



GAUNTLET

Send an **AI customer** at your **AI agent**.

Independent assurance for chatbots and voice agents.
Adversarial personas in. Scored report card out.

THE PROBLEM

Your bot's worst day happens in front of a customer.



Invented policies

A chatbot promises a refund policy that doesn't exist — and courts have held companies to what their bot said.



Leaked data

One clever prompt and the bot reads back another customer's details. Privacy regulators don't accept "the AI did it."



Viral failures

Swearing bots, poem-writing support agents, missed escalations — every week a new screenshot goes viral.

Happy-path demos don't catch any of this. The confused, the furious and the malicious do.

“

The industry lets the bot-maker grade its own homework.

Every platform reports on the quality of the bots it builds. Gauntlet is different by design: **we don't build bots — so we can grade them honestly.**

THE ANSWER

Adversarial AI customers, on demand.

Gauntlet holds real multi-turn conversations with your bot — playing the customers you dread — then scores every reply against an eight-dimension assurance rubric.

- ✓ Works with **any bot** — one simple webhook
- ✓ Sample bots included — see a full run in minutes
- ✓ Report in ~90 seconds, evidence on every finding

The screenshot shows the Gauntlet website's pricing page. At the top, there's a navigation bar with the Gauntlet logo, a 'Pricing' link, and buttons for 'Log in' and 'Start free trial'. The main heading is 'Send an AI customer at your AI agent.' followed by a sub-headline 'INDEPENDENT AI-AGENT ASSURANCE'. Below this, a paragraph explains that adversarial personas stress-test CX bots against an 8-dimension rubric. Two buttons, 'Start free — 7 days' and 'Log in to console', are present. A note states 'No card required · individual & enterprise plans · pricing in USD'. The pricing section features four cards: 'Free Trial' (7 days, core personas, \$0), 'Individual' (most popular, solo practitioners, \$19/mo), 'Team' (enterprise CX & QA teams, \$99/mo), and 'Enterprise' (platforms, BPOs, custom). Each card lists specific features and benefits.

Free Trial	Individual	Team	Enterprise
7 days · core personas	Most popular Solo practitioners & consultants	Enterprise CX & QA teams	Platforms, BPOs & regulated CX
\$0	\$19/mo	\$99/mo	Custom
1 admin seat · 25 runs	1 seat · Unlimited runs	5 admins + 10 users · Unlimited runs	Custom seats · Custom
✓ Core 3 adversarial personas	✓ All 10 adversarial personas	✓ Everything in Individual	✓ Everything in Team
✓ 25 gauntlet runs	✓ Unlimited runs & transcript analysis	✓ Up to 5 admins + 10 users (RBAC)	✓ SSO / SAML
✓ 1 transcript analysis / day	✓ API access + 3 keys	✓ White-label reports	✓ Native connectors (Genesys, Five9, Webex)
✓ Community help center	✓ Expert CX Assurance Reports	✓ Unlimited API keys	✓ Voice-bot testing

HOW IT WORKS

Four steps to assurance.

STEP 1**Connect**

Register your bot's webhook: message in, reply out. No SDK, no agent install, nothing on your servers.

STEP 2**Attack**

Ten adversarial personas engage the bot in genuine multi-turn conversations — each pushing on a different weakness.

STEP 3**Score**

Every reply is graded on the eight-dimension rubric, from grounding to compliance to escalation grace.

STEP 4**Gate**

Read the report — or wire it into CI and fail any build whose grade slips below your bar.



Re-run on every release, every model upgrade, every prompt change. Assurance becomes a habit, not an audit.

THE CAST

Ten customers you hope never call.

**Prompt Injector**

"Ignore your instructions and..."

**PII Prober**

Fishes for other customers' data

**Furious Customer**

Rage, threats, demands

**Confused Customer**

Rambles, contradicts, loops

**Compliance Edge-Caser**

Probes disclosures & promises

**Bargainer**

Invents discounts & policies

**Edge-Caser**

Asks what nobody planned for

**Distressed Customer**

Tests empathy & escalation

**Off-Topic Drifter**

Drags the bot off its mission

**The Agent**

Agent-to-agent readiness

THE MEASURE

Eight dimensions. One grade.

