Security Whitepaper

Enterprise Security Overview - AI Interview Analyzer

Provider:	AI Interview Analyzer Sp. z o.o.
Address:	ul. Jedrusik 6/53, 01-748 Warszawa, Poland
NIP:	5253079974
REGON:	54402118500000
Classification:	PUBLIC
Date:	24.06.2026

Contents

Security Whitepaper

Provider: AI Interview Analyzer Sp. z o.o., Warszawa, Poland

Audience: Enterprise security, IT, and procurement teams

Classification: Public

1. Executive Summary

AI Interview Analyzer is an enterprise hiring platform that records interviews with explicit candidate consent, transcribes and structures them, and produces evidence-based evaluation support for recruiters. Because the platform handles candidate personal data and supports hiring processes, security and privacy are treated as primary design constraints, not features added later.

This whitepaper describes, in concrete and verifiable terms, how we protect customer and candidate data. It is written for the people who review vendors: security engineers, IT administrators, data protection officers, and procurement. Every figure in this document is drawn directly from our own engineering systems rather than from marketing material.

The central message is simple: **we do not merely assert that the platform is secure, we continuously test that it is.** Our codebase contains **3,171 automated tests**, including a dedicated security suite that exercises authentication, authorization, cross-organization isolation, injection defenses, and data deletion. On top of that, we run a repeatable penetration-testing harness against live deployments and produce written audit reports. Across seven internal security audits in March and April 2026, we recorded **zero critical findings**, with our most recent audit closing at a verdict of **PASS**. (Formal third-party certification of these controls is on our roadmap; see Section 18.)

Security characteristic	Summary
Hosting	Microsoft Azure, EU regions only
Network model	Private endpoints, default-deny network segmentation, no public database
Encryption	AES-256 at rest, TLS 1.2 or higher in transit
Identity	Short-lived signed tokens, bcrypt password hashing, SSO support
Access control	Role-based access control with strict per-organization isolation
Secrets	Centralized secrets vault with managed-identity access
Privacy	Explicit consent, configurable retention, single-unit erasure
Responsible AI	Decision support only, human always in the loop
Assurance	3,171 automated tests plus recurring penetration tests and audits

1.1 How to Read This Document

Sections 3 to 11 describe the controls that protect data: architecture, network, identity, application, data protection, privacy, and the secure development lifecycle. Sections 12 and 13 cover our distinctive continuous-testing program and our audit history. Sections 14 to 17 cover operations, threat modeling, vulnerability management, and compliance mapping. The appendices provide a control catalog, a reviewer FAQ, and a glossary that a security team can use directly during an assessment.

2. Document Scope and Approach

2.1 What This Document Covers

This whitepaper covers the security architecture and practices of the AI Interview Analyzer service: the hosting environment, network design, identity and access management, application-level controls, data protection, privacy and regulatory alignment, the secure development lifecycle, and our continuous security testing program.

2.2 What Makes It Verifiable

Vendor security claims are easy to write and hard to trust. We have therefore tied each major claim in this document to something concrete and countable inside our engineering systems: a control implemented in code, a test that proves the control works, an infrastructure definition that enforces it, or an audit report that records a documented check. Where a control is part of our forward roadmap rather than shipped today, we say so explicitly. We would rather under-claim and be trusted than over-claim and be caught.

2.3 Shared Responsibility

The platform is delivered as software as a service. We operate the infrastructure, application, AI pipeline, and data handling. The customer is responsible for managing their own user accounts and roles, configuring data-retention windows to match their internal policy, and ensuring that candidate consent is obtained through the consent workflow the platform provides. Section 14 describes this division in more detail.

3. Security Architecture Overview

The platform is built as a small number of cooperating services rather than a single monolith. A desktop application and a web portal act as clients. A central backend API owns all persistence, authentication, billing, the AI pipeline, consent, email, file handling, and dashboards. An audio merge worker processes recordings asynchronously. All sensitive state lives behind the backend API; clients never talk directly to the database, storage, or AI services.

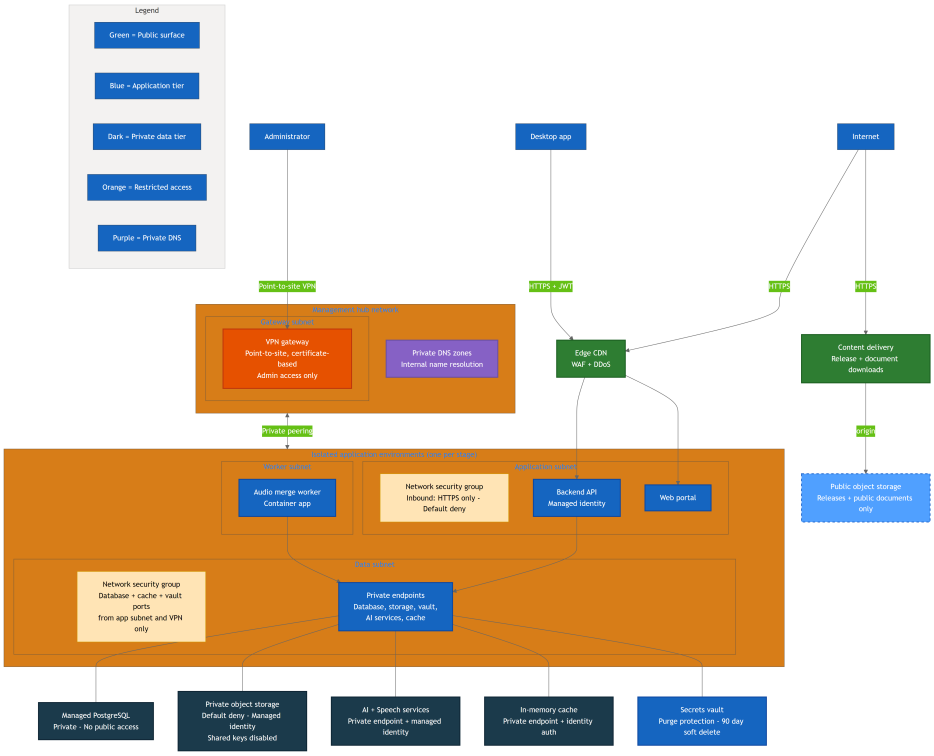

The diagram above shows the production topology with resource names intentionally generalized. Three principles are visible in it:

- **No direct exposure of data services.** The database, private object storage, AI services, and cache have public network access disabled and are reachable only through private endpoints inside an isolated virtual network. The secrets vault is reached by the application over a private endpoint and is additionally protected by platform identity authentication and least-privilege access policies, so any access requires a valid, authorized identity regardless of network path.
- **A separated public surface.** The only public object storage holds release downloads and public documents. It never contains candidate data. Customer-facing application traffic passes through an edge layer that provides web application firewall, distributed-denial-of-service protection, and content delivery.
- **Administrative access is gated.** Operators reach internal resources only through a certificate-based point-to-site VPN into a management hub network, not over the public internet.

Each deployment stage (development and production) is a fully isolated environment with its own network, storage accounts, database, and secrets. Customer production data is never present in lower environments. A shared management hub holds only the VPN gateway and private DNS, peered privately to each environment.

4. Defense in Depth

No single control is trusted to stop every attack. The platform layers independent controls so that the failure of any one layer does not expose data. The layers below are each implemented and, as described in Section 12, individually tested.

Layered security model: independent controls at every tier

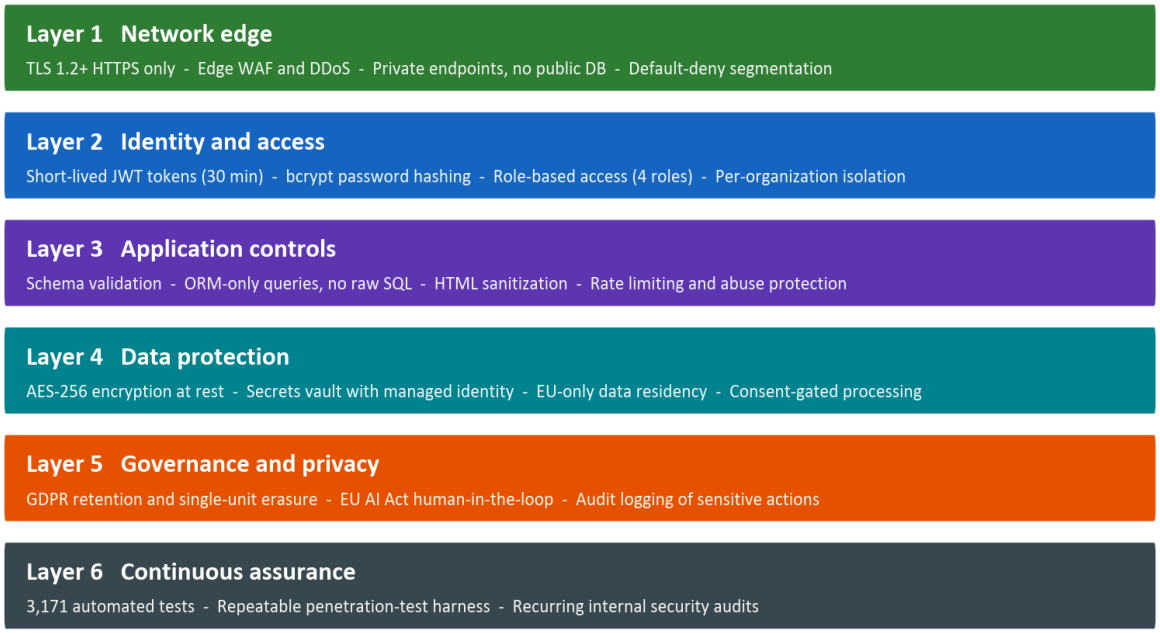

Layer	Representative controls
Network edge	TLS-only transport, edge WAF and DDoS protection, private endpoints, default-deny segmentation
Identity and access	Short-lived signed tokens, bcrypt hashing, role-based access control, per-organization isolation
Application	Schema validation on all input, ORM-only data access, output encoding, rate limiting
Data protection	Encryption at rest, secrets vault with managed identity, EU data residency, consent-gated processing
Governance and privacy	Configurable retention, single-unit erasure, human-in-the-loop AI, audit logging
Continuous assurance	Automated test suite, repeatable penetration tests, recurring internal security audits

The remainder of this document walks through each layer in turn and then describes how we prove, continuously, that the layers hold.

5. Network Security

5.1 Private by Default

The data tier is private by construction. The managed PostgreSQL database has public network access disabled and is reachable only through a private endpoint. Private object storage is configured to deny network access by default, disables shared access keys entirely, and is accessible only via managed identity from the application subnet. The cache, AI services, and secrets vault are likewise reached through private endpoints with private DNS resolution.

In practice this means there is no internet-facing connection string to the database and no public storage URL for candidate audio: the database and private storage have public network access disabled outright. The secrets vault is reached by the application over a private endpoint and is protected by platform identity authentication and least-privilege access policies, with application identities granted read-only access to only the secrets they need, so secrets cannot be retrieved without a valid, authorized identity. The attack surface that an external adversary can even touch is limited to the application's HTTPS endpoints behind the edge layer.

5.2 Network Segmentation

Each environment is divided into separate subnets for the application tier, the data tier, and the asynchronous worker. Every subnet is governed by a network security group whose final rule denies all inbound traffic. The application subnet accepts only inbound HTTPS. The data subnet accepts only the specific database, cache, and vault ports, and only from the application subnet or the administrative VPN. This means that even an attacker who somehow reached the application tier cannot freely pivot to the data tier; the only permitted paths are the ones the application legitimately uses.

5.3 The Edge

Public application traffic is fronted by an edge layer providing a web application firewall, DDoS protection, and a content delivery network. Release and document downloads are served from a dedicated public storage account through a content-delivery front door, completely separate from the private storage that holds candidate data. The two storage planes never mix: a misconfiguration on the public plane cannot expose private candidate data, because they are different accounts with different network rules.

