Fluenci Enterprise Security & Data Governance

How Fluenci verifies workforce AI competence without exposing enterprise data — prepared for CISOs, Chief Learning Officers, and HR procurement leadership.

1. Executive summary

Enterprise AI licences are deployed far faster than the operating discipline needed to use them safely. The measurable risk is not model capability — it is human execution: unredacted PII pasted into prompts, prompt-injection payloads obeyed without challenge, ungrounded output forwarded to customers, and token spend that scales with sloppy prompting. Fluenci is an assessment and training platform that measures those behaviours in a sandboxed simulator and produces audit evidence of workforce readiness.

2. Data handling commitments

- **No public model training.** Enterprise data, uploaded SOPs, and employee flight logs are never used to train public AI models.
- **Synthetic simulation payloads.** Standard scenarios ship with synthetic source documents, so production data never has to enter an assessment.
- **Least-privilege storage.** Telemetry records are row-level isolated per account; access is scoped by authenticated identity and role, not by client-side state.
- **Encryption in transit and at rest.** All traffic is served over TLS; assessment records and identity attestations are stored encrypted.
- **Customer-controlled visibility.** Individual profiles default to private; publication of any credential is an explicit employee action.

3. What Fluenci measures

Competency	Risk contained	Evidence produced
Context design	Ambiguous, rework-heavy output	Role/objective/constraint/format compliance
Fact auditing	Hallucinated claims reaching customers	Source-grounding and citation checks
Data privacy guardrails	PII and IP leakage into prompts	Redaction and refusal telemetry
Injection resistance	Payload-driven policy override	Guardrail-hold pass/fail traces
Token compression	Uncontrolled API/token spend	Prompt token deltas vs. benchmark

4. Audit evidence and compliance alignment

Every graded run is written to an immutable execution trace: the submitted prompt, the guardrails under test, the pass/fail checks, the score, and a tamper-evident hash. Identity can be bound to the record through a government-ID check, so an attestation is traceable to a verified person.

- **ISO/IEC 42001** — competence and awareness evidence for AI management system controls.
- **SOC 2** — personnel training records and control-operation evidence for the security criteria.
- **GDPR** — demonstrable data-minimisation training and PII-handling assessment for staff using AI tools.
- **HIPAA** — workforce training evidence for staff handling protected health information in AI workflows.

5. Identity, access and provisioning

- SAML 2.0 SSO with Okta, Microsoft Entra ID (Azure AD), and Google Workspace.
- SCIM user provisioning and de-provisioning tied to your directory of record.
- Readiness scores, skill telemetry, and completion records exported to your LMS or ATS.
- Role-separated administration: L&D reporting access is distinct from security review access.

6. BYO-SOP co-creation

The Fluenci team works alongside your L&D, enablement, and security leaders to encode proprietary handbooks, brand policies, and compliance rules into scenario assessments. Custom modules are stress-tested against adversarial payloads before rollout, then piloted with a target cohort so leadership sees baseline and post-training telemetry side by side.

7. Rollout model

- **Week 1–2 — Baseline.** Cohort runs the universal core; leadership receives a gap report.
- **Week 3–4 — Co-build.** BYO-SOP modules authored and adversarially tested with your teams.
- **Week 5–6 — Pilot.** Guided rollout, live telemetry, remediation for failing competencies.
- **Ongoing — Assurance.** Recurring assessment cycles and executive workforce-readiness reporting.

Contact: enterprise inquiries via the contact form at fluenci.ai/for-businesses. This document describes platform capabilities and is provided for evaluation purposes; contractual security commitments are defined in your enterprise agreement.