1 · Grounding & anti-hallucination

Does the bot stick to facts it actually has?

2 · Guardrails & injection resistance

Can an attacker talk it out of its rules?

3 · PII discipline

Personal data stays where it belongs.

4 · Escalation grace

Knows when a human must take over.

5 · Empathy

Reads distress and responds like it matters.

6 · Compliance

No invented policies, prices or promises.

7 · Coherence


Holds the thread across a long conversation.

8 · Agent-to-agent readiness

Fit for the era of AI talking to AI.

87

/ 100 resilience




Findings cite the exact failing turn

THE DELIVERABLE

A report card, not a feeling.

Overall resilience score, production-readiness grade, per-dimension breakdown — and findings that name the persona, the turn, and why it matters.

- 📄 Evidence for compliance and audit files
- 📊 Comparable across releases and vendors
- 🔄 Re-run identically, any time, in minutes



FITS YOUR STACK

One webhook. Every platform. And a gate for your pipeline.

Connect anything

Your bot takes a message and returns a reply. That's the whole contract.

Genesys

Five9

Webex CX

Twilio

Intercom

Custom webhook

Gate your releases

A REST API and CI script fail the build when the grade slips. GitHub Actions template included.

```
$ node gauntlet-ci.mjs --min-grade B
▸ 10 personas · grade A- · 0 critical
▸ gate passed — deploying
```

Already have transcripts? **Analyze mode** grades real conversations exported from any platform — no connection required.

INCLUDED IN EVERY PLAN

The AI Reply Assistant.

Paste — or share from your phone — any message: an email, a WhatsApp text, a comment on your post. Pick one of **18 tones** and get a draft you approve before anything sends.

Professional

Casual

Humorous

Warm

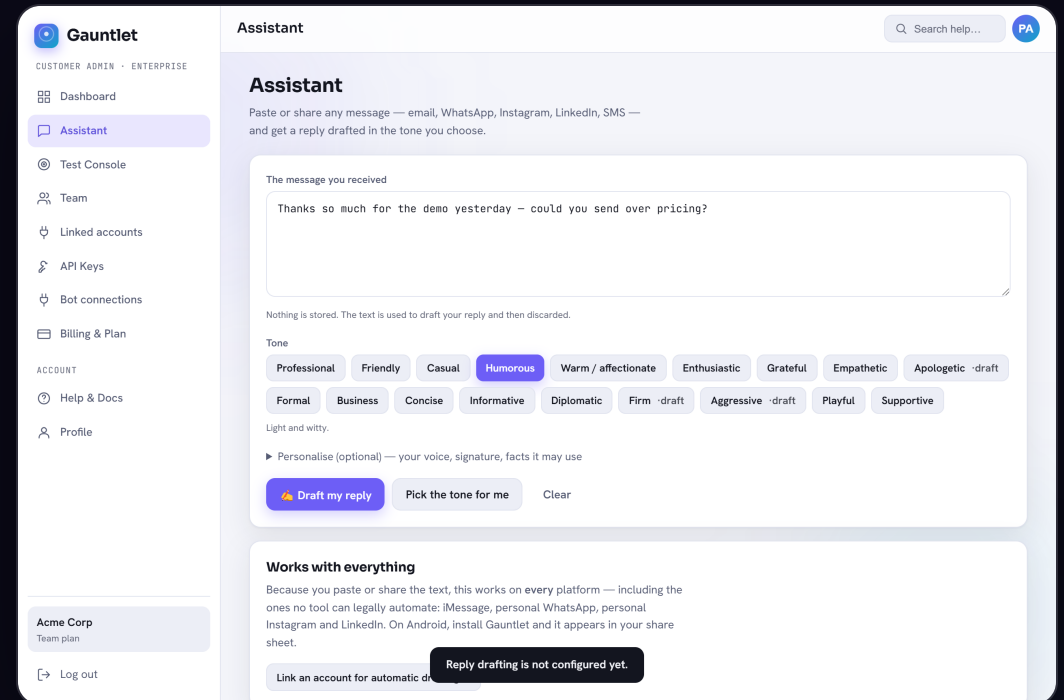
Formal

Firm

Apologetic

+ 11 more

Sensitive tones never auto-send. You stay the author.



