Service Definition Document

EVA AI Screening and Matching Service

Prepared by:

EVA.ai

CONTACT

1. **BEN KAMINSKY** – CEO, EVA.ai (ben@eva.ai)
2. General Queries - contact@eva.ai

This document is provided for information and assurance purposes only. It describes the operational, security, and service management characteristics of the EVA.ai platform based on current service capabilities.

The contents of this document are confidential and intended solely for evaluation by authorised representatives of the receiving organisation. It must not be disclosed, reproduced, or distributed to third parties without prior written consent from EVA.ai.

This document does not constitute a legally binding agreement. Service commitments, performance levels, and remedies are governed by the applicable contract and service level agreement.

1 – AI Screening and Matching Service overview and purpose

1.1 Service overview

EVA.ai provides a cloud-based candidate screening and matching capability that supports recruiters and hiring stakeholders to move from raw applications to a prioritised shortlist in a controlled, auditable way. The service first structures and evaluates candidate information against role criteria (screening). Those structured outputs then feed a matching workflow that ranks and compares candidates against a job query or job requisition context.

The service is delivered as a secure, browser-based SaaS solution. Buyers do not install or manage infrastructure. The service can be used standalone or integrated with recruitment/HR systems so that jobs, candidates, stages and results can flow through existing processes.

1.2 Purpose of the service

The purpose is to help teams apply role requirements consistently across candidate pools and then efficiently surface the strongest-fit profiles for human review. The service supports both high-volume screening and deeper comparison for specialist roles by making criteria visible and outcomes traceable.

The service is decision-support. It does not replace human judgement. Recruiters and hiring teams remain responsible for reviewing outputs and confirming all screening and shortlisting decisions.

1.3 Intended users

The service is designed for:

- Recruitment and talent acquisition teams
- HR/shared services teams managing recruitment operations
- Authorised hiring stakeholders (for example, hiring managers and panel members) with review/oversight responsibilities

Access is restricted to authenticated users and governed by role-based permissions.

1.4 Key service outcomes

The service supports organisations by:

- Turning CV and application data into structured, criteria-based evaluation outputs
- Producing ranked match lists using job context (prompt-based or requisition-based)

- Providing explainable rationale (scores, highlights, and justifications) to support review
- Maintaining traceability through logged configurations and user actions

The service does not guarantee hiring outcomes and does not determine candidate suitability independently of human review.

1.5 Service boundaries

The service:

- Supports screening, comparison, and shortlisting workflows
- Does not automatically hire candidates or replace recruitment governance
- Does not operate as a “black box”; outputs are designed to be reviewable and auditable (for example, scored criteria and written justification)

2 – Scope of service

2.1 Scope of the service

The service supports a two-stage flow that is used iteratively:

1. **Screening:** evaluate each candidate against configured role criteria and generate structured outputs.
2. **Matching:** use those outputs, plus job context, to compare, rank, and prioritise candidates into a shortlist for human review.

This flow can be used for new applicants entering a workflow stage (screening triggered by candidate movement or pipeline actions) and for recruiter-led matching and refinement of shortlists.

2.2 Defining the job context (job query / requisition input)

Matching can be driven in two practical ways:

- **Auto-generated job query from a job requisition or job description:** when a requisition is created or updated, the service can generate a structured matching prompt/query using the full context of the role (responsibilities, required qualifications, desired skills).

- **User-generated prompt or query:** authorised users can search using natural-language prompts to describe the candidate they want to find (for example, skills, experience depth, constraints), and iteratively refine the query.

2.3 Candidate data ingestion

Candidate data may be ingested through:

- Secure upload of candidate documents (for example, CVs)
- Structured application or screening forms
- API-based integration with recruitment/HR systems

Only data required for screening and matching is processed. Candidate data is logically segregated by buyer environment within the platform.

2.4 Screening step: evaluation outputs

Once candidate data is available and screening criteria are defined, the service evaluates candidate information and produces structured screening outputs. Core screening components include:

- **Match evaluation against criteria** (how well the candidate meets each requirement)
- **Conclusion labels** (overall fit labels derived from criteria outcomes)
- **Justification** (a human-readable explanation of what was met/missed and why)

These outputs create consistent evidence for downstream comparison.

