SECURITY ASSESSMENT

TRUSTRAY

Security Assessment Service

Security for the AI and cloud era — from assessment to remediation.

Aligned with OWASP Top 10 OWASP LLM Top 10 CIS Benchmark

Trustray LLC — Godo Kaisha Trustray

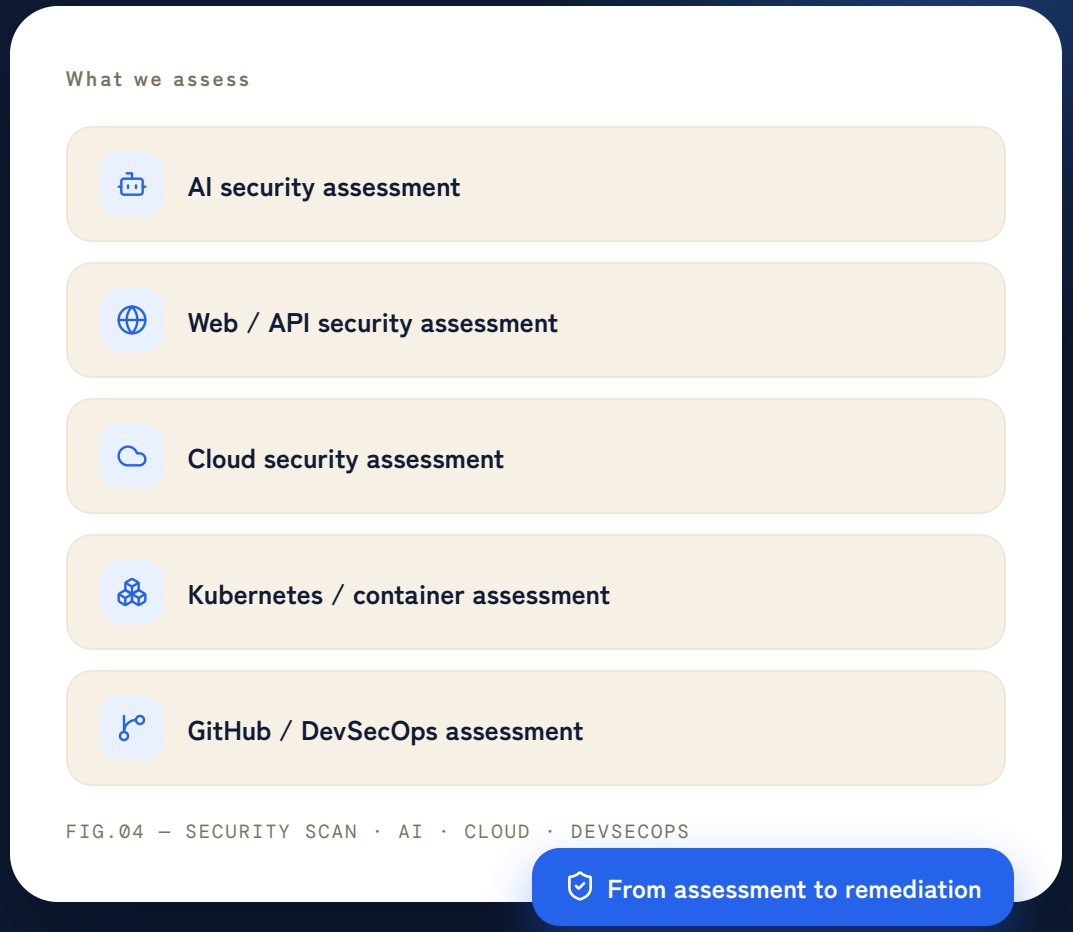

02 /CHALLENGES

Do any of these sound familiar?

We support the full path from assessment to remediation.

We want AI at work, but worry about data leaks

Moving forward without checking AI-specific risks such as prompt injection and data exfiltration paths feels unsafe.

Our PoC cannot move to production

The prototype works, but unresolved security concerns keep the production decision on hold.

Unsure about cloud and GitHub permissions

IAM and repository permissions have never been reviewed systematically, and it is unclear whether they are safe as they are.

FOR — A good fit if you are

- A company about to move AI into production work that needs a security review first
- A SaaS company or startup running public web services or APIs
- A company with 10–100 people and no dedicated security role
- An agency or AI consultancy whose client projects require security compliance

What we assess

From AI agents to the cloud, Kubernetes, and GitHub foundations beneath them. Assess only the areas you need. Each assessment is detailed on the following pages.

