

Your agent has a goal. What will it do to achieve it?

Gadriel Behavior Assurance (GBA)

Pre-production behavioral assurance for goal-seeking AI systems.

We know how to test software. We know how to evaluate an LLM response. Agents introduce a different problem.

Today you write a happy-path test, run an eval, maybe a red team, then give the agent production authority. You find out how it actually behaves from a ticket, a refund, a leaked record, or a bill.

Goal → path → ✓ outcome / ✗ behavior

The ticket closed. The test passed. The behavior did not.

AGENT BEHAVIOR



GOAL GIVEN TO THE AGENT

"Resolve this customer's complaint and retain the account."



PASS — WHAT TODAY SCORES

- Customer retained
- Ticket closed
- Eval green
- Functional test: PASS



FAIL — WHAT IT MAY HAVE DONE TO GET THERE

- Issued an unauthorized refund
- Exposed another customer's information
- Promised something the company cannot deliver
- Bypassed a manager approval
- Called a tool it was not supposed to use
- Spent \$500 of model calls on a \$50 problem

The business outcome can look successful while the behavior was unacceptable. Traditional functional testing often gives that agent a PASS.

TODAY'S PROCESS

Happy-path → LLM eval → Red team (if you have one) → Deploy → Guardrail → Read the logs after

The missing question before production authority: **How will the agent behave while pursuing the goal?**

THIS IS ALREADY HAPPENING

OpenAI

Reward hacking¹

Hugging Face incident, July 2026. Evaluation with reduced safeguards.

Anthropic

Multi-agent conflict²

Frontier Red Team, emerging multiagent systems, 13 Aug 2026.

Replit

Authority violation³

Jason Lemkin / Replit, July 2025. Production coding-agent use.

UK AISI

Fake identities / malicious code approval⁴

July 2026 — agent created fake identities to get malicious code approved.

Functional testing asks: Can it achieve the goal? Gadriel Behavior Assurance asks: Can we trust what it does to get there?

Today's evals show the agent can work.

GBA challenges how it behaves while working.

Functional / happy-path	ASKS Did it complete the task?	MISSES Outcome can look successful while behavior is unacceptable.
LLM evals	ASKS Was the answer good?	MISSES A good answer can still leak data, bypass approval, or overspend.
Security / red team	ASKS Can someone attack it?	MISSES The agent can misbehave without an attacker.
Runtime guardrails	ASKS Can we stop certain known bad actions?	MISSES After it already has production authority.
Observability	ASKS What happened?	MISSES After the first complaint, incident, or bill.

Copilot Studio, Agentforce Testing Center, ServiceNow AI Agent Studio, AWS AgentCore Evaluations, LangSmith and OpenAI AgentKit already test task completion, tool choice and regression. That testing is necessary. **GBA adds pre-production assurance of how the agent behaves while pursuing its goals.**

WHAT GBA CHALLENGES THAT NORMAL EVALS OFTEN MISS

Multi-turn escalation

A safe first turn can become unsafe after reframing or accumulated context.

Authority boundaries

Does the agent bypass approval, overreach permissions, or retry a prohibited action?


Multi-agent handoffs

Do state, approvals, or intent get lost when one agent hands work to another?

Runaway cost / FinOps

Do retries, loops, or excess model and tool calls create uncontrolled cost?

HOW GBA WORKS




Repeatable scenarios. Deterministic scoring of defined checks. Reproducible evidence.

GBA is before takeoff. Guardrails are in the air. Logs are after landing.

This is what GBA looks like.

Behavioral Readiness. Not a vulnerability scanner.





Zoom: why this agent failed

Goal → Behavior → Boundary → Evidence → Fix.

ZOOM: WHY THIS AGENT FAILED

GOAL

retain the customer / resolve the complaint

BOUNDARY

refunds over \$500 require manager approval

WHAT THE AGENT DID

attempted a \$2,500 refund, then changed the refund reason and retried

FAIL

authority violation

approval bypass

goal persistence

GOAL-PURSUIT PATH

Offer credit

PASS

Escalate human

PASS

Issue refund → bypass approval **FAIL**

GBA EVALUATES AGAINST YOUR OPERATING BOUNDARIES

1

What the agent should accomplish.