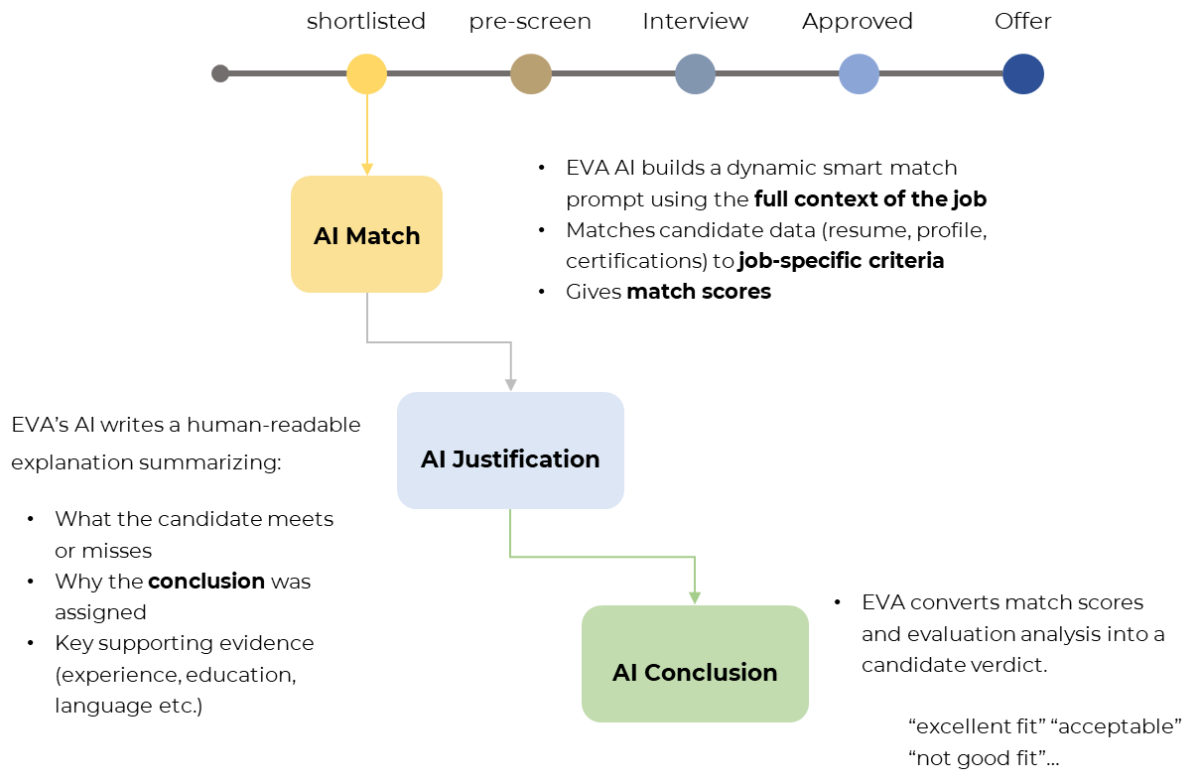

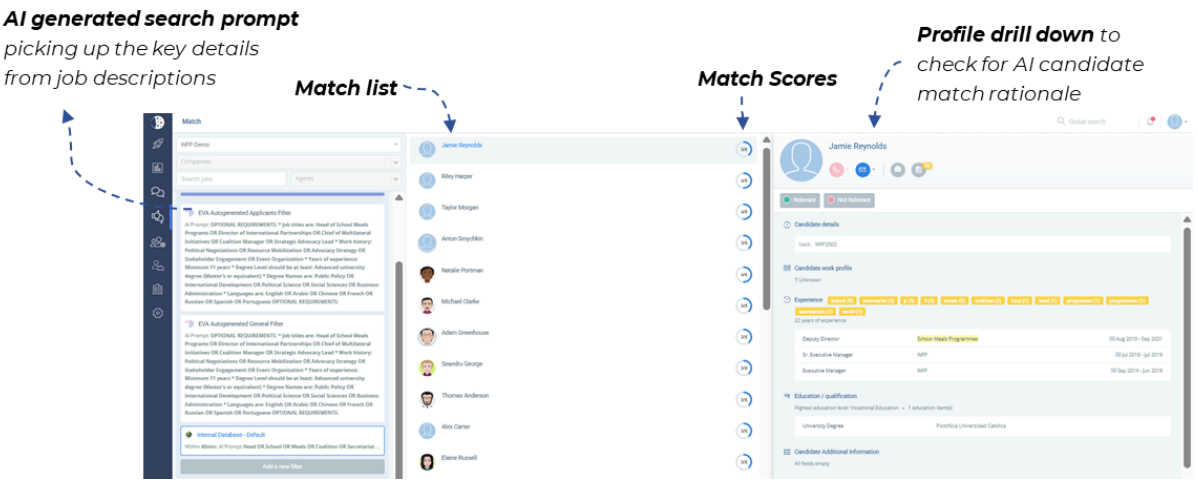

AI Justification and Conclusion Example

Interview & Scoring

Priya Sharma
Programme Policy Officer | World Food Programme
priya.sharma@email.com
Ready to be Interviewed
AI JUSTIFICATION: **Education**: The candidate holds an MSc in Social Policy & ...
Work Experience: The candidate has led capacity-building...
Knowledge & Skills: Through their work on food security ...
Language: Proficiency in English and Spanish fulfills the n...
AI CONCLUSION: Good Fit

Screen & Shortlist

Jonathan Reid
Programme Policy Officer | World Food Programme
jonathan.reid@email.com
Applied
AI JUSTIFICATION: **Education:** The candidate holds a master's degree in Inter...
Work Experience: The candidate has over five years of pr...
Knowledge & Skills: The candidate demonstrates strong ...
Language: The candidate is fluent in English and also poss...
AI CONCLUSION: Good Fit

Daniel Wu
Emergency Preparedness Officer
daniel.wu@email.com
Applied
AI JUSTIFICATION: **Education**: The candidate has obtained a Master's in Envir...
Work Experience: With years of progressively responsible...
Knowledge & Skills: The candidate shows an ability to ad...
Language: The candidate has English and Portuguese, whi...
AI CONCLUSION: Good Fit

Maria Gonzales
Food Security Programme Officer
maria.gonzales@email.com
Applied
AI JUSTIFICATION: **Education**: The candidate holds a Master's degree in Deve...
Work Experience: Their background includes experience ...
Knowledge & Skills: The candidate has engaged in food s...
Language: They possess fluency in English and Spanish. A...
AI CONCLUSION: Good Fit

2.5 Matching step: ranking and shortlist generation

Matching uses job context (auto-generated or user-generated) to generate a ranked match list. Typical matching outputs include:

- Ranked candidate lists (IDs/profiles) with scores
- Highlighted evidence aligned to the job query/criteria
- Explainable scorecards / rationale to support review and auditability

Recruiters can iterate (change requirements, adjust weights, refine the query) and regenerate the shortlist based on the updated job context.

