Orvensa Compliance Overview

Legal review memorandum for recruiting organizations and their counsel.
This document is drafted as a controller-side compliance overview to support internal legal review of the current Orvensa product implementation.

Document status	Implementation-based draft, prepared from the current codebase and internal policy documents.
Intended use	Shared by recruiters or hiring teams with legal, privacy, procurement, or security reviewers.
Regulatory frame	Belgian and wider EU data-protection context, including GDPR accountability and AI governance expectations; recruitment AI is treated in the EU AI Act as a high-risk use case in employment and worker management.
Important limitation	This document describes the current product implementation and control framework. It does not replace customer-specific legal advice, local employment-law review, or a customer-specific records-of-processing analysis.

Prepared for: Orvensa customer legal review workflows

1. Executive Summary

Orvensa is an organization-scoped hiring platform that supports recruiter review of job applicants by combining CV ingestion, structured evaluation outputs, explainability features, audit logging, data-subject-rights workflows, retention controls, and organization-level governance settings. The current implementation is designed around a human-in-the-loop model: the platform produces decision-support outputs, while final hiring decisions remain reserved to human reviewers.

- The recruiting organization is positioned as the primary controller for applicant data processed through the service.
- The platform captures and stores organization-level compliance settings, including legal basis, privacy-notice URL and version, retention window, and AI enablement controls.
- Applicant records snapshot the legal basis at collection, the privacy notice shown, the notice version shown, and the time at which notice was shown.
- The platform supports access/export, erasure, rectification, processing restriction, objection to automated processing, contestation of evaluation outcomes, and manual final decision workflows.
- The current implementation contains data-minimization controls before AI submission, but those controls are partial redactions of obvious direct identifiers rather than full anonymization.
- The current implementation can persist extracted CV raw text in evaluation records; legal notices and customer-facing documentation should therefore not state that raw CV text is never stored.

2. Compliance Positioning and Allocation of Roles

Based on the present architecture and product documents, the expected allocation of responsibilities is as follows:

Role	Expected position	Practical effect
Recruiting organization	Controller	Determines whether applicant data is processed for a particular hiring process, selects legal basis, presents the applicable notice, decides retention, and makes the final hiring decision.
Orvensa	Processor / technology provider	Provides the platform, organizational controls, auditability, AI-assisted evaluation functionality, and supporting documentation and

		incident-management tooling.
OpenAI and other subprocessors	Subprocessor / onward processor role depending on contract chain	Used for model inference and other infrastructure functions; must be covered by the customer-facing DPA and subprocessor disclosures.

This role allocation should be reflected consistently across the privacy policy, terms, data processing agreement, subprocessor list, and any customer onboarding materials.

3. System Description and Core Processing Operations

- Recruiters create organization-scoped accounts and work within an active organization context selected through membership and the X-Org-ID scoping model.
- Jobs are stored per organization and may contain job description text, job-post text, recruiter notes, manager intake notes, optional custom prompts, and optional rubric configuration.
- Applicants are uploaded against a specific job. The system stores the uploaded CV file, a per-organization and per-job file hash for deduplication, and optional intake data.
- The system extracts text from uploaded CV documents and may cache extracted CV raw text and display-formatted text inside the linked evaluation record.
- The AI workflow uses job context plus minimized CV text to generate decision-support outputs such as summary, strengths, risks, interview questions, confidence, score, and fit bucket or rubric subscores.
- A human reviewer can view the evaluation, review an explanation payload, contest the result, block automated use, or record the final hiring decision.

4. Data Categories and Data Flow

Category	Examples in current implementation	Main source	Notes
Identity and contact data	Applicant name fields, applicant email, recruiter account details, billing contact details	Recruiter upload, account creation, organization settings	Applicant profile fields may also be inferred from extracted evaluation data if not stored directly.
Application materials	Uploaded CV file, extracted CV text, job-specific intake information	Recruiter upload	The system deletes CV files on applicant deletion and also cleans up replaced CV files.
Job and review context	Job description, job post, recruiter notes, manager intake, context-file extracts, rubric configuration	Recruiter input	Job context is snapshot into evaluation metadata when processing occurs.