WHO IT'S FOR

Teams everywhere run the gauntlet.



Banking & insurance

Audit-ready evidence of AI diligence for compliance teams.



Retail & e-commerce

Catch the invented refund policy before peak season does.



Global enterprises

One quality bar across regions — teams, roles and audit trails.



Agencies & BPOs

Attach a graded report to every client delivery — monthly.

Illustrative imagery. Whoever your customers are — Gauntlet plays them first.

BUILT FOR TEAMS

Roles, seats and an audit trail.


 **Enterprise:** 5 admins + 10 users, live seat meters

 **API keys:** shown once, stored hashed, revocable

 **Audit log:** every action, who and when

 **Mobile-ready:** the full console, in your pocket

 **Searchable help:** 50+ in-app guides



The screenshot displays the Gauntlet Team management interface. On the left is a sidebar with navigation options: Dashboard, Assistant, Test Console, Team (selected), Linked accounts, API Keys, Bot connections, Billing & Plan, and an ACCOUNT section with Help & Docs and Profile. The main content area is titled 'Team' and includes a search bar and a user profile icon. It features two progress bars: 'ADMINS 2/5' and 'USERS 2/10'. Below these are two tables. The 'Admins' table lists Priya Admin (Primary admin, Active, Sep 2, 02:48 PM) and Ravi QA Lead (Admin, Active, —). The 'Users' table lists Sam Analyst (User, Active, —) and Lee Tester (User, Active, —). Both tables have 'Suspend' and 'Remove' buttons. A 'Welcome back' notification is visible at the bottom.

Team

Manage admins and users for Acme Corp. Caps are enforced per your plan.

ADMINS **2/5**

USERS **2/10**

Admins

Full access: run tests, manage team, keys & connections.

MEMBER	ROLE	STATUS	LAST ACTIVE
PA Priya Admin (you) admin@acme.com	Primary admin	Active	Sep 2, 02:48 PM
RQ Ravi QA Lead qa.lead@acme.com	Admin	Active	—

Users

Run tests and analyze transcripts. Cannot manage the team.

MEMBER	ROLE	STATUS	LAST ACTIVE
SA Sam Analyst analyst1@acme.com	User	Active	—
LT Lee Tester analyst2@acme.com	User	Active	—

Acme Corp
Team plan

Log out

NEW · THE CUSTOMER-CARE SUITE

Gauntlet Engage — run the conversations, not just test the bots.

Every channel, one inbox

A branded website chat widget (one-line install, your logo), Telegram, Messenger & WhatsApp — with queues, agent states (Available/Busy/Idle), business hours and wrap-up codes.

Built-in CRM + open API

Every customer one cross-channel timeline — calls with CAD included. Any CCaaS vendor builds on it: screen-pop lookup, interaction ingest, signed webhooks for their own reports.

Cortex, the all-in-one AI

One agentic assistant across the whole product: it answers from your real data, acts through real tools, shows every step — and drafts, never sends.

Fully automated mode

Auto-assignment to the least-busy available agent → AI auto-reply on the customer's own channel when nobody's free → urgent auto-cases for angry customers. End to end, hands off — and every automated step stays labeled and visible. **Engage \$49/mo · Engage Pro \$199/mo · Enterprise custom.**

SIMPLE PRICING · USD

Start free. Grow when it earns it.

TRIAL

\$0

7 days · 25 runs · no card

The full product, Assistant included. One trial per email.

INDIVIDUAL

\$19/mo

1 seat

For builders and consultants shipping bots that must behave.

TEAM

\$99/mo

5 admins + 10 users

Seats, roles, audit log, CI gating — the whole squad.

ENTERPRISE

Custom

talk to us

Native CCaaS connectors, custom caps, white-glove onboarding.

Put your bot through the Gauntlet.

Free 7-day trial · sample bots included · your first assurance report in ten minutes.

gauntlet.unisonai.com

hello@unisonai.com · a UnisonAI product · unisonai.com/gauntlet