TALENT MARKETPLACE

# Legal and Security Review

A GDPR-first, employment-risk, AI-governance, and light-SOC-2 review of what must be true before Talent Marketplace is offered as an independent B2B beta.

| | |
|---|---|
| **REVIEW DATE** | **August 13, 2026** |
| **CODE SNAPSHOT** | **d1dabf08f38b** |
| **STAGE** | **Beta readiness** |
| **DECISION** | **No-go for real candidate data** |

**CURRENTLY PERMITTED**

Internal or closely supervised demonstration using synthetic candidates, fictional employers, and no employment action.

EXECUTIVE DECISION

# Do not release an external beta that exposes real candidates or influences hiring decisions

Release posture: The prototype displays names, locations, portfolios, evidence, match percentages, capability levels, recommendations, shortlists, and 'not a fit' decisions. The model, candidate source and permission framework, backend authorization, and production controls are outside this repository. Recruiter bearer tokens and hiring decisions are stored in browser localStorage. A real-data beta remains no-go until all twelve launch gates are evidenced.

## Permitted now

- Internal or closely supervised design demonstration
- Synthetic candidate and employer data only
- No outreach, rejection, selection, or other employment action
- No claim that scoring is validated, unbiased, or production ready

## Excluded until gates close

- Real profiles, portfolios, assessments, or rankings
- Automated filtering, rejection, ordering, or decisive recommendations
- Protected, biometric, health, background, or student data
- Regulated geography without current jurisdiction review

## Why the bar is higher than ordinary SaaS

- A suitability score is personal data and likely profiling under GDPR.
- Hiring tools can create disparate-impact and disability-discrimination risk even when a human clicks the final button.
- NYC can require an independent bias audit, public summary, and advance notice.
- If the matching system falls within the EU AI Act, recruitment and candidate selection are an Annex III high-risk use.
- A contract can allocate customer duties, but it cannot erase the product's own privacy, security, accuracy, or discrimination responsibilities.

SOC 2 is not a launch license. The beta needs functioning controls and evidence habits. It does not need an immediate SOC 2 examination, and the product must not be described as 'SOC 2 compliant' or 'certified.'

01 / SCOPE AND EVIDENCE

# What was reviewed and what remains unknown

Findings use the Talent Marketplace frontend at commit d1dabf08f38babc54b78132c2adbb22772060ffe. The platform backend, matching implementation, production runtime, contracts, and operating evidence were not available in this repository.

| Status | Observed product evidence |
|---|---|
| Strength | Invite-only signup, password rules, authenticated routes, shared auth/HTTP packages, typed API contracts, Cloudflare deployment. |
| Strength | Point-in-time production dependency audit reported zero known vulnerabilities on August 13, 2026. |
| Blocker | Bearer access token and expiration are configured in JavaScript-readable localStorage, including a two-week remember-me path. |
| Blocker | Saved and dismissed candidate IDs persist in localStorage without user or tenant namespacing. |
| Blocker | Company, recruiter title, industry, state, target roles, preferred skills, and onboarding state persist in localStorage. |
| Blocker | Invite tokens appear in /signup/:token and a client query-cache key before verification and resend. |
| Blocker | Candidate API data includes name, location, score, recommendation, risk items, skills, evidence, portfolio and capability levels. |
| Gap | No repository-level CSP/security header policy is visible; Vue devtools is in the plugin list for all builds. |
| Gap | LaunchDarkly user context and Cloudflare observability exist; PR E2E remains disabled and CI does not visibly gate on unit/lint/E2E. |

Evidence boundary: Frontend route guards are not proof of server authorization. API tenant enforcement, data sourcing, matching behavior, retention, encryption, logging, rate limits, deletion, and incident practice remain unverified until the backend and deployed service are reviewed.

02 / PRODUCT FLOW

# The marketplace consumes validated evidence and creates a consequential hiring layer

## Intended product boundary

- Skills Validation owns canonical skills, verification status, proficiency evidence, dates, recency, gaps, and reassessment.
- Talent Marketplace owns candidate discovery, search, matching and ranking, profile presentation, shortlists, dismissals, employer workflow, and hiring transaction support.
- A source-evidence error is corrected in Skills Validation and propagated. A display, search, ranking, cache, or employer-workflow error remains Marketplace's responsibility.
- Course completion by itself is not credible professional capability evidence. Matching should use validated, current evidence and show its date and limits.

| Stage | Information and decision |
|---|---|
| Employer input | Company, industry, state, target roles, preferred skills, searches and team context |
| Candidate evidence | Identity, role, location, portfolio, projects, certifications, assessments, skills and dates |
| Marketplace inference | Match score, recommendation, capability level, rank and risk items |
| Employer action | View, shortlist, not a fit, notes, export, contact and eventual selection |
| Feedback loop | Correction, appeal, outcome monitoring, disparity analysis and model change |

## Current implementation concern

The frontend takes the first three available candidates after role filtering and calls them the 'Top 3 recommended.' It also renders match percentages and strong, moderate, or limited capability labels. Whatever the underlying logic, this presentation can materially influence recruiters and must be governed as a decision-support system.

> