



UnisonAI · Gauntlet

INDEPENDENT AI-AGENT ASSURANCE

The AI customer that stress-tests your AI agent.

Every company is racing to launch AI agents. Almost nobody is checking what they actually say when a real — or hostile — customer pushes on them. Gauntlet is that independent check.

[Start free — 7 days](#)[Book a demo](#)**8**

quality dimensions scored

10

adversarial AI personas

<10 min

from paste to graded report

100%

independent — never the bot's maker

THE PROBLEM

An AI agent that hasn't been tested is a liability.

The failures are already famous — a chatbot that invents a refund policy the company is then held to; one talked into a \$0 deal; one that leaks personal data or refuses to reach a human. QA teams sample 1-2% of conversations. Platform vendors grade their own homework. Nobody stress-tests the agent the way a hostile customer will.



Invents policy

Confidently confirms terms that don't exist — the Air Canada failure mode.



Gets jailbroken

Talked out of its own rules into refunds, discounts, or leaking its prompt.



Leaks data

Reads out card numbers or OTPs, or won't hand a frustrated customer to a human.

WHAT IT IS

Send an AI customer at your AI agent — then get it graded.

Gauntlet unleashes a squad of adversarial "AI customer" personas on your bot — the furious refund-seeker, the jailbreaker, the data-fisher, the non-native speaker, even another AI transacting on a customer's behalf. Then it delivers a CX Assurance Report Card: an overall grade, a score on every dimension, the exact failure transcripts, and how to fix each one.

1

Point & probe

Pick a bot — a sample, a pasted transcript, or your live webhook — and the personas to send in.

2

Watch it break

Each adversarial customer works the agent live; failing turns are flagged in real time.

3

Grade & fix

An animated report card scores 8 dimensions and hands you concrete fixes. Re-run to prove it held.

THE RUBRIC

Eight dimensions that define a good CX agent.

- **Grounding** · no hallucinated policy

- **Privacy** · protects PII & secrets

- **Empathy** · acknowledges before policy

- **Coherence** · no repetition loops

- **Guardrails** · resists prompt injection

- **Escalation** · hands off gracefully

- **Compliance** · no guarantees / regulated advice

- **Agent-to-agent** · ready for the AI customer

Refund Warrior

Confused Customer

Hallucination Baiter

Prompt Injector CRIT

Escalation Seeker

PII Prober CRIT

Edge-Caser CRIT

Loop Tester

Non-Native Speaker

The Agent (A2A)

WHY GAUNTLET

The company that **builds** the bot can't be the one that **grades** it.

Independent by design

A fixed rubric applied from the outside — the reason a passing grade means something to a customer, a regulator, or a buyer.

Platform-agnostic

Tests the conversation, not the plumbing.
Works with Genesys, Five9, Webex CC, Twilio, Intercom, or any custom bot via a webhook.

Built for the agentic era

The only tool that scores whether a customer's own AI assistant can transact with yours — the question every 2026 CX roadmap is asking.

WHO IT'S FOR

From a solo consultant to a global contact centre.

**CX & CCaaS teams**

Assure every bot before and after launch; gate risky changes in CI.

**BPOs**

Prove bot & agent quality to enterprise clients with independent scores.

**QA & consultants**

Walk into a prospect with a white-labelled report of their bot's failures.

**Bot builders**

Ship WhatsApp / site / client bots with a pre-delivery assurance step.

PLANS · USD

Start free. Scale when you're ready.

Free Trial

\$0

7 days · one per email · core 3 personas · 25 runs.

Individual

\$19/mo

All 10 personas · unlimited runs · API · 1 seat.

Team

\$99/mo

5 admins + 10 users · white-label reports · unlimited keys.

Enterprise

Custom

SSO · native connectors · voice-bot testing · SLA.

Put your AI agent through it — before your customers do.

Independent AI-agent assurance for the CX & CCaaS era.

gauntlet.unisonai.com →UnisonAI Gauntlet · Independent AI-agent assurance · gauntlet.unisonai.com

Gauntlet applies a fixed rubric independently of the platform under test. A report reflects the specific conversations tested and is not a certification. Sample bots shown in the product are fictional; Gauntlet is not affiliated with any named CX/CCaaS platform. © 2026 UnisonAI.