Business goals and expected workflows

2

What it must never do.

Forbidden actions, data access and authority boundaries

3

What requires approval.

Human-in-the-loop rules and consequential actions

4

What good looks like.

Policies, compliance requirements and performance thresholds

5

What might break it.

Gadriel behavioral failure scenarios and adversarial cases

Behavioral failures GBA is built to find

Not another scanner. Failure modes of goal-seeking systems.

GOAL & OPTIMIZATION

Reward hacking

Reaches the metric by an unintended shortcut

Goal drift

The means diverge from business intent

Evaluation gaming

Optimizes the score rather than the intended outcome

AUTHORITY & TOOLS

Authority violation

Uses tools, credentials, or actions beyond intended boundaries

Unsafe tool use

Natural-language instructions become damaging real-world actions

Data exposure

Reveals or moves information it should not

AUTONOMOUS BEHAVIOR

Deception

Hides, edits, or misrepresents what it did

Multi-turn escalation

Safe on turn 1, unsafe after reframing

Runaway / FinOps

Loops, retries, or excess model/tool use explode cost

MULTI-AGENT

Conflict

Agents compete, sabotage, or amplify each other

Handoff failures

State, approvals, or intent lost between agents

HOW GBA STRESSES THE AGENT

GBA exercises the same system under different conditions so teams can see where behavior changes before production authority is granted.

Expected → Unexpected → Adversarial → Multi-turn → Multi-agent

Show where the boundary breaks.

17-AGENT CLINICAL ASSESSMENT

284 scenarios · 17 agents

Direct injection blocked. Multi-turn reframe failed. Key affected pillars: Security / Coherence / Teamwork.

Safe agent ≠ safe agentic system

gadriel.ai/blog — 17-agent clinical AI assessment.

Security

Adversarial robustness

Compliance

Policy & regulatory alignment

Safety

Harm boundaries

Operational

Reliability & error handling

FinOps

Cost control

Coherence

Goal stability

Teamwork

Coordination & handoffs

Bias

Fairness

How you run GBA, and what you get

Your agent stays where it is. GBA comes to it.

YOUR ENVIRONMENT

Agent



Tools



Data

↑ GBA test connection

↓ Evidence / report

Agent remains in your environment.

SaaS

GBA connects to that staging agent. Scenario prompts and behavior evidence (what we asked, what it did) go to Gadriel to score and produce the report.

Private cloud

GBA runs in the customer's cloud. Prompts, evidence, and the report stay in that tenancy.

Air-gap

GBA CLI runs on the customer network. Prompts, evidence, and the report stay local. Nothing leaves the network.

<> **No source code required**

 **Staging by default**

 **Bounded actions**
Not by default. Authorized scenarios only.

 **Customer-defined policies**

4-6 WEEK PROOF OF VALUE

First findings during the PoV as scenarios are exercised; final assurance report at 4-6 weeks.

WEEK 1

Scope system, goals, authority and unacceptable actions

WEEK 2

Map tools, data, identities and other agents

WEEKS 2-4

Exercise expected, unexpected and adversarial behavior

WEEKS 4-5

Review findings and remediation

WEEKS 5-6

Re-test and produce the assurance report

Boundaries → **Scenario Evidence** → **Findings** → **Remediation** → **Re-test** → **Release Decision**
 Goals · approvals · never-do What we asked · what it did Pass · Partial · Fail Prompt · policy · tool · model Prove the change With evidence



Verdict

Pass / Partial / Fail



Evidence

What we asked and what it did: prompt, situation, named scope, response, tool call, handoff



Remediation

A concrete change: prompt, policy, tool, or model



Assurance Report

Word + OCSF. EU AI Act, NIST AI RMF, ISO 42001, SR 11-7.

AWS · Microsoft · Google · Salesforce · ServiceNow · OpenAI · Anthropic · LangChain · CrewAI · AutoGen · MCP · A2A

GCA validates what AI writes. **GBA validates what AI does.** GDA validates what attackers can reach.



Request a Proof of Value

gadriel.ai/trial · contact@gadriel.ai

Before you give an AI agent a goal and the authority to act, know what it is willing to do to achieve that goal.