EVA Smart Match Prompt based Talent Search and Match

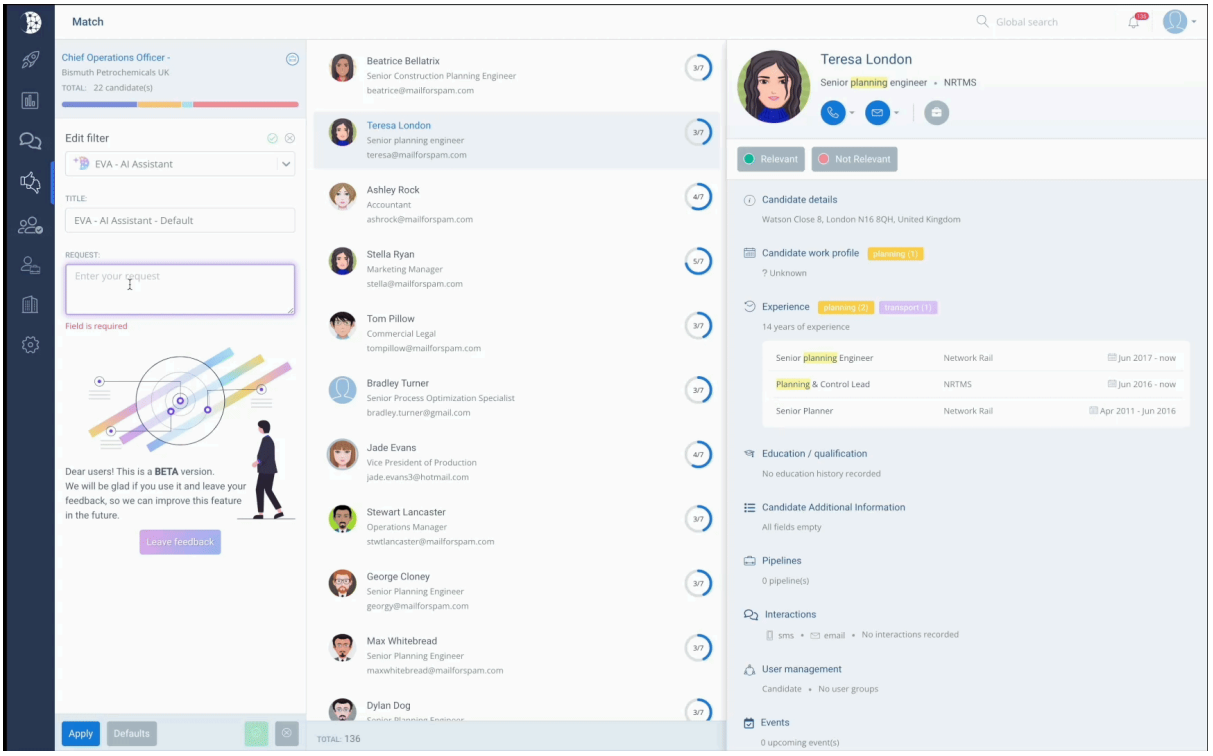

- Smart Match enables recruiters, workforce planners, and managers to query talent data using natural language, delivering fit scores, highlighted insights, and explainable AI justifications.
- It is distinct from background auto-ranking tools, empowering teams to conduct flexible, scenario-driven exploration across workforce management contexts — volume recruiting, internal mobility, succession, contractor sourcing, and workforce planning — while maintaining human-in-the-loop explainability, transparency, and auditability throughout.

EVA.ai

Smart Match Assistant

interpret requests in natural language to surface Candidates, Employees and Contractors for any workforce management use cases

Recruitment Prompt User-generated

EVA - AI Assistant - Default

Keywords: **Hi EVA, I need you to surface all Candidates that updated their profile last year, that have HR Business Partner in their job titles and whose desired salary is less than 100K euros, who have strong interpersonal skills, and who have worked with technologies such as Workday or Successfactors in large public organizations in Europe that speak German or French**

Internal Mobility Prompt Auto-generated

Find candidates who have:
- Grades are P2 or P3
- Quantum Contract Type are FTA OR TA
- Contract Expiry date is less than 2 years from now and greater than now
- Minimum 5 years of experience in security or development
- Skills are Security or safety
- Bachelor Degree or higher
- Language is English

Optional:
- Language is Arabic
- Position is Field Security or related
- Employer is UN
- Full text search: ERP or "crisis environment" or "UN Security" or "risk management"
- Degree name is Social Sciences OR Management OR Security or related

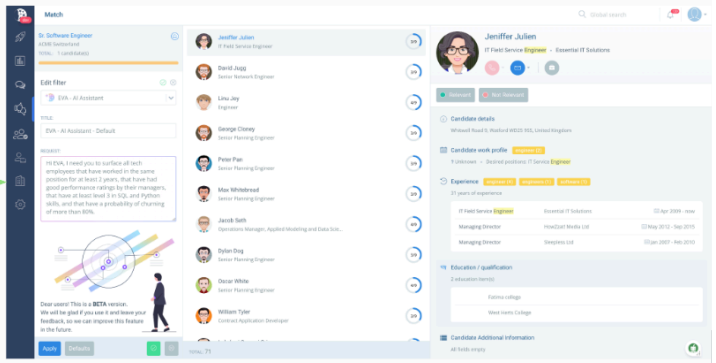

Succession / Promotion Prompt calling 3rd Party data and prediction model

EVA - AI Assistant - Default

Keywords: **Hi EVA, I need you to surface all tech employees that have worked in the same position for at least 2 years, that have had good performance ratings by their managers, that have at least level 3 in SQL and Python skills, and that have a probability of churning of more than 80%.**

