UnisonAI Trust pack

Security, data handling, accessibility and use of AI in Gauntlet and UniCast (public beta). Updated 31 October 2026. The same content is published at www.unisonai.com/trust. Questions: hello@unisonai.com.

1. Trust centre
2. Security practices
3. Your data and our sub-processors
4. Accessibility statement
5. How we use AI

1. Trust centre

How we protect your data and your customers' conversations in Gauntlet and UniCast, what is true today, and what is not done yet.

What we protect

Gauntlet and UniCast hold your organisation's conversations with customers, your contacts and cases, your bot's knowledge, your videos and webinars, and the keys you give us to connect your channels. Your customers trust you with that information; you are trusting us with it.

How, in short

Each organisation kept apart

Every record belongs to one organisation. Automated checks try to read another organisation's data every 10 minutes and must be refused.

Secrets stay secret

Passwords and API keys are stored only as one-way hashes. The keys you give us for your channels are encrypted with a key unique to your organisation and are never shown back to anyone, including our staff.

You control access

Single sign-on, automatic user provisioning, roles with per-module permissions, and an audit log your admins can read.

We test ourselves

Our own platform runs 100+ automated checks every 10 minutes, and a full regression run every night.

Honest status

Gauntlet and UniCast are in public beta. We are a small, founder-run company. We do not yet hold a third-party security certification, we have not yet had an independent penetration test, and the service runs in a single hosting region. Our data processing agreement is a draft that is being finalised. We say this up front so you can decide with the full picture; each page below lists what is in place today and what is not.

In this trust centre

- Security practices: sign-in, encryption, access control, isolation, logging and what we do not have yet.
- Your data and our sub-processors: where data is processed, retention, erasure, export and the companies that process data for us.

- Accessibility statement: our WCAG 2.1 AA aim, captions and keyboard access, check results and known gaps.
- How we use AI: disclosure, hand-off to a person, answers from your own sources, video checks and how we test our own bot.
- Trust pack (PDF): all of the above in one document for your vendor file.

For a security questionnaire or a data-protection review, write to hello@unisonai.com. We answer questionnaires ourselves and will tell you plainly when the answer is "not yet".

2. Security practices

What the product does today to keep your organisation's data safe, and the gaps we have not closed yet.

Everything on this page describes what the product does today. If it is not built, it is in the "not yet" list at the end.

Signing in

- Passwords are stored only as salted one-way hashes (bcrypt). Changing a password needs the current one and is recorded in the audit log.
- Sign-in attempts are rate-limited per network address to slow down password guessing.
- Console sessions expire after 12 hours. Removing a person through single sign-on or automatic provisioning ends their sessions straight away.
- **Single sign-on** for organisations: SAML 2.0 and OpenID Connect, with domain verification. Admins can require single sign-on for everyone, keeping named emergency ("break-glass") admins.
- **Automatic provisioning** with SCIM 2.0: people and groups are created, updated and removed from your identity provider, and groups can map to roles.
- Sign-in with an existing work account (Google, Microsoft, LinkedIn, GitHub) for individuals and small teams. An existing account is linked only when both sides have confirmed the e-mail address.

Keeping secrets and data safe

- The credentials you give us to connect channels and tools are encrypted with AES-256-GCM using a key unique to your organisation. That key is itself encrypted with a master key kept outside the database, so a copy of the database alone does not reveal them.
- Secrets are decrypted only at the moment we call the outside service. Nothing in the product shows a stored secret back to a person, including our own administrators.
- API keys are stored only as hashes; you see the full key once, when you create it.
- All traffic to the service is encrypted in transit (HTTPS).
- Messages we receive from connected services (webhooks) are checked against a signature before they are accepted.

Who can see what

- **Organisation isolation:** every record belongs to one organisation, and every request is checked against the signed-in person's organisation.
- **Roles and permissions** per module (none, read or write). In the CRM a role can see all records, its team's, only its own, or look up one record at a time; fields can be hidden, masked or read-only. Look-ups are rate-limited.
- Exports, imports, deletions and bulk changes are written to an access log, and field changes keep a history.
- **Audit log:** security events such as password changes and support access are recorded, and your admins can read your organisation's log.

- **Our own support access** to your organisation is time-limited (5 minutes to 8 hours), needs a written reason, is recorded in the audit log and is visible to you. Stored secrets stay unreadable even then.

Video and webinars (UniCast)

- Videos are private by default. Share links can have an expiry date, a password and a list of websites allowed to embed them; setting a video back to private turns every link off.
- Video files are delivered through short-lived signed links, not permanent public addresses.
- Nothing is published until its content check passes or a reviewer in your organisation approves it.
- **Authenticity seal:** each organisation has its own digital signing key. A published video is sealed, and anyone can check online that it came from you and has not been altered since.

Payments

- Card details are handled only by our payment provider; we never see or store card numbers. Payment notifications are checked against the provider's signature, and payment actions are rate-limited.