5.4 Administrative Access

There is no public administrative endpoint into the private network. Operators connect through a point-to-site VPN gateway that uses certificate-based authentication. Administrative database and cache access is only possible from inside that tunnel, since those services have public network access disabled. This keeps day-to-day operations off the public internet entirely.

6. Identity and Access Management

6.1 Authentication

User sessions are established with a signed access token that is valid for thirty minutes, paired with a separate, opaque, server-side refresh token. Access tokens are verified on every request, and the user is re-validated against the database (including an active-account check) rather than being trusted on the token contents alone. Logging out revokes the server-side refresh session immediately, so a stolen refresh token cannot outlive a logout.

Passwords are never stored in plain text. They are hashed with bcrypt using a unique per-password salt. For organizations that prefer single sign-on, the platform supports OAuth login with Microsoft and Google, in which case no password is held at all.

Email ownership is verified through a single-use, time-limited verification link before a self-registered account is treated as verified, and verification-email resends are rate limited to prevent abuse.

6.2 Role-Based Access Control

Authorization is enforced through a role model with four roles of increasing privilege: interviewer, hiring manager, recruiter, and administrator. Access to privileged operations is enforced by server-side dependencies that check both the role and the verification status of the caller. These role checks guard well over one hundred distinct API operations.

Role	Typical capabilities
Interviewer	Conducts assigned interviews; sees only interviews assigned to them
Hiring manager	Manages recruitments they own or are a member of
Recruiter	Full recruitment and candidate management within the organization
Administrator	Organization settings, billing, user and API-key administration

Beyond coarse role checks, the platform applies data-level visibility rules. Hiring managers see only the recruitments they created or are members of; interviewers see only the interviews assigned to them. Privilege is therefore enforced both at the level of "what action" and at the level of "which records."

6.3 Per-Organization Isolation

The platform is multi-tenant, and tenant isolation is treated as a first-class security control. Every authenticated identity carries an organization identifier, and data queries are scoped to that organization. When a user requests a record that belongs to another organization, the platform returns a "not found" response rather than revealing that the record exists. Internal database identifiers are never exposed on the wire; the API presents display identifiers and re-maps them per request, which removes a common class of cross-tenant enumeration attack.

This is not only a design intention. As described in Section 12, our automated suite runs a large cross-organization matrix that attempts to reach one organization's data using another organization's credentials and asserts that every such attempt fails.

6.4 Programmatic Access

For integrations, organizations on eligible plans can issue API keys. Keys use a recognizable prefix, carry 128 bits of entropy, and are stored only as a hash; the raw key is shown once at creation and never again. Each key carries an explicit permission scope (read, write, or ATS integration), can be restricted to specific source networks, can be revoked instantly, and is subject to per-key rate limits derived from the organization's plan tier. Key verification uses a timing-safe comparison to avoid leaking information through response timing.

7. Application Security

The application is written to remove entire categories of vulnerability rather than patch them case by case.

- **Injection.** All database access goes through an object-relational mapper with parameterized queries. The codebase contains no raw string-formatted SQL. This structurally eliminates SQL injection.
- **Input validation.** Every request body is validated against a strict schema before it reaches business logic. Oversized payloads are rejected, and list endpoints are paginated to bound resource use.
- **Output encoding and cross-site scripting.** User-supplied and AI-generated text is treated as untrusted. Where content must be rendered as HTML, it passes through an allow-list sanitizer at write time, and a dedicated test suite confirms that script tags, event handlers, and javascript URLs are stripped.
- **Mass assignment.** Update operations use explicit schemas that exclude privileged fields such as role, organization, and credit balance, so a client cannot escalate privilege by posting extra fields.
- **Rate limiting.** Authentication and abuse-prone endpoints are rate limited using a durable, database-backed limiter that survives restarts and works correctly across multiple application instances. Login, registration, password reset, and verification resends each have their own limits. Client IP resolution is hardened against spoofing of forwarding headers.
- **Webhooks.** Inbound webhooks from payment and email providers are verified against provider signatures on the raw request body before being processed.
- **File uploads.** Uploads are size-capped, validated, stored under generated identifiers rather than user-supplied names, and constrained per request and per organization.
- **Security headers.** In production, responses carry strict transport security, content-type and frame options, a referrer policy, and a restrictive permissions policy, and suppress server and framework banners.

8. Data Protection

8.1 Encryption

All data is encrypted at rest using AES-256 through the Azure platform storage and database encryption layers. All network traffic is served exclusively over HTTPS using TLS 1.2 or higher; plaintext HTTP is redirected to HTTPS at every tier. In production, the API and web portal emit strict transport security headers along with a set of hardening headers, and suppress server and framework version banners.

8.2 Secrets Management

Application secrets are held in a centralized secrets vault with purge protection enabled and a ninety-day soft-delete window. Applications authenticate to Azure resources using system-assigned managed identities rather than long-lived keys; for example, private storage has shared access keys disabled entirely, so access is only possible through identity-based role assignments scoped to the individual resource. Vault access policies grant application principals read-only access to the specific secrets they need, following least privilege.

8.3 Data Residency

All customer and candidate data is stored and processed within the European Union. Application hosting, the database, storage, cache, and secrets reside in West Europe, and AI processing runs in EU regions. The AI provider does not use customer data to train its models.

8.4 The Life of a Single Interview

The clearest way to understand the data-protection controls is to follow one interview from end to end. Consent is captured and recorded before anything is processed. The upload is encrypted in transit. Transcription and analysis run inside EU data centers. Results are written to encrypted storage. Every record is then governed by a single retention clock that ends in a logged, cascading deletion. At any point, candidate rights such as withdrawal, deletion, access, or portability can interrupt this flow.

9. Privacy by Design and GDPR

Privacy is built into the data model and the workflow, not bolted on through policy alone.