Pre-built Prompts

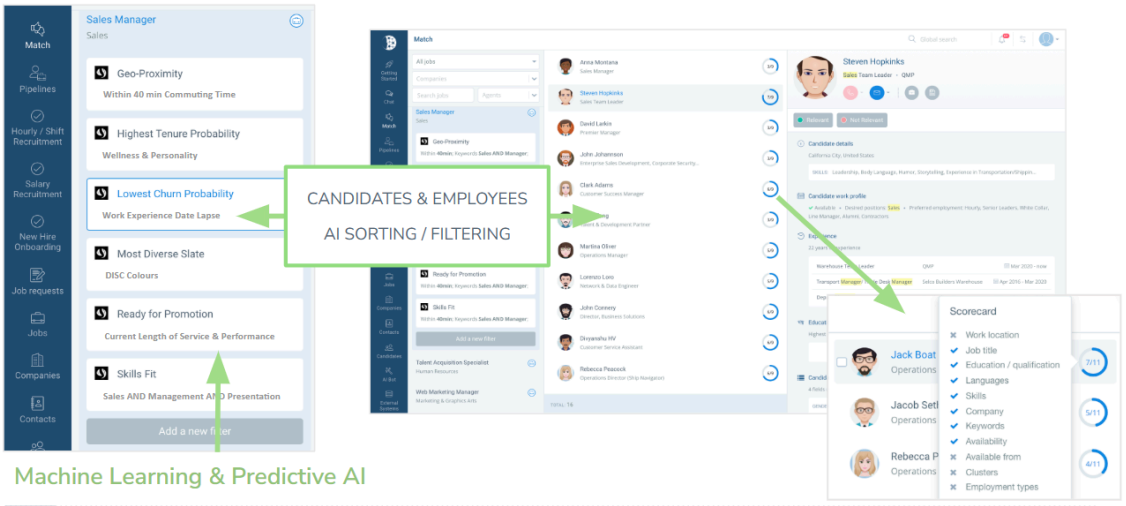

Machine Learning & Predictive AI

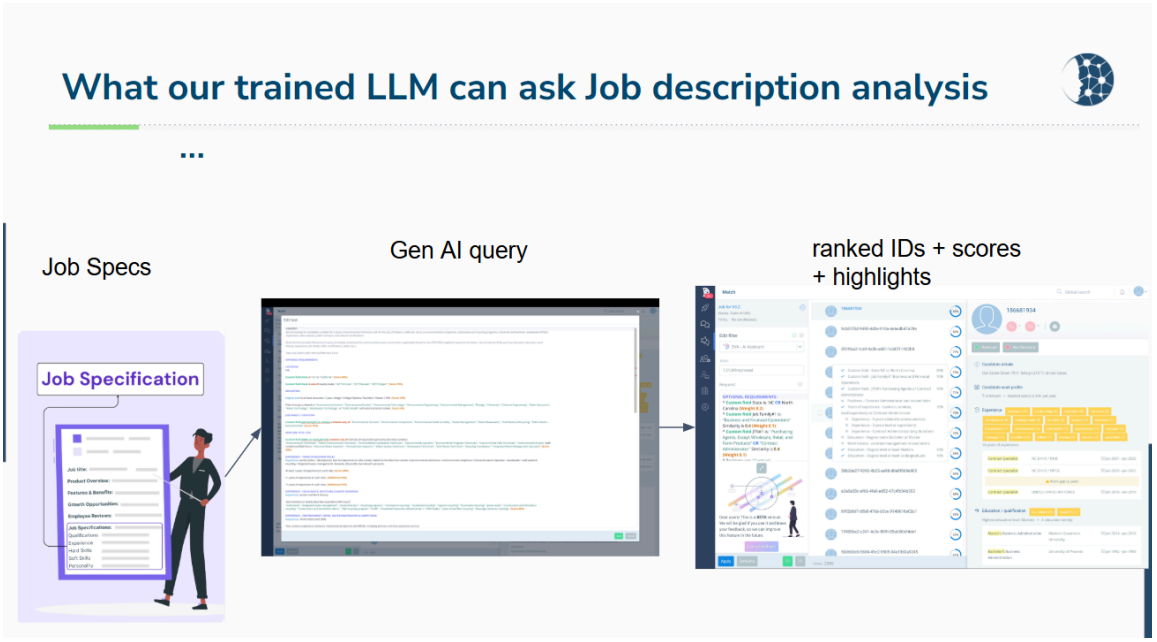

2.6 Human review and confirmation

Authorised users review screening and matching outputs, drill into evidence, and apply professional judgement before confirming any shortlist outcomes. No candidate is automatically progressed or rejected without human involvement.

2.7 Audit and traceability

The service logs relevant activity for governance and audit purposes, including:

- Screening/matching criteria configurations and changes
- Candidate evaluation events and shortlist generation
- User actions, confirmations, and timestamps

3 – User roles and access model

3.1 User roles

AI Screening and Matching supports multiple user roles to ensure appropriate access control and segregation of responsibilities. User roles are assigned based on organisational needs and are configurable within the service.

Typical roles include:

Administrators
Responsible for managing user accounts, assigning roles, configuring service-wide settings, and overseeing access permissions.

Recruiters / screening and matching users
Responsible for configuring screening and matching criteria, defining job context or queries, reviewing screening and matching outputs, and confirming shortlisting outcomes.

Read-only or oversight users
Responsible for reviewing screening and matching activity, outcomes, and audit information without the ability to modify configurations or decisions.

The exact role structure may vary depending on buyer requirements. All access is governed by role-based access controls.

3.2 Authentication and access control

Access to the EVA platform requires user authentication. Users are authenticated using secure credentials and, where configured by the buyer, multi-factor authentication.