Derived assessment outputs	Score, subscores, fit bucket, confidence, summary, strengths, risks, interview questions, contest flags, final decision metadata	AI workflow and recruiter actions	These outputs are expressly framed as decision support, not as autonomous hiring decisions.
Governance and accountability data	Audit logs, DSAR records, retention settings, security incident records, fairness snapshots	System operation and admin workflows	Used for accountability, compliance review, and internal governance.

5. Lawful Basis and Transparency Controls

The present implementation includes explicit controller-side configuration points for legal basis and transparency. At organization level, Orvensa stores the configured legal basis, privacy-notice URL, privacy-notice version, retention enablement, retention period, and AI enablement settings. At applicant level, the system snapshots the legal basis at collection, notice URL shown, notice version shown, notice timestamp, and consent status where the organization uses consent as its legal basis.

- Supported legal-basis values in the current model are legitimate interest, contract, and consent.
- Where the organization has selected consent, applicant upload requires a consent flag before processing is accepted.
- The policy layer blocks AI processing if the privacy notice shown or privacy notice version is missing.
- The policy layer also blocks AI processing if the organization legal basis is missing or, where consent is required, consent has not been recorded.
- The compliance serializer exposes 'blocking reasons' such as missing legal basis, missing privacy-notice URL, missing privacy-notice version, or missing retention configuration.

From a legal-review perspective, these controls are useful because they create evidence of transparency and basis selection at the point of collection. Customer counsel should nevertheless confirm that the privacy notice actually presented to applicants contains the full disclosures required for the specific recruitment process and jurisdictions concerned.

6. AI Use, Human Oversight, and Article 22 Positioning

The platform is currently designed and described as AI-assisted decision support rather than a fully automated decision-making system. This distinction is significant under GDPR and under broader Belgian and EU supervisory expectations. The implementation contains several safeguards intended to reinforce that the final hiring decision is taken by a human reviewer rather than by the model alone.

- Organization-level kill switches allow AI screening and AI evaluations to be disabled.
- Evaluations set human_review_required to true and retain a separate final_decision field completed by a human user.
- Applicants can object to automated processing; once objected, AI processing is blocked and manual review is required.
- Applicants can be marked as processing restricted; once restricted, AI processing is blocked and the retention purge skips the affected record.

- Evaluations can be contested; contested evaluations are excluded from automated processing and reset to human-review posture.
- The explanation endpoint returns the model metadata, automation version, strengths, risks, and the separate human decision status.

Customer-facing legal texts should therefore describe the system as a decision-support tool with meaningful human review. They should not say that the system autonomously accepts or rejects candidates. Recruiters should also be trained not to treat model output as determinative.

7. Data Minimization and Current Data Persistence Reality

The codebase contains a minimization function that redacts obvious email addresses and phone numbers before CV text is sent for AI inference. The implementation expressly states that this is a pragmatic minimization control and not perfect anonymization. That is the correct description for legal purposes and should be retained.

- Email addresses are replaced with a redaction token before AI submission.
- Phone-like strings with sufficient digits are replaced with a redaction token before AI submission.
- Whitespace is normalized before submission.
- The minimization layer does not remove all direct or indirect identifiers and should not be described as anonymization or irreversible de-identification.

Separate from minimization before inference, the current implementation can persist extracted CV raw text in evaluation metadata. This follows from the task and evaluation views, which write raw_text into evaluation.extracted where text extraction succeeds. Legal and security materials should therefore state the narrower and accurate proposition: Orvensa applies minimization before AI inference, but extracted CV text may be retained within evaluation records for product functionality and review workflows unless and until deleted under retention, manual deletion, or DSAR erasure.