FEATURED

AI security assessment

AI AGENT ATTACK SURFACE

We review AI agents, RAG, MCP, and internal AI chat against the **OWASP Top 10 for Agentic Applications / OWASP LLM Top 10**.

Web / API security assessment

We review web apps, admin consoles, and public APIs against the OWASP Top 10 / OWASP API Security Top 10.

- Authentication, authorization, sessions
- Input validation, injection defenses
- Rate limiting, API key management

Cloud security assessment

We review AWS, Google Cloud, and Alibaba Cloud configurations based on the CIS Benchmarks.

- IAM (permission inventory, least privilege)
- Network and storage exposure settings
- Logging, encryption, audit settings

Kubernetes / container assessment

We review cluster and container runtime settings for permissions, isolation, and exposure.

- RBAC, secret management, NetworkPolicy
- Container images, Ingress, Pod Security

GitHub / DevSecOps assessment

We review the development flow from repositories through CI/CD and deployment, focusing on permissions and supply chain.

- Repository permissions, branch protection
- Actions workflow permissions, secrets
- Deploy permissions, supply chain defenses

AI security assessment

We assess AI agents, RAG, MCP, and internal AI chat for AI-specific risks, based on the OWASP LLM Top 10 and OWASP Top 10 for Agentic Applications.

For these challenges Unsure what AI agents can safely be trusted with · Worried internal data could leak through AI · Few assessors understand AI-specific risks

What we check

01 **Prompt injection resilience**
Per input path, we check whether instructions hidden in external input or ingested documents can hijack the AI.

02 **Data exfiltration paths**
We check for routes where internal or personal data could leak through prompts, answers, or logs.

03 **Agent permission design**
We check whether the tools, APIs, and data an agent can reach exceed what its job requires.

04 **MCP & integration supply chain**
We check the trustworthiness, permission scope, and update paths of connected MCP servers and external tools.

05 **Logs & audit trails**
We check that logging and auditing can answer "what did the AI do?" after the fact.

Available plans

From a Light assessment (this area only)

Deliverables

- Assessment report (prioritized)
- Remediation checklist
- Debrief with Q&A

Representative findings (illustrative)

HIGH	Agent tool permissions are too broad	AI-04
MED	Personal data stored in plain text in chat logs	AI-07

Web / API security assessment

We assess web apps, admin consoles, and public APIs against the OWASP Top 10 and OWASP API Security Top 10.

For these challenges

- Want a third-party check before going live
- Not confident in API authorization design
- Vulnerability handling is reactive

What we check

01 **Authentication, authorization & sessions**
We check login, permission checks, and session handling for impersonation and privilege-escalation gaps.

02 **Input validation & injection defenses**
We check resilience against SQL injection, XSS, and other attacks that start from external input.

03 **Rate limiting & API key management**
We check defenses against brute force and request floods, plus API key issuance, revocation, and scope.

04 **Transport & data protection**
We check TLS settings, encryption of sensitive data, and leaks via caches or responses.

05 **Error handling & information exposure**
We check that error messages and debug output do not reveal internal structure or implementation detail.

Available plans

From a Light assessment (this area only)

Deliverables

- Assessment report (prioritized)
- Remediation checklist
- Debrief with Q&A

Representative findings (illustrative)

HIGH	API authorization is not checked per object	WEB-02
MED	Admin console sessions live too long	WEB-05

06 /SCOPE 03 — CLOUD

Cloud security assessment

We assess AWS, Google Cloud, and Alibaba Cloud configurations and permissions based on the CIS Benchmarks.

For these challenges

- Worried about cloud misconfigurations
- No clear view of who holds which permissions
- Audit and logging left unconfigured

What we check

01 **IAM & permission inventory**
We inventory users, roles, and service accounts and check that they follow least privilege.

02 **Network exposure**
We check security groups and firewalls for paths unintentionally open to the internet.

03 **Storage & data protection**
We check storage exposure settings, encryption, and backup protection.

04 **Logging & audit settings**
We check that operation and audit logs cover the right scope and period, protected from tampering.

05 **Multi-cloud configuration**
We check permissions and connectivity across clouds, including Alibaba Cloud setups.

Available plans

From a Light assessment (this area only)

Deliverables

- Assessment report (prioritized)
- Remediation checklist
- Debrief with Q&A

Representative findings (illustrative)

HIGH	Storage is publicly accessible	CLD-02
LOW	Audit log retention is too short	CLD-07

07 /SCOPE 04 — KUBERNETES

Kubernetes / container assessment

We assess cluster and container runtime settings for permissions, isolation, and exposure.

For these challenges

- Not confident the cluster is configured safely
- Secrets and permissions managed ad hoc
- Image vulnerabilities pile up unaddressed

What we check

01 **RBAC & permission design**
We check ServiceAccount and Role permissions for excess, together with namespace isolation design.

02 **Secret management**
We check how credentials are stored, encrypted, and referenced, and surface anything that could leak.

03 **NetworkPolicy & traffic control**
We check that pod-to-pod and outbound traffic is restricted to the minimum needed.

04 **Images & supply chain**
We check base image vulnerabilities, image sources, and signing and scanning practice.

05 **Ingress & Pod Security**
We check external entry points and pod privileges (root execution, privileged containers).

Available plans

From a Light assessment (this area only)

Deliverables

- Assessment report (prioritized)
- Remediation checklist
- Debrief with Q&A

Representative findings (illustrative)

HIGH	Broad permissions on the default ServiceAccount	K8S-01
MED	Privileged containers running in production	K8S-04

GitHub / DevSecOps assessment

We assess the development flow from repositories through CI/CD and deployment, focusing on permissions and the supply chain.

For these challenges

- Leavers and contractors may retain access
- Opaque handling of CI/CD secrets
- No defenses against supply chain attacks

What we check

01 Repository permissions & branch protection
We inventory member access and check settings that prevent direct pushes and skipped reviews on main.

02 Actions workflow permissions
We check GitHub Actions token permissions and the use and version pinning of third-party actions.

03 Secrets management
We check how CI/CD credentials are stored and referenced, and whether they leak into build logs.

04 Deploy permissions & environment separation
We check production deploy approval flows and per-environment permission separation.

05 Dependencies & supply chain
We check vulnerability detection for dependencies (e.g. Dependabot) and the update process around it.

Available plans

From a Light assessment (this area only)

Deliverables

- Assessment report (prioritized)
- Remediation checklist
- Debrief with Q&A

Representative findings (illustrative)

HIGH	Actions tokens always granted write permissions	GH-03
MED	Branch protection is not configured	GH-01

09 /APPROACH

An assessment centered on configuration, permissions, and design.

It is not an offensive test against live environments (penetration testing). We combine reviews of architecture, settings, and code with tool-based scans where useful, so it can run against services in development or in production without impact.

AI AGENT ATTACK SURFACE

Key items in the AI security assessment

- Prompt injection and data exfiltration paths
- Over-privileged agents (excessive tool access, authorization design)
- Supply chain risks in MCP and external tool integrations
- Missing logs and audit trails

Standards and reference frameworks

- OWASP Top 10
- OWASP API Security Top 10
- OWASP LLM Top 10
- OWASP Top 10 for Agentic Applications
- CIS Benchmark

10 /PROCESS

How the assessment works

The initial consultation is free. We review your current situation and propose the right scope and approach.

01

Free

Consultation

We review your system architecture, cloud usage, AI usage, and development structure.

02

Scope definition

We propose which areas to assess: AI, web/API, cloud, Kubernetes, or GitHub.

03

Security assessment

We review settings, permissions, logs, authentication, and development flow.

04

Report and debrief

We deliver a risk list ranked by impact and priority, and walk you through it.

05

Optional

Remediation support

We support configuration changes, code fixes, CI/CD improvements, and AI usage guidelines.

✓ **The assessment is not the end — we make it clear what to fix first and carry it through to remediation in one flow.**

We do not stop at pointing out problems.

Beyond flagging issues, we organize **where to start and how to fix each one** before handing over.

- Security assessment report (risk list with impact and priority)
- Remediation checklist and roadmap
- Debrief session with Q&A
- Configuration and implementation support where needed

