



Forensics Manual

Status: approved. Version: 1.2.0.

0. Why we publish this

A forensic report is system-generated and expert-approved: the mechanical steps - reading the on-chain transaction data, calculating a verified loss, checking the pattern against our own record of past exploits - are automated, but no report reaches you unreviewed. We publish the review procedure itself so you can see what "expert-approved" actually means before you ever need a report: any qualified reviewer following it independently reaches the same conclusion from the same evidence, and every determination has a recorded rationale that holds up if you dispute it.

0.1 Who reviews

A report is only as trustworthy as the person checking it. The reviewer meets a fixed bar, regardless of who holds the role:

- Hands-on experience reading a transaction trace or call graph directly - not through a summarizing tool alone.
- The ability to read Solidity or raw EVM bytecode well enough to confirm what a contract actually did.
- Familiarity with the mechanics of common exploit classes - reentrancy, price-oracle manipulation, access-control failures, flash-loan-enabled attacks, and the other techniques that show up across real DeFi incidents.

This bar does not move. Whoever reviews a report meets it already - it is never lowered to fill the role.

1. Principles that apply to every report

- **We sell analysis; you make the decision.** A report states what the evidence shows and whether it's consistent with your own policy language - it never tells you what to do about it.
- **We never overstate certainty.** Every report states what could not be established from the evidence available, whether the loss figure is a total or a floor, and how far apart the two independent measurements of that loss landed. What a report does not carry is a confidence score. A number out of a hundred invites you to weigh your claim against it, which makes it a valuation of your incident rather than a statement of what the evidence shows - and the evidence, with its gaps named, is what you are buying.
- **Observational, not instructive.** A report says "the on-chain evidence shows X, consistent with policy clause Y" - never "this claim qualifies for payment."
- **Not legal advice.** A report addresses whether the evidence is consistent with your policy's text as written. It is not an opinion on how a court or arbitrator would rule, and says so explicitly.

2. What review starts from

Before review begins, the reviewer has:

- **What you submitted:** the exploit contract address, the exploit transaction hash(es), your own policy text, and your contact details.
- **The system's draft:** the parsed transaction data, a calculated loss amount, and a match against Foreshock's own record of real, classified DeFi exploits - disclosed to the reviewer as low-confidence and starting-point only, never as a conclusion. The draft explicitly states it is not a determination: reasoning and a confidence level are supplied by the reviewer, not the system.

The reviewer always sees the underlying evidence the system used, never just its conclusion - that's what makes independent verification possible.

3. The five independent checks

Every report, regardless of how confident the system's draft was, goes through the same five checks before a reviewer approves, amends, or escalates it:

1. **The transaction(s) exist and match the claim.** The submitted hash(es) resolve on-chain, on the chain you specified, at approximately the time claimed. A claim citing a transaction that doesn't exist, or exists on a different contract than the one submitted, is an automatic escalation (§4) - not a low-confidence approval.
2. **The loss calculation is independently reproducible.** The reviewer re-derives the verified loss figure from the raw transaction data rather than trusting the system's arithmetic. The system's number is a draft, not evidence.
3. **The pattern match is substantively correct, not just superficially similar.** The reviewer confirms the identified technique actually describes what happened in this transaction, and checks whether it's a new incident or a follow-on claim on one already on record.
4. **The contract address belongs to the protocol named in your policy.** A mismatch - the wrong contract, a fork, an unrelated protocol with a similar name - is an automatic escalation.
5. **Your policy text actually addresses this scenario.** The reviewer reads your submitted policy language directly, never the system's paraphrase of it, and checks whether what happened falls inside or outside what your policy covers, in your policy's own words.

None of this is skipped because the system reported high confidence. Confidence affects how much investigation happens beyond this baseline - never whether the baseline happens at all.

4. When a report gets escalated

Some cases are held for a further, more senior review pass before anything is delivered, rather than resolved on the first reviewer's own authority. This happens when:

- The submitted transaction or contract doesn't verify against the claim.
- The attack doesn't match any pattern we've seen before at reasonable confidence - a genuinely novel technique.

- Your policy text is ambiguous, self-contradictory, or silent on whether this type of loss is covered at all.
- The claim is large enough that getting it wrong carries outsized consequences.
- You dispute a delivered report (§6).
- Anything else a reviewer's judgment flags as outside what this protocol anticipated - the protocol is a floor on rigor, not a ceiling on judgment.

Escalation isn't a sign that something went wrong with your case. It's what the process is supposed to do when a case genuinely calls for more scrutiny. Resolving a genuinely novel technique also becomes a new addition to what we check future cases against - the library of known patterns is built from real engagements, not fixed in advance.

5. What a delivered report contains

Every report, whatever its outcome, states:

- What the on-chain evidence shows, in plain, observational language.
- The verified loss amount and how it was derived - including whether it is a total or a floor. Some losses sit outside the affected address: bad debt left behind in a lending market, or deposits already invested elsewhere. Where the loss is of a kind we do not yet reconstruct, the report states the figure as a floor and says so, rather than presenting an incomplete total as a complete one.
- How far apart the two independent measurements of that loss landed. The loss is measured twice - the net change in the affected address's balances, and the sum of the transactions attributed to the incident - and the disagreement between them is reported, not smoothed over. On a case with no published total to compare against, that gap is the check available to you.
- Whether the evidence is consistent with the specific clause(s) of your own policy text - quoting or closely paraphrasing the relevant language, not a bare yes/no.
- The limitations of the analysis - what could not be established from the evidence available, and what would change the finding. This section is never empty.
- A closing statement that the report is Foreshock's independent analysis of the evidence against your policy text - evidence for your own coverage decision, not a substitute for it, and not legal advice on how your policy would be interpreted in a dispute.

A report never simply says "this qualifies" or "this doesn't." It always connects evidence to your policy's own clauses and leaves the decision to you.

6. Amendments and disputes

If independent verification changes the outcome: when the reviewer's own checks produce a different loss figure, pattern match, or conclusion than the system's draft, the delivered report reflects the reviewer's finding - not the system's - and records what changed and why.

If you dispute a delivered report: your case gets a fresh, independent review pass, starting again from the five checks in §3 - not a negotiation over the existing conclusion. The result of that re-review is recorded the same way as any other outcome, including if it confirms the original finding.

7. What's kept on record

For every report, regardless of outcome, we retain: the system's original draft, the reviewer's independent verification notes against each of the five checks in §3, the final determination and its reasoning, and - for anything amended or escalated - what changed from the draft and why. That's what makes a report defensible if you dispute it, and reproducible by a different reviewer later.