Testing, backups and changes

- **Continuous testing:** 100+ automated checks run against the live service every 10 minutes, using test organisations only. They cover sign-in, roles, keys, isolation between organisations, the audit log, billing and every product area, and alert us when something fails.
- A full regression run, including tests too risky for the live service, runs every night on a throwaway server.
- A consistent copy of the database is taken before every release, without stopping the service.
- Changes to sign-in, access control and payments get a separate review before release, and fixes from those reviews are made before the change ships.

What we do not have yet

- No third-party security certification (for example ISO 27001 or a SOC 2 report) yet.
- No independent penetration test yet.
- A single hosting region; no second region to fail over to.
- No built-in two-step verification for password sign-in yet. Organisations that need it can require single sign-on and use their identity provider's multi-factor sign-in.
- Public beta: we are a small team, and there is no round-the-clock security operations centre.

Found a security problem? Write to hello@unisonai.com with the details. Please give us a reasonable time to fix it before you share it publicly.

3. Your data and our sub-processors

Where data is stored and processed, how long we keep it, how to erase or export it, and who processes it for us.

> **Status:** the facts on this page come from our Data Processing Agreement, draft v1 of 21 October 2026. It is being finalised with our lawyer and is not yet a binding document. Ask us for the current copy.

Who is responsible

You, the customer, are the **controller** of the personal data you put into Gauntlet. UnisonAI is the **processor**: we process it only on your documented instructions, which are the way you configure the service and our Terms.

What we process for Gauntlet

ITEM	GAUNTLET
Purpose	AI agent assurance testing; customer-care inbox, CRM, AI drafting, reports
Data subjects	Customer's staff; customer's end-users who contact them
Categories of data	Names, contact details, conversation content, case notes, usage data
Duration	For the subscription term; deleted or returned within 30 days after termination unless law requires longer

Where data is stored and processed

Service data (servers, storage and backups) is hosted in: **India (Mumbai)**. Some sub-processors below work in other locations; each one's location is listed.

International transfers: where personal data from the EU or UK is processed in India or by sub-processors outside the EU or UK, we rely on the European Commission's Standard Contractual Clauses (Module 2, controller-to-processor) and the UK Addendum. India's DPDP Act obligations apply to processing in India.

Retention, erasure and return

- **Retention:** we keep your data for the subscription term. After it ends, we delete or return it within 30 days, unless the law requires us to keep it longer.
- **Erasure:** the CRM "Erase" function removes a person and all their linked records, so you can answer erasure requests yourself.
- **Requests from people:** we help you answer requests for access, correction and erasure.
- **Return:** at the end of the service we delete or return your personal data, as you choose.

Our commitments as processor

- Our staff and contractors are bound by confidentiality.
- We use only the sub-processors listed below and give you 30 days' notice of any change, so you can object.
- We tell you without undue delay after we become aware of a personal-data breach, with a target of 48 hours.
- We share the information you need to show compliance and allow audits on reasonable notice (or provide third-party audit reports once available).
- Personal data is never used to train our own AI model; that training uses only curated Q&A pairs that are filtered to remove personal data.

Sub-processors

Sub-processors, as listed in Annex A of the DPA draft v1

SUB-PROCESSOR	PURPOSE	LOCATION	DATA
Hostinger International	VPS hosting of Gauntlet, FolioLyt and unisonai.com (servers, storage, backups)	India (Mumbai)	All service data
Google (Gmail SMTP)	Sending transactional email, alerts and newsletters	Global	Name, email address, message content
Hostinger Mail	Mailboxes hello@unisonai.com / hello@foliolyt.cloud	Global	Correspondence
Razorpay Software Pvt Ltd	Payments, subscriptions, invoices/receipts (card data handled only by Razorpay)	India	Name, email, phone, billing country, payment metadata
Groq, Inc.	AI inference (chat answers, drafting, triage)	USA	Conversation text sent for an answer
Mistral AI	AI inference (fallback)	EU (France)	Conversation text sent for an answer
Google (Gemini API)	AI inference (fallback)	Global	Conversation text sent for an answer
Cloudflare, Inc. (Workers AI)	AI inference (fallback)	Global	Conversation text sent for an answer
UnisonAI (own model on the same server)	AI inference on our own server — no third party	Same as hosting	Conversation text
Kaggle (Google)	Training our own model on curated, non-personal Q&A pairs	Global	No personal data (filtered before upload)

SUB-PROCESSOR	PURPOSE	LOCATION	DATA
Let's Encrypt (ISRG)	TLS certificates	USA	Domain names only

For the full DPA, or to be told about sub-processor changes, write to hello@unisonai.com.

4. Accessibility statement

Our aim for the Gauntlet app, the UniCast video player and the webinar room, what we have checked, and the gaps we know about.

Our aim

We aim to meet the **Web Content Accessibility Guidelines (WCAG) 2.1, level AA** for the Gauntlet app, the UniCast video player (shared video and playlist pages, webinar registration, certificates and seal checks) and the webinar room. We do not yet meet it everywhere; the known gaps are listed below. We have not had an independent accessibility audit.