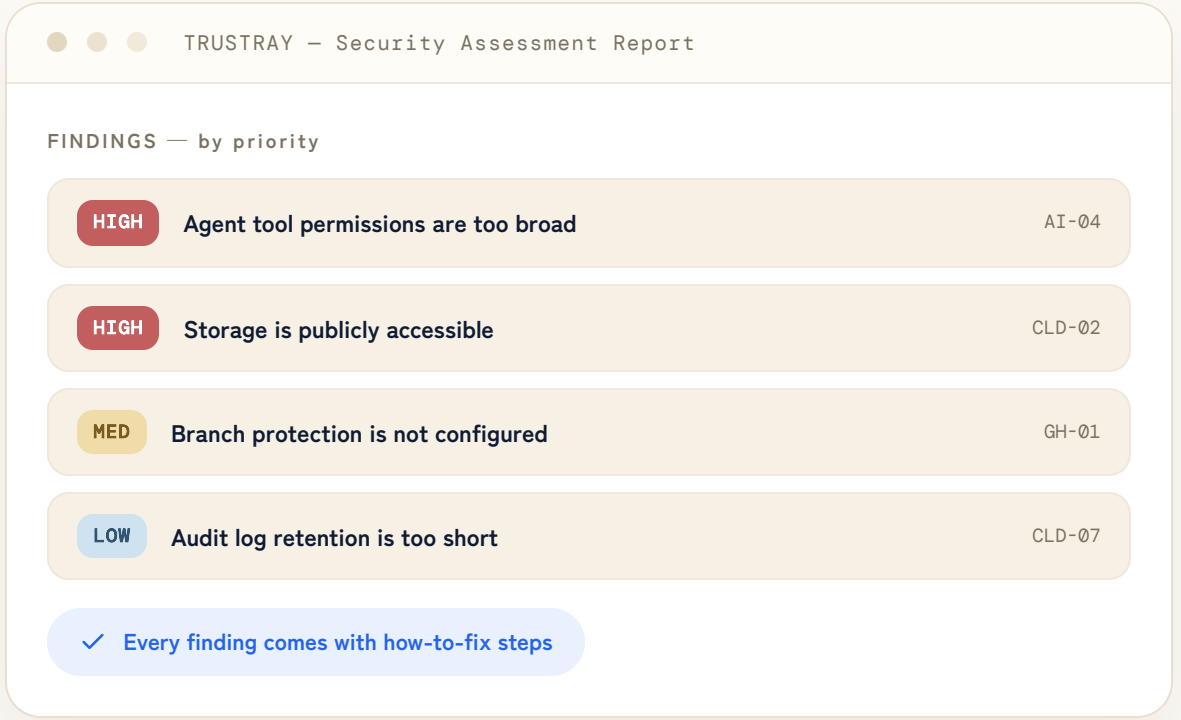

Assessment plans

We review how you use AI, web, cloud, and GitHub / CI/CD today, then propose the scope that fits best.

LIGHT

Light assessment

Focus on a single area and review its main risks.

Scope

A focused review of one area

Best for

Companies that want to start by checking part of their risks

Deliverables

- Summary assessment report
- Key risk list
- Remediation checklist

STANDARD

Recommended

Standard assessment

Combine multiple areas such as web / API, cloud, and GitHub.

Scope

A standard review combining 2–3 areas

Best for

Companies ready to review their development and cloud environments properly

Deliverables

- Assessment report
- Prioritized remediation list
- Recommended response policy

FULL

Full assessment

A cross-cutting review of AI, web, API, cloud, Kubernetes, and DevSecOps.

Scope

A comprehensive review across the whole stack

Best for

Companies reviewing AI adoption, cloud operations, and the dev platform together

Deliverables

- Detailed assessment report
- Remediation roadmap
- Implementation support where needed

We start by reviewing your current situation and propose the right scope and approach.

What makes Trustray different

Few assessment services can cover both AI security and cloud security in one engagement.

01

AI and cloud, assessed together

We review AI-agent-specific risks and the cloud, Kubernetes, and CI/CD foundations beneath them in a single engagement.

02

Alibaba Cloud support

We also assess Alibaba Cloud configurations — still rare in Japan. Ask us about services for the Chinese market as well.

03

We do not slow development down

Instead of ideal-but-impractical controls, we propose prioritized improvements your current team can realistically adopt.

04

Beyond the assessment

We do not stop at the report. We stay with you through configuration changes, implementation, and operational guidelines as needed.

Company

Trustray LLC is a technology company focused on AI, cloud, and consulting. We help companies clarify business issues and support the full path from AI strategy to system development, cloud foundations, and operational improvement.

Company name	Trustray LLC
Legal name in Japan	Godo Kaisha Trustray
Business	AI adoption consulting / AI system development / Cloud infrastructure / Operational improvement
Focus areas	Generative AI & AI agents / Cloud / Security / Distributed systems / Business application development

CONTACT

Start by understanding where you are.

If you have concerns about the security of your AI, web apps, APIs, cloud, or GitHub environment, get in touch. The initial consultation is free — we review your situation and propose the right scope and approach.

Talk to us for free →

© 2026 Trustray LLC – AI × Cloud × Consulting