccessible citations separately. High-consequence alerts should identify the changed evidence, affected decision, uncertainty and next verification step, with a named approver.

Sources & research notes

01 / EEAS / Fourth FIMI Threat Report (March 2026)

https://euvsdisinfo.eu/eeas-4th-fimi-threat-report-march-2026/
Incidents detected during 2025; selective monitoring sample, not a global census. Channels are not people reached.

02 / C2PA / Technical Specification 2.5

https://spec.c2pa.org/specifications/specifications/2.5/specs/C2PA_Specification.html
Content provenance and integrity mechanisms do not establish that an underlying claim is true.

Independent TACITUS analysis. Research checked 11 November 2026. Published findings retain their dates and limits. Proposed curricula, tests, schedules and worked calculations are labelled analytical examples, not client results.