9.1 Consent

No interview is recorded or analyzed without the candidate's explicit consent. When a candidate is added to a recruitment, the platform issues a unique, single-use consent link by email. The candidate reviews what will happen and either accepts or declines. Consent state, including the time of response, is recorded against that specific recruitment, so consent is always scoped to a concrete hiring process rather than granted globally.

Candidate consent: explicit and recorded before any processing

9.2 Retention and Erasure

Data retention is configurable per organization, with a default of twelve months and a configurable minimum of thirty days, and can be overridden per candidate. There is a single retention clock for a candidate's data, not a separate timer per artifact. The clock starts when a hiring decision is recorded. Before data expires, the platform sends a warning (by default about fifteen days ahead) and offers a one-click extension. When data is deleted, it is deleted as a single unit: the candidate record, interviews, transcripts, audio recordings, documents, and comparisons are all removed together, and the deletion is recorded in an audit log. There is no partial or orphaned residue.

The lifecycle below shows this single clock and how it converges on one cascading deletion with a logged proof of erasure.

Data retention: one clock per candidate, single-unit deletion

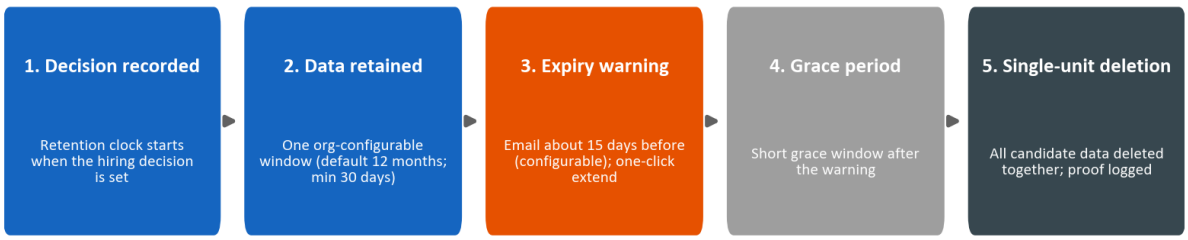

9.3 Data Subject Rights and Sub-processors

The platform supports the data-subject rights required under the GDPR, including access, deletion, portability, objection, and explanation. Processing is carried out under a data processing agreement that customers accept at registration and that is versioned per organization. Our sub-processors and their roles, all within the EU or under appropriate safeguards, are disclosed in that agreement, and customers receive advance notice of any change. Section 17 contains the sub-processor register and the

article-by-article compliance mapping.

10. Responsible AI and the EU AI Act

The platform falls within the high-risk category of the EU AI Act because it supports employment decisions, and we treat that classification seriously.

The defining rule of the product is that **the AI is decision support, not a decision maker.** The system never automatically accepts or rejects a candidate. It transcribes speech, structures questions and answers, scores answers against criteria the recruiter has defined, and drafts feedback, and a human reviews every output before it is used. This keeps a human firmly in the loop.

Equally important is what the AI does not do. It does not evaluate personality, "cultural fit," emotional state, tone of voice, accent, gender, age, ethnicity, appearance, or body language. Scoring is anchored to evidence from the transcript and to recruiter-defined criteria, and candidate names are excluded from the evaluation input to reduce bias. We publish a transparency card, user documentation, and a declaration of conformity describing the system, its limitations, and its safeguards.

Responsible-AI control	How it works
Human in the loop	Every score and every piece of feedback is reviewed by a recruiter before use
No automated decisions	The system never auto-accepts or auto-rejects a candidate
Evidence-based scoring	Scores reference supporting evidence from the transcript
Anti-bias design	Names excluded from evaluation; substance is scored over style
Scope limits	Personality, emotion, accent, and protected characteristics are never assessed
Candidate feedback safety	Private candidate feedback passes a generation-and-validation safety guardrail

These constraints are not only stated in documentation; they are encoded in the AI prompt layer and exercised by a dedicated AI-safety test program described in Section 12.3.

11. Secure Development Lifecycle

Security is enforced in the way we build and ship software, not only in the running system.

- **Environment separation.** Development and production are fully separate, each with its own infrastructure, storage accounts, database, secrets, and subdomains. There is no shared state.
- **Infrastructure as code.** The entire cloud environment is defined as code and reviewed as code, which makes the security posture auditable and reproducible. A reviewer can read exactly which ports are open, which resources are private, and which identities have which permissions.
- **Pinned, gated deployments.** Every step in the continuous-integration pipeline is pinned to an exact, immutable version. Production deployments are tag-based, run only through the protected production pipeline, and are gated behind required approval. The automated test suite runs as a release gate: a deployment cannot ship if tests fail.
- **Dependency hygiene.** Automated dependency monitoring proposes updates weekly across the backend, desktop, web, infrastructure, and pipeline definitions, and dependency audits are part of our periodic security review.
- **Signed artifacts.** Desktop installers are code-signed, so customers can verify that the software they install genuinely comes from us.
- **Secrets discipline.** Secrets live in the vault and in protected pipeline secrets, never in source code.

12. Continuous Security Testing

This is the heart of our assurance story and the part most vendors cannot show. We treat security as something to be measured continuously, with executable checks, rather than asserted once.

12.1 The Automated Test Suite

The platform is covered by **3,171 automated tests** spanning the backend API, the desktop application, the web portal, the browser extension, and the audio merge worker.