8. Data Subject Rights Workflows

Right / control	Current implementation	Observations for legal review
Access	Applicant export and DSAR export endpoints build structured payloads including applicant data, job data, evaluation outputs, transparency information, and relevant audit logs.	The export deliberately excludes embedded CV file contents and includes CV metadata instead. Counsel should decide whether that scope is appropriate for operational handling or whether CV delivery should sometimes be handled separately.
Erasure	DSAR erasure deletes evaluations, deletes the applicant record, and	The erasure function is idempotent and records audit events even

	relies on cascade plus pre-delete file cleanup to remove the CV file from storage.	where the applicant had already been deleted.
Rectification	DSAR rectification supports update of key applicant identity fields and intake-data patches; it stores a 'before' snapshot and stamps the application time.	This is positive from an accountability standpoint because it documents what changed.
Restriction	DSAR restriction and manual restriction mark the applicant as processing_restricted, record timestamp and reason, and block future automated processing.	Retention enforcement skips restricted applicants, which helps avoid deleting a record while restriction is active.
Objection to automation	Applicants can be flagged as objecting to automation; once flagged, automated processing is blocked and manual review is required.	The platform should make this route operationally available where the customer commits to it in notices.
Contestability	Evaluations can be contested; contesting clears final decision and forces human-review posture.	Useful for both fairness and Article 22 risk mitigation.

9. Retention and Deletion Controls

Retention is configurable per organization and constrained in the current governance model to a range of 6 to 24 months. Automatic retention enforcement deletes applicants older than the configured threshold where retention is enabled, but it expressly skips records marked as processing restricted.

- Retention is enabled or disabled at organization level and is linked to a retention_months setting.
- Automatic retention enforcement calculates a cutoff date, deletes the stored CV file before row deletion, and writes an audit event for each automatic deletion.
- Applicant pre-delete signals also remove CV files before row deletion, and pre-save signals clean up replaced CV files to reduce orphaned files.
- Manual delete and hard delete flows are available to admins or owners, and DSAR erasure uses an idempotent execution path.
- Deletion of an organization cascades across related organization data.

The retention range currently hard-coded into the product should be aligned with the text used in the privacy notice, DPA, and any customer-facing retention policy. If Orvensa intends to support longer retention periods for some customers, the model constraints and legal materials should be updated consistently.

10. Access Control, Organizational Isolation, and Accountability

The platform uses authenticated organization scoping and role-based access control. An active organization is resolved from the X-Org-ID header where supplied, with fallback to the user's active or most recent organization membership. API middleware rejects access where no active organization is selected for organization-scoped routes, where the organization does not exist, or where the user is not a member of that organization.

- IsOrgMember, IsOrgAdminOrOwner, and IsOrgOwner permission classes gate access to sensitive routes.
- Sensitive actions such as exports, final decisions, governance changes, manual deletions, and billing actions are reserved to elevated roles.
- Audit logging records key events such as applicant views, evaluation views, explanation views, final decisions, DSAR creation and updates, erasure, rectification, processing restriction, retention deletions, governance changes, and security-incident actions.
- Audit logs capture organization ID, user, event type, relevant record IDs, timestamp, and request IP address where available.

11. Security Controls and Incident Handling

The current codebase includes basic but relevant security and resilience measures, together with a dedicated incident register. From a legal-review perspective, these controls support the accountability narrative, but some deployment-level controls will still need to be confirmed in production before strong statements are made externally.

Control area	Current implementation evidence
Authentication	JWT-based API authentication; password reset, email verification, email-change confirmation, and password-change notification flows are implemented.
File handling	CV file type restrictions and size limits are enforced on upload; CV files are cleaned up on delete and replacement.
Background resilience	Celery background processing includes retries for transient AI-provider failures and stale-processing protection.
Webhook integrity	Stripe webhooks are signature-verified, stored idempotently, and processed exactly once by event ID.
Incident handling	Security incidents can be created, updated, and marked reported at organization level, with audit logs for these actions.
Secrets and service credentials	Settings are environment-driven for Stripe and SendGrid secrets; production hardening must ensure DEBUG is disabled and secret management follows deployment policy.

12. Third-Party Processing and International Transfer Considerations

The internal processor inventory identifies OpenAI for AI inference, SendGrid for transactional email, Stripe for billing, and a cloud-hosting provider for infrastructure. Legal materials should clearly distinguish which processing operations occur within the service itself and which are carried out by subprocessors.

- OpenAI receives job context and minimized CV text for inference.
- SendGrid receives email addresses and transactional message content for account, billing, and operational notices.
- Stripe receives customer and billing metadata for subscription management and payment handling; card data is not stored in Orvensa.
- The cloud-hosting environment necessarily processes stored application data, logs, and backups.