Features that help

- **Captions** on recorded videos, made automatically, which you can correct.
- **Translated captions**, so viewers can follow in another language, and a searchable transcript next to the video.
- **Live captions in webinars**, with live translation for viewers, and "listen in my language" (clearly labelled as an AI voice).
- **Resizable caption text** (small, medium, large) and a show / hide captions switch.
- **Screen-reader announcements:** live captions and the room's status are announced politely through live regions, without stealing focus.
- **Keyboard access:** the video player uses the browser's own controls, which work with a keyboard; page controls are standard buttons, links and form fields.
- Works in any modern browser on a computer, tablet or phone, with no horizontal scrolling on small screens, and nothing to install.

How we checked

On 31 October 2026 we ran an automated check (the open-source axe-core rules for WCAG 2.0 and 2.1, levels A and AA) on 14 views: the six UniCast public pages, the sign-in and sign-up pages, and six screens of the signed-in app including the webinar room. Automated checks find only part of the possible barriers, so we also need, and have not yet completed, a full manual keyboard and screen-reader review.

Automated check results, 31 October 2026

AREA	RESULT
UniCast video, playlist, webinar registration, joining, certificate and seal-check pages (6 views)	No failures found
App home, UniCast library, chatbot and inbox screens	No failures found
Sign-in and sign-up pages	1 issue: some links in sentences are shown by colour only

AREA	RESULT
Video page in the app	2 issues: one link shown by colour only; the transcript panel scrolls but cannot be reached with the keyboard
Webinar room (host view)	1 issue: low colour contrast on some participant buttons (such as "Remove")

Known gaps

- The four issues in the table above; we are fixing them.
- Live captions need the speaker to use Chrome, Edge or Safari; on other browsers the room says so and captions are not produced.
- Captions and translations are made by AI and can contain mistakes until someone corrects them.
- The website chat window and the Agent Desktop have not yet been checked in this round.
- No manual screen-reader test (for example with NVDA, JAWS or VoiceOver) has been completed yet.

Report a barrier

If something in Gauntlet or UniCast is hard or impossible for you to use, write to hello@unisonai.com with the page, what you were trying to do, and the browser or assistive technology you use. A person reads every message; we reply and tell you when it is fixed. If you need the content of a video in another format, we will help you or the organisation that shared it.

These features help organisations meet duties such as the UGC guidelines in India, ADA Title II in the United States and the European Accessibility Act. Using Gauntlet or UniCast does not on its own make an organisation compliant.

5. How we use AI

Where AI is used in Gauntlet and UniCast, how people stay in control, and how we test our own bot.

AI helps in Gauntlet and UniCast, but people stay in charge: your team decides what the bot knows, can take over any conversation, and approves anything that a check holds back.

Your website chatbot

- **Answers from your own sources:** the bot answers from the knowledge you give it (your website pages, files and answers you write). It is told never to invent prices, features, dates, discounts or commitments, and to say honestly when a person will follow up instead.
- **A real person, one click away:** visitors can ask for a person at any time. The bot also hands over when it cannot answer, when a message is very negative, or on words you choose. After hand-off the bot stays silent; your team sees why it was handed over and what the bot used.
- **No person free?** Visitors can leave a voice message or ask for a call-back.
- **Disclosure:** AI replies appear under your bot's name, separate from your team members' names, and the visitor is told when they are back with the assistant. You choose the bot's name and welcome message; we recommend making clear that it is an automated assistant.

Testing bots before they speak

Gauntlet sends hostile AI "customers" at a bot (trying to extract personal data, slip it new instructions, refuse it a person, or make it promise things it should not) and gives a report card with the exact sentence that failed. **We test our own website bot with Gauntlet too.**

Video: checked before anyone watches

- Every UniCast video is checked before it can be published: frames for nudity, and the title, description, transcript, comments and live chat for hate, harassment, violence, self-harm and other harmful content.
- The rules run first; AI can only make a verdict stricter, never more lenient. A video that could not be checked is never passed automatically; it waits for a person in your organisation to review it.

Captions and translation

- Transcripts, captions, translated captions and live webinar captions are made by AI. They can contain mistakes, and you can correct captions.
- "Listen in my language" is labelled as an AI voice, so viewers know it is not the speaker.

Your data and AI

To answer, the text of a conversation is sent to an AI provider or to our own model. The providers and their locations are listed on the data and sub-processors page. Personal data is never used to train our own model.

Rules these features help with

Gauntlet and UniCast **help you meet** these rules; they do not make you compliant on their own, and the decisions stay with you:

- **India's DPDP Act:** personal-data tests for your bot, erasure of a person's records, and a processor agreement.
- **EU AI Act transparency duties for chatbots:** your bot's name and welcome message are where you tell visitors they are talking to an AI, and Gauntlet tests that the bot hands over to a person when asked and does not invent facts.
- **European Accessibility Act, ADA Title II (United States) and UGC accessibility guidelines (India):** captions, live and translated captions and an accessible player. See our accessibility statement for known gaps.