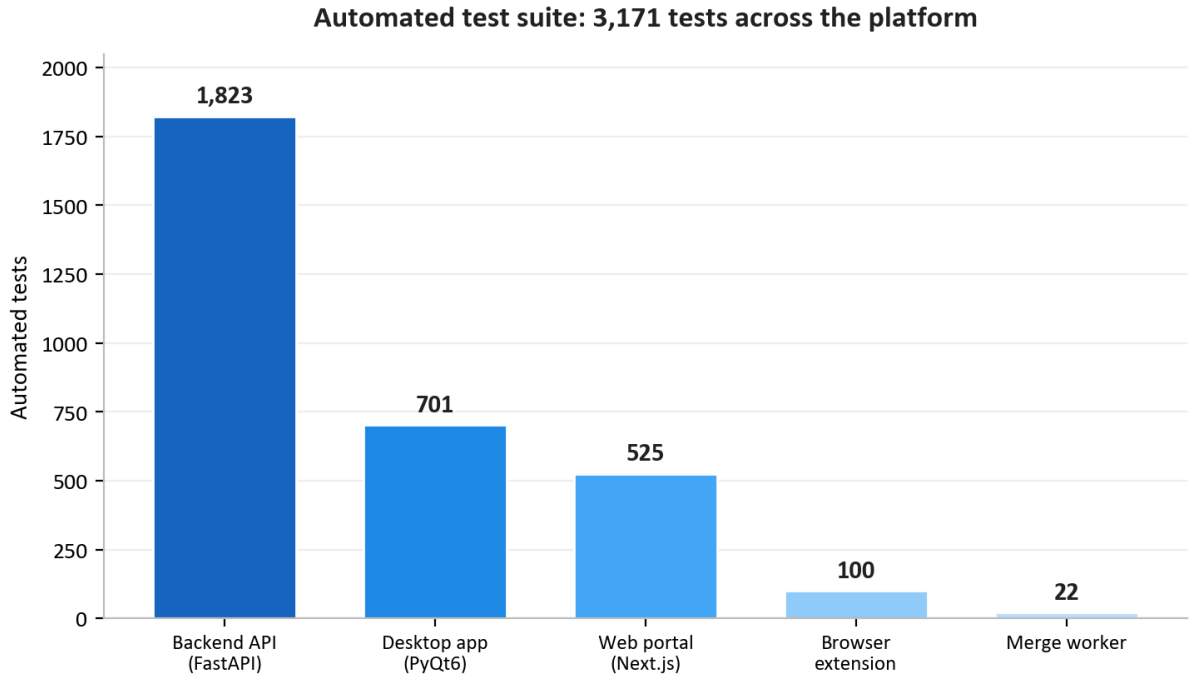

These are not only functional tests. A substantial, dedicated security suite exercises the controls described earlier in this document. The chart below breaks down the security-specific tests in the backend API by domain.

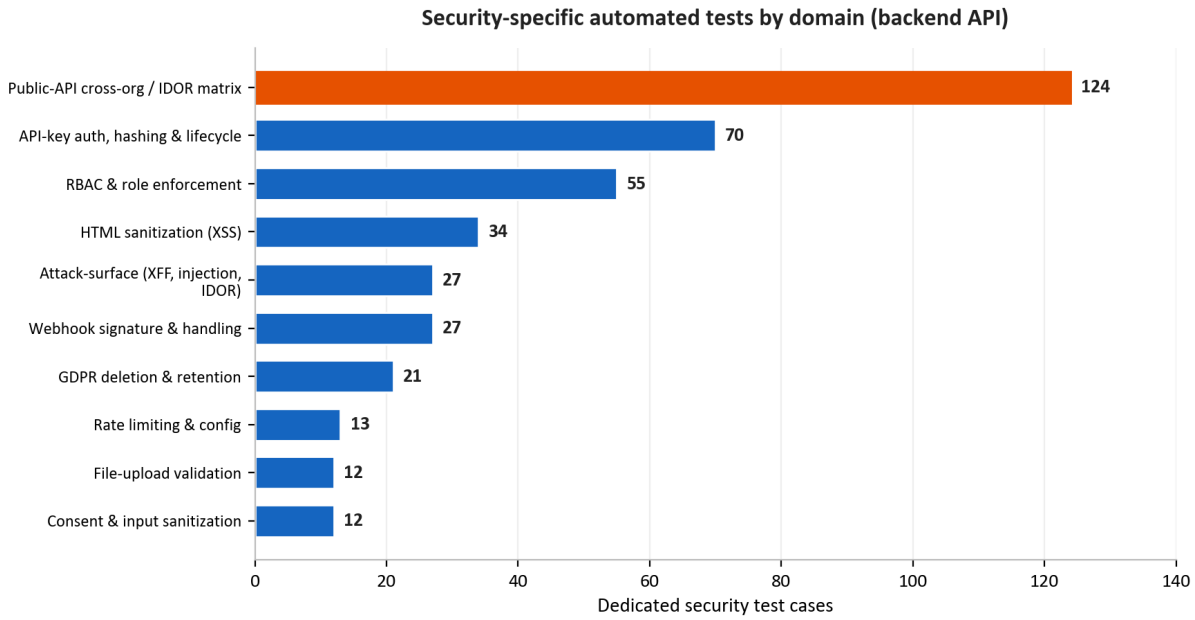

Among many others, this suite includes a large public-API matrix that runs each endpoint as a legitimate user, as the organization's own API key, and as a rival organization's API key, asserting that every cross-organization attempt is blocked. It includes dozens of adversarial attack-surface tests for forwarding-header spoofing, header injection, and identifier leakage, a focused HTML-sanitization suite for cross-site scripting, role-enforcement tests for the full role model, and tests that prove candidate data is genuinely deleted as a unit. Because these tests run as a release gate, a regression that weakened any of these controls would stop the release rather than reach customers.

12.2 Live Penetration Testing

Automated unit tests prove that controls behave correctly in isolation. To prove they hold together in a real deployment, we maintain a repeatable penetration-testing methodology that runs real attack scripts against a live environment. It is organized into six phases:

Phase	Focus	Examples of what is exercised
1. Static analysis	Source code	Secrets, injection patterns, dangerous functions, missing auth, unsafe HTML
2. Architecture review	Infrastructure	Private endpoints, segmentation, TLS, secrets configuration
3. Attack-vector analysis	Source control and cloud	Branch protection, identity scope, public exposure
4. Live penetration testing	Running environment	Unauthenticated probing, cross-org access, injection, token tampering, SSRF, rate-limit bursts
5. Enterprise scoring	Maturity	Sixteen security categories scored against an enterprise baseline
6. Dependency and supply chain	Third-party risk	Dependency CVE audit, pinned pipeline actions, lock-file integrity

Phase 4 is genuine adversarial testing against a deployed system, not a checklist. It probes protected endpoints without credentials and confirms they refuse access; it registers two organizations and attempts to reach one organization's records with the other's account; it injects cross-site-scripting and server-side-template payloads and confirms they are neutralized; it tampers with authentication tokens and confirms they are rejected; it attempts server-side request forgery against cloud metadata endpoints; and it bursts authentication endpoints to confirm that rate limiting actually triggers in the live environment, not only in theory.

