

AI Operational Evidence Platform

Verifiable control for AI actions.

GeoClear™ is an AI-first product company building an independent operational evidence layer for AI and agent actions. It helps organizations understand their agent estate, determine which actions require stronger control, evaluate whether a specific consequential action has the required authority, policy, evidence and review before a customer-designated system accepts it, and preserve the decision-to-outcome record under customer control. GeoClear works alongside existing models, gateways, identity, policy, workflow, cloud, data, security and mission systems.

The GeoClear operational assurance lifecycle

DISCOVER	ASSESS	PREPARE	CONNECT	GOVERN	VERIFY	PRESERVE	AUDIT & LEARN
Agents, tools, actions, owners, boundaries	Materiality, authority, data, dependencies	Readiness, gaps, remediation, profiles	Agent Gateway, APIs, MCP/A2A, identity	Policy Foundry, Evidence Contracts, review	Exact action, authority, evidence, five outcomes	Customer-held evidence, acceptance, outcome	Search, replay, reports, control insight

Customer systems enforce. Customers retain the evidence.

OPERATIONAL IMPACT GeoClear helps organizations identify high-consequence action boundaries, assess and remediate readiness gaps, prevent unsupported actions within configured customer boundaries, preserve evidence for review and corrective action, and measure operational control effectiveness without replacing existing platforms.

PLATFORM CAPABILITIES

- **Agent Discovery & Registry:** discover or register agents, tools, consequential actions, owners, environments and receiving-system boundaries.
- **Agent Assessment & Readiness:** classify action risk, assess authority, data, reversibility, dependencies and evidence gaps, then produce a remediation and rollout path.
- **Policy Foundry & Evidence Contracts:** turn authoritative requirements, exceptions and expert judgment into reviewed, versioned and testable action-level requirements, including sequence-aware, time-aware and cumulative conditions.
- **Runtime Verification & Action Control:** evaluate the exact agent, represented principal, action, target, authority and required evidence; return Accept, Hold, Reject, Block or Escalate.
- **Agent Gateway (L1-17):** connect agent, model, MCP/A2A, API and tool traffic through governed action boundaries while customer systems remain the enforcement point.
- **Evidence Vault & Independent Verification:** preserve customer-controlled operational evidence, policy/profile references, acceptance, actual outcome and approved verification assets.
- **Audit Vault & Analytics:** search, reconstruct and replay action chains; investigate non-Accept outcomes and measure control effectiveness.
- **Model, Runtime & Supply-Chain Assurance:** preserve provenance and operating restrictions for models, runtimes, tools, RAG sources and software dependencies supporting an action.

WHERE IT APPLIES

- Agent onboarding, inventory, ownership, delegated authority and multi-agent lineage
- System access, role changes, privileged operations and entitlement workflows
- Procurement, approvals, obligations, payments, vendor actions and separation of duties
- Software, cloud, infrastructure, database, export and destructive actions
- Enterprise knowledge, RAG sources, publishing and regulated communications
- Cyber-defense agents, incident response, high-risk tool use and human escalation
- Cross-platform agent-to-tool and agent-to-agent workflows across gateways and systems
- Mission support, GEOINT, edge, physical AI and disconnected operations under deployment-specific profiles

WHY GEOCLEAR

- **Action-specific:** model or platform assurance does not by itself establish whether this exact action should be accepted.
- **Identity is not authority:** actor identity remains distinct from delegated authority, action, target, environment and time.
- **Customer-controlled:** customer systems enforce; customer policy, approvals, evidence custody and final responsibility remain with the customer.
- **Independent and portable:** operational evidence can be reviewed beyond the originating runtime, gateway or workflow, including through local verification.
- **Precise about scope:** GeoClear verifies the approved evidence path; it does not certify substantive model correctness, legal compliance, mission validity or risk-free operation.

17 governed capability families

NOTE Capability availability varies by release and deployment scope.

L1-01 Agent Discovery & Registry

Know which agents, tools, actions, owners and boundaries are in scope.

L1-02 Agent Identity & Access

Establish who or what is acting, for whom and within which delegated limits.

