

The six essential practices (AI6)

The National AI Centre's Guidance for AI Adoption (October 2025) condensed the ten guardrails of the Voluntary AI Safety Standard into six essential practices for organisations that develop or deploy AI. This datasheet takes every action in the guidance, says whose action it is when you deploy Senua AI's software, and what we provide. The software is narrow AI: it judges requests to your web services against what it has learned is normal for them. We do not use models. No LLMs. No SLMs. Of any kind.

Every action, and what we provide

Action	Whose	What Senua AI provides
1.1 Assign a senior leader as AI governance owner	Yours	A named accountable contact at Senua AI for the software
1.2 Create an AI policy	Yours	The facts on this datasheet, for your policy
1.3 Make a person accountable for every AI system	Yours	The supplier side of that accountability: one company for the whole system
1.4 Train your accountable people	Yours, with our support	Product documentation
1.5 Clarify supply chain accountabilities	Ours	One supplier for the whole system, with no third-party model or AI service inside; the shared responsibility model in the procurement datasheet
1.6 Turn the AI policy into a governance framework	Yours	Our mappings to ASD Essentials, the ISM and ISO/IEC 27001 at senuaai.com/technology/standards/
2.1 Carry out a stakeholder impact assessment	Yours, with our support	The people affected are visitors to your web services; report-only mode lets you see every verdict before any visitor is blocked
2.2 Create contestability channels	Yours, with our support	Every block carries a reference to its evidence record, written before the block is answered, so a challenged decision can be looked up and reviewed
2.3 Engage stakeholders throughout the lifecycle	Yours	What the software does, in plain words, on this datasheet and senuaai.com/technology/standards/ai6/
2.4 Find systemic unwanted impacts from feedback	Yours, with our support	Evidence records and the audit log, to look for patterns
3.1 Screen AI systems for risk	Yours, with our support	Narrow AI, one purpose, report only by default, every security feature in every edition
3.2 Manage traditional IT, narrow, general purpose and agentic AI differently	Yours	Its classification: narrow AI; not general purpose; not agentic
3.3 Assess risks for each use case	Yours	The standards datasheets and the cryptographic register
3.4 Apply controls to match the risk	Yours, with our support	Built-in controls: multi-factor sign-in, roles, tamper-evident audit, encryption at rest, report-only mode
3.5 Investigate and learn from AI incidents	Both	Our disclosure policy and our notice to affected customers within 24 hours of confirming a flaw or incident (both from November 2026)
4.1 Keep an AI register	Yours, with our support	The register entry: purpose, data source (your own traffic only), location (your servers), accountability, human control, contestability

Action	Whose	What Senua AI provides
4.2 Disclose your use of AI	Yours	The description on this datasheet, for your disclosure
4.3 Document capabilities and limitations	Ours	What it does and does not do: one task, no generated content, no tools, no agency; an AI statement for each product from December 2026
4.4 Be transparent across the AI supply chain, including training data sources	Ours	Every component is Senua AI's own; no third-party model; the only data it learns from is your own traffic; every algorithm in the public cryptographic register
4.5 Explain AI outcomes	Ours	The evidence behind every block, found by its reference
5.1 Ask the supplier for proof of testing	Ours	Our measured results on public datasets, each with its pass mark set before the run, at senuaai.com/technology/evidence/
5.2 Test before you deploy	Yours, with our support	Report-only mode: the software judges your real traffic and records every verdict without blocking anything
5.3 Monitor after you deploy	Yours, with our support	Verdicts and evidence recorded; events delivered, signed, to your log platform
5.4 Extend data governance and cybersecurity to AI	Ours	What the software learns is encrypted at rest under your key; our ISM, Essentials and ISO mappings
5.5 Stress-test the system	Both	Your own testing is welcome; our independent penetration testing from March 2027
5.6 Consider independent testing	Ours	An accredited tester's report under NDA from March 2027, then every six months
6.1 Ensure meaningful human oversight	Yours, with our support	Report only by default; every administrative action in a tamper-evident log
6.2 Build in human override points	Ours	Blocking is switched on, and off, by an administrator at any time, and each change is audited; the service can be stopped
6.3 Train the people who oversee AI	Yours, with our support	Product documentation
6.4 Keep alternative pathways for critical functions	Yours, with our support	The software works in front of your services beside your existing controls; in report-only mode it never touches the traffic

About this datasheet

Senua AI's own mapping to the National AI Centre's Guidance for AI adoption: foundations (October 2025). The guidance is voluntary and this datasheet implies no endorsement by the National AI Centre. Dated items are on the assurance roadmap, senuaai.com/technology/standards/roadmap/. Questions: sales@senuaai.com.