Customer counsel will typically ask where these providers process data and what transfer safeguards apply. Those points should be completed in the DPA, privacy notice, and subprocessor

schedule with current vendor-specific locations and transfer mechanisms rather than left at a general level.

13. EU AI Act and Belgian/EU Governance Context

Recruitment and worker-management AI is treated under the EU AI Act as a high-risk use case in employment-related contexts. That does not by itself answer all GDPR questions, but it does raise the governance bar for documentation, human oversight, logging, risk management, transparency, and post-market monitoring. Belgian supervisory guidance likewise stresses that AI systems processing personal data must remain accountable, transparent, and lawful under the GDPR.

- The current codebase already contains a risk register, fairness policy, DPIA documentation, AI transparency material, model-governance notes, and EU AI Act mapping documents.
- The product stores model name, prompt version, automation status, human-review requirement, and contestability metadata per evaluation.
- Fairness snapshots and incident records create a basis for internal governance, although the customer-specific governance process around those outputs should still be documented contractually and operationally.

For external legal review, Orvensa should present itself as a provider that supports customer compliance with documentation, controls, and evidence, while making clear that the customer remains responsible for lawful deployment in its own recruitment process.

14. Customer Responsibilities That Should Be Made Explicit

- Select and document the appropriate legal basis for the relevant recruitment process.
- Provide a compliant applicant privacy notice and keep the published notice URL and version current.
- Use the platform in a manner consistent with local employment law, non-discrimination law, works-council obligations if applicable, and sector-specific rules.
- Ensure human reviewers are trained to use AI outputs as decision support rather than as determinative decisions.
- Configure retention in line with the customer's own legal and litigation-hold requirements.
- Handle DSARs and objections within statutory timelines and with appropriate identity verification procedures where necessary.
- Review and approve the subprocessor list, DPA, security materials, and any international-transfer safeguards before go-live.

15. Implementation Gaps or Drafting Issues to Correct Before External Release

Issue	Why it matters	Recommended correction
Some internal policy documents state that raw CV text is not persisted.	The code currently stores raw_text in evaluation.extracted in several flows. Overstating data minimization creates legal and trust risk.	Revise all external and internal legal texts to state the actual position: minimization before inference is applied, but extracted CV text may be retained in evaluation records until deleted.
Current Django settings show DEBUG=True and a development-style secret key in source.	These settings are not acceptable as production security statements.	Ensure production configuration is separated and never represented externally as current production posture.
Certain audit action labels used in code are broader than the declared model action choices, and some deletion/audit labels appear inconsistent across files.	Inconsistencies can undermine evidence quality and implementation confidence.	Normalize audit action taxonomy before external assurance claims are finalized.
The legal materials should not imply that all GDPR compliance is fully automated by the platform.	Controllers remain responsible for their own notices, lawful basis assessment, and recruitment practice.	Frame Orvensa as an enabling control environment rather than as a substitute for customer legal compliance.

16. Recommended External-Facing Positioning Statement

A careful and supportable external summary for customers and their counsel would be along the following lines:

Orvensa provides AI-assisted candidate-review tooling designed for human-in-the-loop recruitment workflows. The platform supports organization-level legal-basis and notice configuration, applicant-level transparency snapshots at collection, objection and restriction controls, DSAR handling, configurable retention, audit logging, and explainability features. AI outputs are intended as decision support; final hiring decisions remain with the recruiting organization.

17. Annex - Source Material Reviewed

- Backend code covering organizations, applicants, evaluations, compliance, audit, incidents, billing, jobs, and account flows.
- Internal governance and policy documents including AI transparency, risk register, DPIA, data retention, fairness policy, security controls, and third-party processor documentation.
- Official legal and supervisory sources referenced for orientation, including GDPR, the EU AI Act, Belgian DPA guidance on AI and the GDPR, and EDPB guidance on automated decision-making and profiling.

References for legal orientation

- GDPR - Regulation (EU) 2016/679.
- AI Act - Regulation (EU) 2024/1689.
- Belgian Data Protection Authority, Information brochure on artificial intelligence systems and the GDPR.
- EDPB guidance on automated individual decision-making and profiling.