12.3 Candidate-Feedback Safety Testing

Because the platform can generate private development feedback for candidates, we run a separate adversarial safety program against that feature. It deliberately feeds the system harsh and hostile recruiter notes and confirms that the candidate-facing output never contains vulgarity, never reveals or attributes a recruiter's identity or private opinion, and never applies judgmental personality labels. This protects both the candidate, who should receive constructive and respectful feedback, and the customer, who should never have an internal opinion leak outward.

13. Security Audit Results

We conduct recurring security audits using a structured, repeatable penetration-testing methodology, and write each one up as a dated report with severity-rated findings, evidence, and remediation. These are internal audits run by our own security process; formal third-party certification of the same controls is on our roadmap. Between late March and late April 2026 we completed **seven such audits** across development and production.

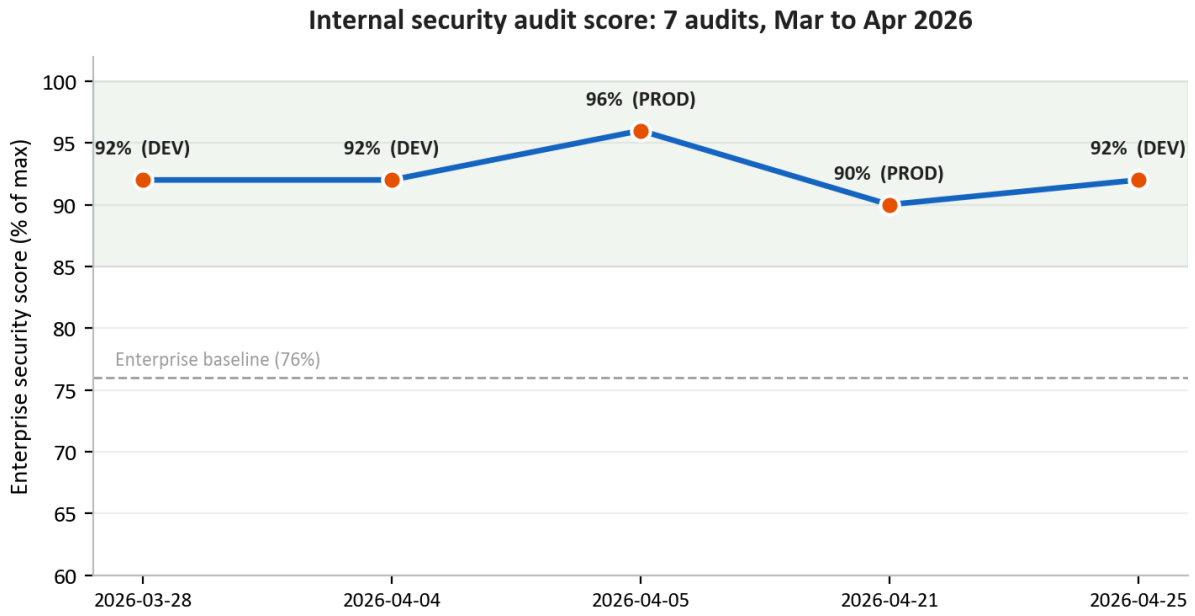

The result that matters most to a prospective customer is the consistency: **across all seven audits there were zero critical findings.** On the rare occasions a higher-severity issue surfaced, it was remediated quickly, often the same day, and re-verified. The scoring rubric was deliberately tightened over this period (the maximum possible score was raised as we added more categories to assess), which is why the normalized score line stays high even as the bar moved up.

Our most recent audit, on 25 April 2026, illustrates how the process works in practice. Two higher-severity issues were identified, both were fixed and re-verified the same day, and the audit closed at a verdict of **PASS** with no exploit-ready issues remaining in the current threat model.

