40° South Guard

Guard alongside Guardrails AI

Different problems, different products

Version 1.0 | May 2026

Why this comes up

Guardrails AI is an open-source framework for validating LLM outputs against structured specifications. It lets developers define validators (using RAIL specs or Python) that check whether model outputs meet quality, format, and safety requirements.

The question we hear is: "We are using Guardrails AI to validate our model outputs. Why do we need Guard as well?"

The short answer: Guardrails AI is a validation framework. Guard is a compliance product. Guardrails AI answers "does this output match the spec I defined?" Guard answers "can we prove to a regulator that our controls are active and our evidence is tamper-proof?"

Side by side

Capability	40 South Guard	Guardrails AI
Primary purpose	Per-call compliance evidence and attestation for all AI providers	Open-source framework for validating LLM outputs against structured specs
Architecture	Inline API proxy. One API change, every provider covered.	Library integrated into application code. Customer builds and maintains validators.
Australian PII	Purpose-built detection with checksum validation for TFN, Medicare, ABN, and ACN	No built-in Australian PII detection with checksum validation. Custom validators possible.
Regulatory mapping	Every call mapped to CPS 234, the AI Safety Standard, and the Privacy Act	No regulatory mapping. Validators check output quality and format.
Per-call attestation	Cryptographically signed attestation per API call with tamper-evident 7-year audit trail	No attestation. No evidence vault. Validation results stay in application logs.

AI-specific controls	Prompt injection detection, model inventory, shadow AI blocking, policy enforcement	Output validation (format, type, quality), re-asking on failure, custom validator hub
Data residency	Australian-hosted infrastructure. Data never leaves AU jurisdiction.	Runs wherever the application runs. No managed infrastructure.

Could you run them together?

Yes. Guardrails AI is useful for ensuring your LLM outputs meet your own quality standards. Guard is the compliance infrastructure that generates regulatory evidence.

A practical deployment would use Guardrails AI within your application to validate output quality and structure, and Guard as the inline proxy to generate per-call attestation and compliance evidence. Guard is managed infrastructure. Guardrails AI is a library.

Suggested next step

We can run a 30-minute proof of value showing Guard generating compliance evidence and attestation on the same AI calls where Guardrails AI is validating outputs. You will see the difference between output validation and regulatory compliance evidence.

Contact

hello@40south.au

40south.au

Sources

1. APRA Prudential Standard CPS 234 Information Security, July 2019. www.apra.gov.au/sites/default/files/cps_234_july_2019_for_public_release.pdf
2. Guardrails AI, "Introduction — Guardrails AI documentation." guardrailsai.com/guardrails/docs
3. Australian Government, "Voluntary AI Safety Standard," August 2024. www.industry.gov.au/publications/voluntary-ai-safety-standard
4. OAIC, "Chapter 8: APP 8 — Cross-border disclosure of personal information." www.oaic.gov.au/privacy/australian-privacy-principles/australian-privacy-principles-guidelines/chapter-8-app-8-cross-border-disclosure-of-personal-information
5. 40 South Guard technical documentation. 40south.au