L1-03 Agent Assessment

Determine which actions need stronger control before production.

L1-04 Agent Readiness

Turn findings into a clear readiness, remediation and rollout path.

L1-05 Policy Foundry

Translate authoritative requirements into reviewed, versioned and testable policy, including sequence-aware, time-aware and cumulative requirements.

L1-06 Evidence Profiles & Contracts

Define exactly which authority, context, evidence, freshness, review and outcome data each action requires.

L1-07 Runtime Verification & Action Control

Evaluate the exact consequential action before downstream acceptance and return one of five governed outcomes.

L1-08 Evidence Vault

Keep the decision, acceptance and outcome record under customer control.

L1-09 Audit Vault & Analytics

Investigate, reconstruct, replay and report on action chains and control effectiveness.

L1-10 Operational Control & Case Management

Move non-Accept outcomes into accountable review, remediation and closure workflows.

L1-11 Integration Hub

Connect models, agent frameworks, identity, policy, ITSM/SIEM and customer enforcement through vendor-neutral adapters.

L1-12 Assurance, Simulation & Continuous Validation

Exercise policy, contracts, adapters and failure paths through simulation, replay and adversarial testing.

L1-13 Model, Runtime & Supply-Chain Assurance

Track provenance, qualification, restrictions, substitution and reassessment across supporting components.

L1-14 Developer Platform

Provide APIs, SDKs, schemas, fixtures, conformance tooling and local verification support.

L1-15 Enterprise Administration, Security & Deployment

Support tenancy, RBAC, dedicated environments, keys, observability, backup, rollback and lifecycle.

L1-16 Edge, Physical AI & Disconnected Operations

Support local action boundaries, cached trust, device identity, safe-state behavior and reconciliation.

L1-17 Agent Gateway

Provide a governed path for agent, model, MCP/A2A, API and tool traffic while preserving action context and customer enforcement boundaries.

DEMONSTRATED CAPABILITY & READINESS

- **Technology readiness:** TRL 5 - integrated working prototype validated in a representative test environment; next milestone is a fixed-scope customer evaluation in a relevant environment.
- **Demonstrated paths:** Accept, Hold, Reject and Escalate, plus tamper-integrity detection, customer-held evidence, outcome binding and local-verification paths. Additional outcomes and integrations remain release- and evaluation-specific.
- **External demonstration:** working platform demonstrated to federal-sector systems-integrator leadership under confidentiality.
- **IP posture:** patent-pending operational evidence technologies for AI actions; U.S. provisional patent applications filed.

HOW AN EVALUATION STARTS

1. Select one consequential action, named owner, receiving system and customer enforcement point.
2. Discover or register the relevant agents, tools, authority, dependencies and boundaries.
3. Assess risk and evidence gaps; produce a readiness and remediation plan.
4. Approve one Evidence Contract and integrate the pre-action evaluation path.
5. Run positive, negative, tamper, replay, expiry, revocation, dependency-failure and outcome-binding scenarios.
6. Deliver the evaluation report, operational evidence export, verification results and production recommendation.

DEPLOYMENT & INTEGRATION

- **Managed minimized evaluation:** GeoClear evaluates approved attributes, commitments and declared results without requiring full business objects.
- **Customer-controlled deployment:** GeoClear components can run inside a customer or authorized-integrator boundary where sovereignty or isolation requires it.
- **Customer-computed / bind-only:** the customer computes the approved result and GeoClear binds the declared result to the action, acceptance and outcome.
- **Integration points:** agent platforms, MCP/A2A ecosystems, identity providers, policy engines, gateways, workflows, ITSM/SIEM, cloud, on-premises and edge systems.

ROUTES TO MARKET

- Paid fixed-scope Operational Evidence Evaluation
- Direct enterprise or government delivery
- Prime / systems integrator or approved marketplace delivery
- Managed, dedicated, customer-controlled and deployment-specific licensing

PUBLIC-SAFE BOUNDARY

This statement describes customer-facing capabilities and product boundaries. It does not disclose unpublished algorithms, internal schemas, patent claim strategy or counsel-sensitive implementation details.