Access to platform functionality and data is controlled through role-based access controls (RBAC). Users can only access features and data explicitly permitted by their assigned role. This ensures that administrative, configuration, screening, and matching activities are restricted to authorised users.

Privileged access to management functions is limited to named roles and reviewed periodically.

3.3 Management and support access

Access to management interfaces is restricted to authorised EVA.ai personnel with defined operational responsibilities. Administrative access is protected using secure authentication mechanisms and is monitored for security and audit purposes.

Support access to buyer environments is provided only where required to resolve support requests. Such access is time-bound, logged, and restricted to the minimum level necessary to perform support activities.

3.4 Audit and monitoring of access

User access and actions within the service are logged. Audit records include:

- User authentication events and access attempts
- Configuration changes
- Screening and matching activity and shortlist confirmations

These logs support governance, security monitoring, and compliance requirements. Buyers can access audit information relating to their users through the service interface. Audit information relating to supplier actions is available on request through the support process.

3.5 Principle of least privilege

Users are granted the minimum access required to perform their responsibilities. Access rights can be updated or revoked by authorised administrators as roles or responsibilities change.

4 – Data handling and processing

4.1 Types of data processed

The service processes candidate and recruitment-related data submitted by the buyer for screening and matching purposes. Data types may include:

- Candidate CVs and resumes
- Application form responses
- Screening question responses
- Skills, experience, and qualification information
- Recruitment workflow metadata (for example, role identifiers or status indicators)
- Job or role-related information (for example, job descriptions or requisition details) used to define matching context

Only data required to perform configured screening and matching activities is processed. The buyer determines what data is submitted to the service and remains responsible for the completeness and accuracy of that data.

4.2 Data roles and responsibilities

EVA.ai acts as a data processor on behalf of the buyer. The buyer acts as the data controller and determines the lawful basis and purpose for processing candidate data.

Data processing is performed strictly in accordance with buyer instructions and applicable data protection legislation, including the UK GDPR and EU GDPR where relevant.

4.3 Data processing activities

Candidate data is processed to:

- Evaluate candidate information against buyer-configured screening and matching criteria
- Generate structured screening outputs, matching scores, and ranked candidate lists
- Present outputs for authorised user review and confirmation
- Record audit information relating to screening, matching, and shortlisting activity

Customer data is not used for product training, profiling outside the service scope, or secondary purposes. Data is not shared with third parties except where required to operate the service or to meet legal obligations.

4.4 Data segregation and access control

Customer data is logically segregated within the platform to prevent access between buyers.

Access to data is enforced through authenticated user accounts and role-based access controls. Users can only access data associated with their organisation and assigned permissions.

Access by EVA.ai personnel is restricted to operational or support purposes, is time-bound where applicable, and is logged for audit purposes.

4.5 Data retention

Candidate data is retained in accordance with buyer instructions and contractual terms.

Retention settings may be configured within the service or managed through buyer-initiated deletion requests, depending on service configuration. Audit logs and operational records are retained to support security monitoring, governance, and compliance requirements.

4.6 Data export and portability

Authorised users can export candidate data and screening and matching outputs during the contract term and within an agreed exit period.

Exports are provided in commonly used, machine-readable formats. Once exported, data is under the buyer’s control.

4.7 Data deletion

At contract termination, or upon buyer instruction, customer data is securely deleted following the agreed retention period.

Deletion processes are designed to prevent data recovery and are aligned with EVA.ai's security and data protection controls. Deletion activities are logged to support audit and assurance requirements.

5 – Backup, restore, disaster recovery, and business continuity

5.1 Backup approach

EVA.ai is hosted on cloud infrastructure designed with resilience and redundancy. Customer data is backed up regularly to support service continuity and data protection requirements.

Backups are stored securely within the European Economic Area (EEA) and protected using encryption and access controls. Backup processes are managed and monitored to ensure data integrity.

Backup data is used solely for recovery purposes and is not accessed for operational use.

5.2 Data restoration

In the event of data loss, corruption, or service disruption, restoration procedures are in place to recover data from backups.

Data restoration activities are initiated based on incident severity and operational requirements. Restoration efforts aim to minimise disruption to service availability while maintaining data integrity.

Buyers may request support for data restoration where appropriate, subject to the terms of the support agreement.

5.3 Business continuity and Disaster Recovery

Business continuity management approach

EVA.ai maintains a documented Business Continuity Plan (BCP) designed to support the continued delivery of business-critical services during and following disruptive incidents. The plan provides a structured framework for responding to emergencies or incidents that may disrupt normal operations, including technical failures, loss of staff availability, loss of facilities, or external events.

The objective of the BCP is to minimise disruption to service delivery, protect customer data, and ensure the timely restoration of business-critical functions. The plan is reviewed periodically and updated where significant organisational or operational changes occur .

Identification of business-critical services

Business-critical services are defined as operational and support functions that are required to maintain continuity of recruitment screening, matching, and shortlisting activities. For this service, business-critical services include:

- Availability of the core application supporting candidate screening, matching, and shortlist generation
- Access to candidate data, job or role data, and configured screening and matching criteria
- Secure, authenticated access for authorised users performing screening, matching, and review activities
- Availability of audit logging and traceability functions supporting governance and compliance
- Operation of support, monitoring, and incident response functions required to restore service availability or functionality

These services are prioritised for recovery during any disruptive incident.

Incident response and continuity phases

The business continuity approach considers multiple recovery phases following an incident:

- **Immediate response (first 24 hours):**
  Focus on incident assessment, staff safety, communication, and stabilisation of essential services.
- **Short-term continuity (2–7 days):**
  Restoration of core service functionality, support operations, and access to essential systems.
- **Extended recovery (8–14 days):**
  Full restoration of services, validation of system integrity, and transition back to normal operating conditions.