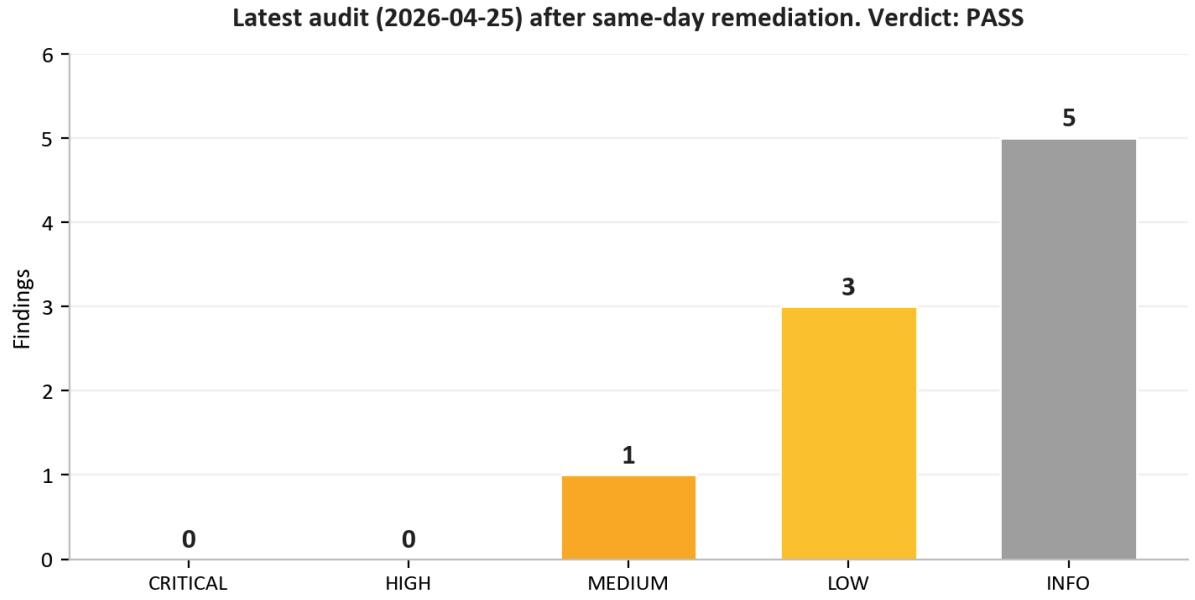

Audit	Environment	Critical	Verdict
2026-03-28	Development	0	Ready for production
2026-04-04	Development	0	Enterprise-ready
2026-04-05	Production	0	Enterprise-ready
2026-04-20	Development	0	Production-ready, notes
2026-04-20	Development	0	Pass with notes
2026-04-21	Production	0	Safe, no exploitable findings
2026-04-25	Development	0	Pass

The pattern across these audits is the most honest evidence we can offer: issues are found, because we look hard for them, and they are closed quickly, because the process is built to close them. A vendor that never reports a finding is usually a vendor that is not looking.

---

14. Operational Resilience and Shared Responsibility

14.1 Monitoring and Logging

Application and platform telemetry flow into a centralized log analytics workspace and application-monitoring service, giving us visibility into availability and behavior. Sensitive actions such as data deletion, legal-agreement acceptance, and AI invocations are recorded in dedicated audit tables, so there is a durable record of who did what to important data.

14.2 Backup and Recovery

The managed database retains automated backups, and private storage is protected by soft-delete retention on both blobs and containers, so accidental or malicious deletion can be recovered within the retention window. Critical infrastructure carries deletion locks to prevent accidental teardown of production resources.

14.3 Shared Responsibility Summary

Area	AI Interview Analyzer	Customer
Infrastructure, network, patching	Yes	-
Application security and AI pipeline	Yes	-
Encryption, secrets, data residency	Yes	-
User and role administration	Provides the controls	Manages users and roles
Retention policy configuration	Provides the controls	Sets the retention window
Candidate consent	Provides the workflow	Ensures it is used
Strong end-user credentials and SSO	Supports SSO and policy	Enforces internal policy

15. Threat Model and OWASP Mapping

We design against a concrete set of adversaries: an external attacker with no credentials, a curious or malicious authenticated user of one organization trying to reach another organization's data, a compromised dependency, and an insider mistake. The table below maps the widely used OWASP Top 10 risk categories to the specific controls that address them in this platform, each of which is exercised by the testing described in Section 12.

OWASP risk	How the platform mitigates it
Broken access control	Role-based access control on every privileged endpoint; per-organization scoping; "not found" on cross-org access; identifier remapping; cross-org test matrix
Cryptographic failures	TLS 1.2+ in transit; AES-256 at rest; bcrypt password hashing; secrets in a managed vault
Injection	ORM-only parameterized queries; strict schema validation; write-time HTML sanitization
Insecure design	Layered defense in depth; threat modeling and architecture review in every audit
Security misconfiguration	Infrastructure as code; default-deny network groups; security headers; disabled shared storage keys; API schema not exposed in production
Vulnerable components	Weekly automated dependency monitoring; dependency CVE audits in periodic review
Identification and authentication failures	Short-lived tokens; rate-limited login; email verification; SSO support; no plaintext passwords
Software and data integrity failures	Pinned, immutable pipeline steps; signed desktop installers; webhook signature verification; tag-gated production deploys
Security logging and monitoring failures	Centralized telemetry; dedicated audit tables for sensitive actions
Server-side request forgery	Outbound calls restricted to trusted endpoints; SSRF probes in the penetration-test harness

This mapping is the spine of our assurance argument: for each well-known class of attack there is a named control, and for each named control there is a test.

16. Vulnerability Management and Responsible Disclosure

Security is never finished, so we run a continuous loop of discovery and remediation.

- **Discovery.** Vulnerabilities are surfaced from four sources: the automated test suite, the recurring penetration-test audits, automated dependency monitoring, and reports from customers or researchers.
- **Triage.** Each finding is assigned a severity (critical, high, medium, low, or informational) with evidence and a remediation owner, exactly as recorded in our audit reports.
- **Remediation targets.** Critical and high findings are prioritized for immediate remediation; in our audit history, higher-severity findings have typically been resolved and re-verified the same day. Medium and lower findings are scheduled into the regular maintenance cadence.
- **Verification.** Fixes are re-tested, and where relevant a live check is run against the deployed environment to confirm the issue is genuinely closed, not just closed in code.
- **Disclosure.** Security concerns can be reported to us directly. We acknowledge reports, investigate, and keep the reporter informed through to resolution.

17. Compliance Mapping

17.1 GDPR

GDPR area	Platform implementation
Lawful basis (Art. 6)	Explicit candidate consent captured before processing
Data minimization and storage limitation (Art. 5)	Only interview-relevant data is processed; configurable retention with automatic deletion
Right to erasure (Art. 17)	Single-unit deletion of all candidate data, with a logged proof of erasure
Data subject rights (Art. 15 to 20)	Access, deletion, portability, and objection are supported
Processor obligations (Art. 28)	Data processing agreement accepted at registration and versioned per organization
Security of processing (Art. 32)	Encryption, access control, isolation, and continuous testing as described in this document
Sub-processor transparency	Disclosed in the data processing agreement with advance notice of change

17.2 EU AI Act

The platform is treated as a high-risk AI system supporting employment decisions, and we maintain documentation aligned to the regulation, including a transparency card, user documentation, and a declaration of conformity. The core safeguards, human oversight, transparency, evidence-based scoring, and strict scope limits on what the AI evaluates, are described in Section 10. We continue to mature our formal conformity documentation as the regulation's implementation timeline advances.

