

SOLUTIONS · AI AGENT INSURANCE

# Claim-grade evidence for AI agent insurance.

For insurers, MGAs, brokers, and enterprises deploying autonomous agents.

AI agents can approve refunds, send messages, move money, and make commitments. When something goes wrong, screenshots and vendor logs are not enough. GeoClear gives underwriters better evidence and gives claims teams verifiable proof at three moments: before coverage, during coverage, and after an incident.

**GeoClear preserves evidence. The insurer determines coverage.**

GeoClear does not determine coverage. We preserve the evidence insurers need to review what happened.

## How GeoClear helps — the same substrate at three insurance moments

**1 · BEFORE COVERAGE**

**Evidence for underwriting and authorization scope.**

A signed Agent Scope Certificate captures what the agent is authorized to do at policy bind: tools, action types, spending limits, data access, human-approval thresholds, agent operating policy version. Underwrite against a machine-checkable scope, not a free-text questionnaire.

**2 · DURING COVERAGE**

**Receipts for high-risk decisions.**

When the agent issues a refund, sends a contractual response, calls a financial API, or acts without required human approval, a Runtime Decision Receipt records what it did, in what context, against what agent operating policy, and when. Receipts are the equivalent of telematics for agents.

**3 · AFTER AN INCIDENT**

**Claim Evidence Bundle for review.**

Scope certificate + runtime receipts + policy snapshot + tool-call record + verification report, packaged into one portable, customer-held ZIP. Insurer, insured, and broker each verify it independently.

## Five evidence artifacts

| ARTIFACT                     | INSURANCE MOMENT     | WHAT IT CAPTURES  |
|------------------------------|----------------------|---|
| Agent Scope Certificate      | Pre-bind             | Authorized + restricted actions, spending limits, approval thresholds, agent operating policy version       |
| Underwriting Evidence Bundle | Pre-pricing          | Governance assessment, adversarial probe summary, security posture, tool access map, escalation rules       |
| Runtime Decision Receipt     | Per high-risk action | Agent identity, tool invoked, arguments, result, policy at action time, approval status, timestamp          |
| Claim Evidence Bundle        | Post-incident        | Scope cert + runtime receipts + policy snapshot + tool calls + approvals (or absence) + verification report |
| Renewal Evidence Summary     | Pre-renewal          | High-risk action count, escalations, exceptions, drift indicators, control improvements, evidence coverage  |

## Example claim — was the agent inside or outside scope?

A customer-support agent is authorized to issue refunds up to \$500 and create tickets. Refunds above \$500 require human approval. A customer threatens legal action; the agent promises a \$7,500 refund, accepts liability on the company's behalf, and calls the refund API without approval. *Was this covered behavior or outside authorized scope?*

|   |   |   |
|---|---|---|
| <p><b>STEP 1 · AT POLICY BIND</b></p> <p><b>Agent Scope Certificate</b></p> <p>authorized                      refund ≤ \$500</p> <p>restricted                      legal admissions</p> <p>approval req.                      refund &gt; \$500</p> <p>agent op policy                      v1.7</p> <p><b>verified at bind</b></p> | <p><b>STEP 2 · AT ACTION TIME</b></p> <p><b>Runtime Decision Receipt</b></p> <p>agent                              refund \$7,500 +</p> <p>action                              admission</p> <p>tool call                              refund_api(7500)</p> <p>human approval                      missing</p> <p>timestamp                              14:32:18 UTC</p> <p><b>outside scope</b></p> | <p><b>STEP 3 · INSURER VIEW</b></p> <p><b>Claim Review</b></p> <p>finding                              outside approved scope</p> <p>evidence                              limit exceeded · approval missing</p> <p>verification                              signature valid · scope unchanged</p> <p>bundle                              portable .zip</p> <p><b>evidence preserved</b></p> |
|---|---|---|

**Tamper test:** if the reviewer changes the refund amount in the bundle from \$7,500 to \$500, the signed receipt no longer matches the canonical payload — verification fails. A screenshot does not have this property.

## Buyer value

|   |   |   |
|---|---|---|
| <p><b>INSURERS · MGAS</b></p> <ul style="list-style-type: none"> <li>Faster claim review</li> <li>Better underwriting evidence</li> <li>Less ambiguity on scope</li> <li>Stronger renewal data</li> </ul> | <p><b>AI STARTUPS</b></p> <ul style="list-style-type: none"> <li>Easier procurement</li> <li>Better risk posture</li> <li>Stronger insurance application</li> <li>Faster incident review</li> </ul> | <p><b>ENTERPRISE RISK TEAMS</b></p> <ul style="list-style-type: none"> <li>Clear accountability</li> <li>Reviewable agent behavior</li> <li>Vendor risk evidence</li> <li>More confidence deploying agents</li> </ul> |
|---|---|---|

## What GeoClear is not · talk to us

**GeoClear is not an insurer, broker, claims adjuster, coverage-determination service, risk-pricing engine, or legal-conclusion engine.** What we do: produce signed, customer-held Evidence Bundles that an insurer, insured, broker, regulator, or auditor can independently verify against a public key. **The insurer determines coverage. We preserve the evidence.**

» Live page: [geoclear.io/solutions/ai-agent-insurance](https://geoclear.io/solutions/ai-agent-insurance)

» Interactive 7-min demo (real client-side cryptography; nothing leaves your browser; downloads a real Evidence Bundle ZIP): [geoclear.io/demo/agent-claim-review](https://geoclear.io/demo/agent-claim-review)

» Talk to us: [hello@geoclear.io](mailto:hello@geoclear.io)

Page 2 of 2 · GeoClear for AI Agent Insurance · GeoClear preserves evidence. The insurer determines coverage.