Recovery activities beyond this period are managed as part of ongoing service restoration where required .

Staff continuity and remote working

The business continuity plan assumes the ability for staff to work remotely in the event of loss of access to physical premises. EVA.ai operates a distributed workforce model, which reduces reliance on any single location.

In scenarios such as pandemics or severe local disruption, staffing levels may be reduced temporarily. The plan accounts for prioritisation of business-critical roles and the reallocation of resources to maintain essential services .

Backup and data protection

Customer data is backed up regularly using secure, automated processes. Backups are stored within the European Economic Area (EEA) and protected using encryption and access controls.

Backup data is used exclusively for recovery purposes and is not accessed for routine operations.

Disaster recovery and system restoration

Disaster recovery arrangements are designed to support the restoration of service following major incidents such as loss of infrastructure or data corruption.

Recovery measures include:

- Redundant infrastructure components
- Geographic separation of primary and backup environments
- Defined restoration procedures for essential IT systems and records

Disaster recovery procedures are tested periodically to validate effectiveness. Detailed recovery objectives and restoration timelines are available to buyers on request.

Communications during disruption

Clear communication procedures are defined within the business continuity plan. These include:

- Internal staff communication and coordination
- Notification to buyers where service delivery is affected
- Escalation to senior management where required

All significant actions and decisions taken during an incident are logged to support audit and post-incident review .

Review and continuous improvement

Following significant incidents or tests, the business continuity plan is reviewed to identify lessons learned and improvement actions. This supports continuous improvement of resilience, preparedness, and service reliability.

6 – Security and compliance controls

6.1 Information security framework

EVA.ai platform services operates within a formal information security management framework aligned with recognised industry standards. EVA.ai maintains policies and procedures covering information security, risk management, incident response, access control, and data protection.

The organisation's security governance is supported by independent assurance, including **ISO/IEC 27001** and **SOC 2 Type II** compliance.

6.2 Data security

Data security controls are applied throughout the service lifecycle.

Data in transit

- All data transmitted between users and the service is protected using **TLS 1.2 or above**.
- Secure communication protocols are enforced for integrations and APIs.

Data at rest

- Customer data is encrypted at rest using industry-standard encryption.
- Access to stored data is restricted through role-based access controls.

6.3 Access control and authentication

Access to the service is restricted to authenticated users. Authentication mechanisms include secure credentials and, where configured, multi-factor authentication.

Role-based access controls are used to ensure users can only access data and functionality appropriate to their role. Privileged access is restricted to authorised personnel and reviewed periodically.

6.4 Monitoring and logging

Security monitoring is in place to detect and respond to potential security events. Logs are maintained for:

- User access and activity
- Configuration changes

- Administrative actions

Logs are protected against unauthorised access and retained in line with governance and compliance requirements.

6.5 Vulnerability management and penetration testing

The service follows defined vulnerability management processes. These include:

- Regular vulnerability scanning
- Assessment of identified risks
- Timely remediation based on severity

Independent penetration testing is conducted at least annually to identify and address potential security weaknesses.

6.6 Incident management

Security incidents are managed through defined incident response procedures. Incidents are assessed, prioritised, and escalated based on severity.

Where incidents affect customer data or service availability, buyers are notified in accordance with contractual obligations and applicable regulatory requirements.

6.7 Compliance with data protection requirements

The service supports compliance with data protection regulations, including GDPR. Measures are in place to support data subject rights, such as access, rectification, and deletion, where applicable.

EVA.ai acts as a data processor and processes data only in accordance with buyer instructions and contractual terms.

7 – Onboarding and implementation support

7.1 Onboarding approach

EVA.ai provides onboarding support to help buyers begin using the platform effectively. Onboarding activities are designed to support a smooth transition into service use while minimising disruption to existing recruitment processes.

Onboarding is typically delivered remotely and may be tailored based on buyer requirements and service scope.

7.2 Initial service setup

Initial setup activities may include:

- Creation of the buyer's secure service environment
- Configuration of user accounts and roles
- Assistance with initial screening and matching configurations
- Validation of access and permissions

Setup activities are coordinated with authorised buyer contacts to ensure alignment with organisational requirements.

7.3 Training and guidance

EVA.ai provides guidance to support users in understanding and operating the service. This may include:

- User documentation
- Configuration guidance
- Remote walkthrough sessions

Training focuses on enabling users to configure screening and matching criteria, define job context or queries, interpret screening and matching outputs, and manage core service features independently.

7.4 Integration support

Where integrations with external systems are required, EVA.ai provides guidance on supported integration approaches and available APIs.

Integration activities are scoped and agreed in advance. Buyers remain responsible for the configuration and operation of third-party systems.

7.5 Ongoing support during onboarding

During the onboarding phase, buyers have access to support channels to raise questions or issues. Support requests are logged and managed in line with the defined support model.

Onboarding is considered complete once users can independently access and operate the service.

8 – Configuration, customisation, and service constraints

8.1 Configuration within standard service functionality

EVA platform is configurable within its standard functionality to align with buyer requirements. Configuration is performed through the secure web-based service interface by authorised users.

Configurable elements may include:

- Screening criteria such as skills, experience, and qualifications
- Weighting and prioritisation of screening and matching criteria
- Role- or campaign-specific screening and matching configurations
- Job context or matching queries used to compare and rank candidates
- User roles and access permissions

Configuration changes apply prospectively and do not alter historical screening or matching records, supporting auditability and traceability.

