

nVisionIT

Responsible AI Use Policy: summary

Policy version 2.0, in effect since July 2026 · Summary issued September 2026 · www.nvisionitgroup.com/responsible-ai.html

AI assists. Humans own.

nVisionIT has a Responsible AI Use Policy that governs every use of AI in a client engagement, from a coding assistant to an agentic tool. It is owned by the directors, approved by the CEO, reviewed every six months, and enforced through tool approval, data classification and a named human reviewer on every deliverable. This is the summary. The full policy is client-facing and we will send it to you on request.

Eight commitments that apply to every engagement

Our programme is aligned to, and interpreted consistently with, the Microsoft Responsible AI Standard, ISO/IEC 42001 for AI management systems, the NIST AI Risk Management Framework, and the technology and information governance principles of the King V Report on Corporate Governance.

- **Accountability.** We remain fully accountable for any AI system used in your environment. People review, verify and approve every AI-generated output.
- **Fairness.** Structured processes to minimise bias in model selection, prompt design, data handling and output evaluation.
- **Inclusiveness.** AI-assisted design and interface output is checked for accessibility to at least WCAG 2.2 AA before delivery.
- **Privacy.** Data minimisation, purpose specificity, consent, secure processing and no exposure of confidential data to external AI platforms without your written approval, under POPIA and the Mauritius DPA.
- **Transparency.** You get clear visibility into how AI functions in your engagement, how outputs are generated and validated, and where AI was used in a deliverable.
- **Human oversight.** Where AI materially informs a decision with legal or similarly significant effects on a person, a competent human makes or meaningfully reviews that decision first. Affected people can request human review.
- **Security.** AI-enabled processes are resilient, governed, and protected from misuse, manipulation and leakage.
- **Validity.** All AI-generated material is quality-checked, tested and verified for accuracy and business relevance before it reaches a deliverable.

The Zero Exposure Rule

No confidential, regulated, personal or client-owned data may be submitted to an external AI platform unless a contractual agreement exists, the platform meets your security requirements, and explicit written approval has been obtained. There is no informal path around this.

What our people may not do

- Put client, citizen, confidential or third-party data into a public AI tool without explicit approval.
- Ship AI-generated code without engineer review, testing and named sign-off, or deploy unvalidated code into your environment.
- Use AI in a way that breaches law, contract, privacy requirements or industry standards.
- Generate deceptive, fraudulent or misleading content.

Agentic tools: AI that acts on its own

Agentic tools plan, generate, test, iterate and execute multi-step work on their own. They carry more risk than a coding assistant, so they carry additional controls.

- **Only where you have agreed in writing.** Agentic tools are deployed in a client environment only where that has been explicitly agreed with you in writing. There is no default.
- **Never wired to live systems without authorisation.** They are never connected to production systems or to data stores holding personal or confidential information without your written authorisation and a security review signed off by our Engineering Lead.
- **Same review, same gates.** Anything an agentic tool produces for a deliverable goes through the same human review, validation and quality gates as any other AI-assisted output.
- **You can ask what it touched.** You may request details of any agentic tool used on your engagement, including the scope of access it was granted and the controls applied to it.

You are told where AI was used

We provide a standard Client AI Disclosure whenever AI generated more than incidental content in a deliverable, such as a business requirements document, solution architecture document, specification or design. It is not held back until you ask. It covers which AI tools were used and their approved status in our registry; the controls applied, including human review, quality gates, security scanning and the Zero Exposure Rule; how your data was protected during the work; and the human accountability chain, meaning who reviewed it, who approved it and who is answerable for it.

Who enforces it, and how

- **Board and executive oversight.** Consistent with King V, our governing body is accountable for the governance of AI, information and technology. The CEO approves material changes to the policy once it is in force.
- **A governed tool registry.** Only tools on our approved registry may be used on client work. Each entry records security posture, licensing, data residency, risk tier, whether the vendor trains on your data, approved use cases and the approval date.
- **Named human ownership.** Every AI-generated artefact delivered to you is approved by a named engineer or consultant, version-controlled with provenance tagging, and validated against our Definition of Done.
- **Continuous monitoring.** We review model performance and output quality, emerging security risk from AI tooling, compliance under POPIA and the Mauritius DPA, and the effectiveness of the controls themselves, on a set cadence.

Incidents and review

Any suspected or confirmed AI-related incident, including data leakage, incorrect output, harmful content or unauthorised access, is escalated immediately for containment and, where required, client notification. We investigate under our POPIA breach protocol, or the Mauritius DPA protocol where Mauritius-based data subjects are involved, which includes notifying the Data Protection Office within 72 hours where feasible. You may request a written incident report once an investigation concludes.

The policy is reviewed every six months, and sooner if the regulatory environment shifts, your contractual requirements change, a new AI capability introduces a risk the current policy does not address, or an incident reveals a gap.

Ask us for the full policy. The full Responsible AI Use Policy is client-facing. We will send it to you, with the Client AI Disclosure template, the tool-registry entries and the security evidence relevant to your engagement: info@nvisionit.co.za.