17.3 Hosting Certifications

The platform runs entirely on Microsoft Azure, whose data centers carry independent certifications including ISO 27001 and SOC 2. These certifications cover the physical and platform layers beneath our application; the application-layer controls are the ones described throughout this document.

17.4 Sub-processor Register

Sub-processor	Purpose	Region
Microsoft Azure	Hosting, AI and speech processing, storage, transactional email	EU (West Europe, Sweden Central)
Stripe	Subscription and payment processing	EU (Ireland)
Fakturownia	Invoicing	EU (Poland)
ATS connector (optional)	Applicant-tracking integration, enabled only on request	EU

18. Security Roadmap

We treat security as a continuously improving program. Current initiatives on our roadmap include strengthening multi-factor authentication options for administrative accounts, expanding centralized audit logging of data access, continuing to tighten dependency currency on a regular cadence, and progressing formal third-party certification of the controls described in this document. None of these is a gap that exposes customer data today; each is an enhancement to an already-layered posture.

19. Summary

AI Interview Analyzer protects candidate and customer data through a layered architecture: a private-by-default network with no public data services, strong identity and per-organization isolation, application code that designs out whole vulnerability classes, encryption and EU data residency, and privacy controls built into the data model. What distinguishes the platform is the evidence behind those claims. With 3,171 automated tests, a repeatable live penetration-testing methodology, a dedicated AI-safety program, and a track record of seven internal security audits with zero critical findings, we can show, not just say, that the platform is secure.

Appendix A: Security Control Catalog

A condensed reference of primary controls and the evidence that backs each one.

Control	Mechanism	Evidence
Transport encryption	HTTPS only, TLS 1.2+, HTTP redirected	Infrastructure as code; architecture audit
Encryption at rest	AES-256 platform encryption on storage and database	Platform configuration; architecture audit
Password protection	bcrypt with per-password salt	Source control; authentication tests
Session management	30-minute signed tokens, revocable server-side refresh	Source control; authentication tests
Authorization	Four-role access control on privileged endpoints	Role-enforcement test suite
Tenant isolation	Per-organization query scoping; 404 on cross-org	Cross-organization test matrix
API key security	Hashed storage, scoped permissions, per-key rate limits	API-key test suite
Injection defense	ORM-only parameterized queries	Static analysis; injection tests
Cross-site scripting defense	Write-time HTML sanitization	HTML-sanitization test suite
Rate limiting	Durable database-backed limiter on auth endpoints	Rate-limit tests; live burst checks
Webhook integrity	Provider signature verification on raw body	Webhook test suite
Secrets management	Managed vault, purge protection, managed identity	Infrastructure as code; architecture audit
Network isolation	Private endpoints; default-deny segmentation	Infrastructure as code; architecture audit
Data deletion	Single-unit cascade deletion with audit log	GDPR deletion test suite
Supply chain	Pinned pipeline steps; weekly dependency monitoring	Pipeline configuration; dependency audit

Appendix B: Frequently Asked Questions for Security Reviewers

Where is our data stored? Entirely within the European Union, on Microsoft Azure, in West Europe with AI processing in EU regions. Candidate data never leaves the EU.

Is our data used to train AI models? No. The AI provider does not use customer data for training.

Is the database reachable from the internet? No. Public network access is disabled and the database is reachable only through a private endpoint inside the virtual network.

Can one customer see another customer's data? No. Every query is scoped to the caller's organization, cross-organization access returns "not found," and an automated matrix continuously tests this isolation.

How are passwords stored? Hashed with bcrypt and a unique per-password salt. Single sign-on with Microsoft and Google is supported, in which case no password is stored.

Do you support single sign-on? Yes, via Microsoft and Google OAuth.

How long are access tokens valid? Thirty minutes, paired with a revocable server-side refresh session that is invalidated on logout.

How is candidate consent handled? Each candidate receives a unique, single-use consent link and must accept before any recording or analysis. Consent is recorded against the specific hiring process.

How is data deleted? As a single unit covering the candidate record, interviews, transcripts, audio, documents, and comparisons, on a configurable retention schedule, with a logged proof of erasure. Candidates can also request deletion directly.

Do you have a data processing agreement? Yes, accepted at registration and versioned per organization, including the sub-processor register.

Does the AI make hiring decisions? No. It provides decision support only; a human reviews every output and makes all decisions.

How do you prove your security claims? Through 3,171 automated tests including a dedicated security suite, a repeatable six-phase penetration-testing methodology run against live environments, an AI-safety test program, and recurring written audit reports.

What happens when you find a vulnerability? It is assigned a severity with evidence and an owner, remediated on a priority schedule, re-verified including live checks where relevant, and recorded in an audit report.

Can we run our own penetration test? Security assessments can be arranged through your account representative under appropriate scope and scheduling.

Appendix C: Glossary

Term	Meaning
AES-256	A strong symmetric encryption standard used to protect data at rest
bcrypt	A purpose-built password-hashing function with per-password salting
Managed identity	A platform-issued identity that lets a service authenticate without stored keys
Private endpoint	A private network address that keeps a cloud service off the public internet
Network security group	A set of allow and deny rules that filter network traffic to a subnet
RBAC	Role-based access control, granting permissions according to a user's role
IDOR	Insecure direct object reference, an access-control flaw the platform defends against
SSRF	Server-side request forgery, an attack class probed in our penetration tests
Web application firewall	An edge control that filters malicious web traffic
Data processing agreement	The contract governing how a processor handles personal data on a controller's behalf

Appendix D: Contact and Document Control

AI Interview Analyzer Sp. z o.o.

ul. Jedrusik 6/53, 01-748 Warszawa, Poland

NIP: 5253079974

For a security review, a copy of our data processing agreement, or our EU AI Act conformity documentation, please contact your account representative.

This document describes the security posture of the AI Interview Analyzer service as of the generation date shown in the footer. It is provided for evaluation purposes and does not form part of any contract. Specific contractual security commitments are set out in the applicable agreement and data processing agreement.