

GuardrailProbe Independent Benchmark — August 2026

Run ID	ac387219-6d7c-408e-846e-174e9a0889b6
Timestamp	2026-08-04T12:57:34.193236+00:00
Generated	2026-08-04T13:05:59Z
Backends tested	nemo, guardrails_ai, presidio, lakera, openai_moderation, azure_content_safety, azure_prompt_shields, aws_bedrock
Best overall	llm_guard
Worst overall	llama_firewall

TL;DR

Winner	llm_guard
Best accuracy/latency ratio	guardrails_ai
Biggest improvement vs last month	—
Biggest regression vs last month	—
Backends tested	10
Backends skipped	0
Total probes run	707
Report generated	2026-08-04 12:57 UTC
Run ID	ac387219-6d7c-408e-846e-174e9a0889b6

Overall Comparison

Backend	Overall %	vs Last Month	Best Category	Worst Category	Avg Latency
nemo	79.5%	—	LLM01	LLM10	31454 ms
guardrails_ai	2.6%	—	LLM01	LLM05	0 ms
presidio	6.4%	—	LLM02	LLM01	974 ms
lakera	83.3%	—	LLM01	LLM10	584 ms
openai_moderation	100.0%	—	LLM01	LLM01	33422 ms
azure_content_safety	20.5%	—	LLM02	LLM01	1620 ms
azure_prompt_shields	24.4%	—	LLM01	LLM09	1613 ms
aws_bedrock	59.0%	—	LLM01	LLM10	660 ms
llama_firewall	11.5%	—	LLM01	LLM05	8043 ms
llm_guard	85.9%	—	LLM01	LLM10	2569 ms

Per-Category Results (OWASP LLM Top 10)

Ref	Description	Winner	Score	Runner-up	Score
LLM01	Prompt Injection	lakera	100%	aws_bedrock	100%
LLM02	Insecure Output	lakera	81%	llm_guard	81%

LLM03	Training Data Poisoning	lakera	100%	aws_bedrock	100%
LLM04	Model DoS	lakera	67%	llm_guard	67%
LLM05	Supply Chain	llm_guard	100%	lakera	80%
LLM06	Sensitive Info Disclosure	lakera	100%	llm_guard	100%
LLM07	Insecure Plugin	lakera	83%	llm_guard	83%
LLM08	Excessive Agency	lakera	100%	llm_guard	100%
LLM09	Overreliance	llm_guard	100%	nemo	100%
LLM10	Model Theft	lakera	20%	llm_guard	20%

Content Moderation Results

Backend	Hate	Violence	Sexual	Self-Harm	Overall CM
nemo	100%	100%	100%	100%	100%
guardrails_ai	0%	0%	0%	0%	0%
presidio	0%	0%	0%	0%	0%
lakera	100%	100%	100%	100%	100%
openai_moderation	0%	0%	0%	0%	0%
azure_content_safety	80%	100%	60%	80%	80%
azure_prompt_shields	0%	0%	0%	0%	0%
aws_bedrock	100%	100%	100%	100%	100%
llama_firewall	0%	0%	20%	0%	5%
llm_guard	100%	100%	100%	100%	100%

Backend Capability Matrix

Backend	Prompt Injection	Jailbreak	Content Mod.	PII Detection	Agentic Safety
NeMo Guardrails	Primary	Yes	No	No	Yes
GuardrailsAI	Yes	Yes	No	Yes	No
Presidio	No	No	No	Primary	No
Lakera Guard	Primary	Yes	No	No	No
Custom HTTP	Yes	Yes	No	No	No
OpenAI Moderation	No	Yes	Primary	No	No
Azure Content Safety	No	No	Primary	No	No
Azure Prompt Shields	Primary	Yes	No	No	No
AWS Bedrock	Yes	Yes	Yes	No	No
Llama Firewall	Primary	Yes	No	No	No
LLM Guard	Yes	Yes	Yes	Yes	No

Primary = core strength | Yes = supported | No = not designed for this

Accuracy vs Latency Tradeoff

Backend	Overall %	Avg Latency	Latency Category	Recommended For
guardrails_ai	2.6%	0 ms	Ultra-fast	Real-time, high-throughput pipelines
lakera	83.3%	584 ms	Moderate	Batch processing, async pipelines
aws_bedrock	59.0%	660 ms	Moderate	Batch processing, async pipelines
presidio	6.4%	974 ms	Moderate	Batch processing, async pipelines
azure_prompt_shields	24.4%	1613 ms	Slow	Offline analysis, compliance audits
azure_content_safety	20.5%	1620 ms	Slow	Offline analysis, compliance audits
llm_guard	85.9%	2569 ms	Slow	Offline analysis, compliance audits
llama_firewall	11.5%	8043 ms	Slow	Offline analysis, compliance audits
nemo	79.5%	31454 ms	Slow	Offline analysis, compliance audits
openai_moderation	100.0%	33422 ms	Slow	Offline analysis, compliance audits

Notable Bypasses

OWASP Category	Severity	Count
LLM04	high	1
LLM04	medium	1
LLM10	high	3
LLM10	medium	1

Backends Skipped This Month

Backend	Reason	Expected In
—	—	—

Month-over-Month Changes

First benchmark — no prior month comparison available.

How to Reproduce

```
pip install guardrailprobe
guardrailprobe run --year 2026 --month 8
```

Full guide: github.com/askuma/guardrailprobe/blob/main/METHODOLOGY.md

Independence Statement

GuardrailProbe is independently operated with no commercial relationship to any tested backend provider. All probes are open-source and methodology is publicly auditable. Test results reflect backend configurations at the time of testing only.