8.2 Customisation boundaries

The service supports configuration rather than unrestricted customisation. Buyers can tailor how the service is used through available settings but cannot modify the underlying service architecture or core functionality.

Bespoke development, non-standard workflows, or changes outside standard service capabilities are not included as part of the core service. Where such changes are required, they may be considered separately by agreement and subject to additional cost.

8.3 Service constraints

The service operates within the following constraints:

- Access requires internet connectivity and a modern web browser
- Planned maintenance is scheduled outside UK business hours where possible
- Maintenance windows are communicated in advance
- The service is delivered as a shared, multi-tenant SaaS solution
- No specialist hardware or client-side software is required

8.4 Buyer responsibilities

Buyers are responsible for:

- Managing user accounts, roles, and access permissions within their organisation
- Defining and maintaining screening and matching criteria in accordance with organisational policies
- Managing internal recruitment processes, review activities, and final decision-making

9 – Service levels, availability, performance, and support

9.1 Service availability

The service is operated on a 24×7 basis, excluding scheduled maintenance windows.

- **Availability commitment:** 99% uptime, measured monthly
- **Measurement basis:** Application-level availability, excluding pre-notified maintenance
- **Monitoring:** Continuous automated monitoring of service endpoints and core platform components

Tolerance:
Monthly uptime below 99.5% but above 99.0%.

Remedy / service behaviour:
Where availability falls below the committed threshold, service credits are applied in line with the agreed service level agreement. Availability incidents are reviewed to identify root causes and preventive actions.

9.2 Planned maintenance

Planned maintenance is carried out to support security patching, platform updates, and infrastructure improvements.

- Maintenance is scheduled **outside UK business hours where possible**
- Buyers are notified **in advance** of planned maintenance windows
- Maintenance is designed to avoid data loss and minimise user disruption

Emergency maintenance may be performed where required to address urgent security or stability risks. Emergency actions are logged and reviewed post-implementation.

9.3 Service performance

The platform is designed to support sustained recruitment workloads, including high-volume candidate processing for screening and matching.

Operational characteristics include:

- Browser-based interaction with no local processing dependencies
- Horizontally scalable processing of candidate data to support screening and matching activities
- Asynchronous handling of screening and matching workloads to maintain interface responsiveness

Performance indicators are monitored continuously. Capacity is adjusted based on observed usage patterns rather than fixed thresholds, enabling the service to absorb demand spikes without manual intervention.

9.4 Support model

Support is provided via email and an online ticketing system. All support requests are logged and tracked.

Requests are categorised by severity, including:

- Service unavailability
- Material degradation of service functionality
- Functional defects
- General usage or configuration queries

Response and resolution targets are defined contractually and enforced through internal support procedures.

9.5 Support hours

- **Standard support:** UK business hours, excluding public holidays
- **Extended support:** Available by prior agreement under enhanced service tiers

Support coverage and escalation paths are defined as part of the agreed SLA.

9.6 Escalation and communication

Support tickets are escalated based on severity and elapsed time against SLA targets.

Buyers are informed of:

- Incident acknowledgement
- Progress updates for high-impact issues
- Confirmation of resolution

For significant incidents, a post-incident summary may be provided upon request.

9.7 Service review and improvement

Service availability, incident trends, and support performance are reviewed internally on a recurring basis. Outcomes from reviews are used to inform platform improvements and operational adjustments within the standard service scope.

10 – Incident, outage, and maintenance management

10.1 Service monitoring

The platform is subject to continuous monitoring across multiple layers, including:

- Application availability
- Infrastructure health
- Error rates and abnormal system behaviour
- Security-relevant events

Automated alerts are generated when predefined thresholds are breached, triggering investigation by operational staff.

10.2 Incident identification and classification

Incidents may be identified through:

- Automated monitoring alerts
- Buyer-reported support requests
- Internal operational checks

Incidents are classified based on **impact and scope**, which determines response priority and escalation level.

10.3 Incident response and resolution

Incident management follows a defined operational process:

1. Incident triage and impact assessment
2. Containment and mitigation where required
3. Root cause investigation
4. Corrective action and validation

All incidents are tracked through the ticketing system to ensure traceability and accountability.

10.4 Buyer communication

Where incidents affect service availability or buyer operations, buyers are notified through established communication channels.

Communications typically include:

- Confirmation of incident identification
- Progress updates for ongoing incidents
- Resolution confirmation

Post-incident summaries may be provided for high-impact events.

10.5 Outage management

Service outages are treated as high-severity incidents.

Outage handling prioritises:

- Restoration of service availability
- Protection of data integrity
- Verification of platform stability before closure

Outages are logged and reviewed to identify corrective and preventive measures.

10.6 Maintenance management

Maintenance activities include:

- Security patching

- Platform updates
- Infrastructure changes

Planned maintenance is communicated in advance where possible. Emergency maintenance may be performed to address critical risks. All maintenance actions are logged for audit and operational review.

10.7 Operational learning

Lessons learned from incidents, outages, and maintenance activities are reviewed and used to refine operational processes and controls.

11 – Hosting architecture and data residency

11.1 Hosting model

The service is delivered as a **multi-tenant SaaS platform** hosted on public cloud infrastructure and managed entirely by EVA.ai.

Buyers do not provision, manage, or maintain infrastructure components. Platform operations, patching, monitoring, and scaling are managed centrally.

11.2 Cloud infrastructure

The platform is hosted on **Amazon Web Services (AWS)** using standardised, security-hardened configurations.

Infrastructure is designed to support:

- Horizontal scalability
- High availability
- Operational resilience

Customer environments and data are logically segregated within the shared infrastructure.

11.3 Data storage and processing locations

Customer data is stored and processed within the **European Economic Area (EEA)**.

- **Primary hosting:** Ireland
- **Backup and disaster recovery:** Additional EEA locations

Data is not transferred outside the EEA except where contractually required or legally mandated. Buyers cannot select individual data centre locations.

11.4 Data residency enforcement

Data residency is enforced through:

- Region-restricted infrastructure deployment
- Controlled access to data and systems
- Monitoring of data access and movement

These controls ensure customer data remains within agreed geographic boundaries.

11.5 Resilience of the hosting environment

The hosting environment incorporates:

- Redundant infrastructure components
- High-availability deployment patterns
- Continuous monitoring and alerting

These measures support availability commitments and operational resilience.

11.6 Buyer responsibilities

Buyers remain responsible for:

- Lawful collection and submission of candidate data
- Compliance with applicable data protection obligations
- Appropriate internal governance over service use

12 – Technical and system requirements

12.1 User access requirements

The service is accessed through a secure, browser-based web interface.

To access the service, users require:

- A modern web browser supporting current web standards
- Internet connectivity capable of supporting encrypted HTTPS traffic
- Authenticated user credentials issued by the buyer's authorised administrator

No client-side software, plugins, or browser extensions are required.

12.2 Supported environments

The service is designed for use on standard desktop and laptop operating systems capable of running modern web browsers, as commonly deployed within public sector environments.

Mobile devices may be used for limited access (for example, viewing information), but the primary service functionality and administrative features are optimised for desktop use.

12.3 Security-related system requirements

End-user devices must comply with the buyer's internal security policies.

This typically includes:

- Up-to-date operating systems and browsers
- Endpoint security controls such as antivirus or managed device protection
- Secure handling of authentication credentials

All communication between users and the service is encrypted using industry-standard protocols (TLS 1.2 or higher). Buyers are not required to manage encryption keys or certificates.

12.4 Integration and API requirements

Where integrations are used, external systems must support secure data exchange using documented APIs or supported data transfer mechanisms.

Authentication credentials, API keys, and access tokens used for integrations must be managed securely by the buyer. EVA.ai provides technical documentation to support integration configuration and operation.

12.5 Buyer responsibilities for system compatibility

Buyers are responsible for:

- Submitting accurate and appropriately formatted data
- Ensuring network access to required service endpoints
- Managing user accounts, roles, and access permissions

EVA.ai is not responsible for service issues arising from unsupported browsers, misconfigured buyer networks, or failures in third-party systems.

13 – After-sales support and service management

13.1 Support channels

After-sales support is provided through the following channels:

- Email
- An authenticated online ticketing system (JIRA helpdesk)

All support requests are logged in the ticketing system to ensure traceability, prioritisation, and auditability. Live web chat support is not provided.

13.2 Support request handling

Support requests are categorised by severity and operational impact, including:

- Service unavailability
- Material degradation of service functionality
- Functional defects
- Usage or configuration queries

Requests are assigned to operational or technical teams in line with defined support procedures. Response and resolution expectations are governed by the agreed service level agreement.

Typically, requests are assigned based on severity and impact. Critical service issues receive an initial response within **2 hours** during business hours and are targeted for resolution within **4 hours** of reporting. Non-critical requests are responded to within agreed SLA response windows. Where resolution exceeds defined tolerances, escalation and service remedies are applied in accordance with the service level agreement.

13.3 Escalation and incident linkage

Where a support request indicates a broader service issue, it is escalated into the incident management process.

Escalation ensures:

- Timely involvement of appropriate technical or operational personnel
- Alignment with incident severity classification
- Tracking through resolution and closure

Buyers are kept informed of progress for escalated issues until resolution.

13.4 Service communications

EVA.ai communicates with buyers regarding:

- Service incidents or outages
- Planned maintenance activities
- Operational updates that may affect service availability or use

Communications are delivered through agreed channels and are intended to support operational awareness and planning.

13.5 Service management and review

Service usage, incident trends, and support activity are reviewed internally on a recurring basis. These reviews inform operational improvements and prioritisation of platform enhancements within the standard service scope.

14 – Offboarding, data export, and exit management

14.1 Contract end and exit planning

At contract termination, buyers retain full ownership of their data. EVA.ai supports an orderly exit in accordance with contractual terms and agreed timelines.

Exit processes are designed to minimise operational disruption and support continuity for the buyer.

14.2 Data export

Authorised users can export their data during the contract term and within an agreed exit period.

Exportable data may include:

- Candidate and recruitment data
- Service configurations relevant to the buyer
- Audit records associated with buyer activity

Data is provided in commonly used, machine-readable formats to support portability.

14.3 Exit support

Standard exit support includes:

- Guidance on available data export mechanisms
- Clarification of data formats
- Confirmation once data extraction is complete

Additional exit assistance beyond standard support may be agreed separately and may incur additional cost, depending on scope and effort.

14.4 Data deletion and sanitisation

Following confirmation that data has been successfully exported, and after the agreed retention period, customer data is securely deleted.

Deletion processes are designed to prevent data recovery and are aligned with EVA.ai's security and data protection controls. Deletion actions are logged to support audit and assurance requirements.

15 – Accessibility considerations

The service is delivered as a browser-based web application and accessed using standard modern web browsers. It does not require proprietary software, plugins, or live audio or chat interaction.

The user interface is designed in alignment with **WCAG 2.2 Level AA accessibility principles**, including use of standard web components, semantic markup, and keyboard-accessible interaction patterns. The service has not undergone formal accessibility certification or independent compliance testing.

The service relies on browser and operating system accessibility capabilities. No additional accessibility configuration is currently provided at the application level. Accessibility considerations are reviewed as part of ongoing service development.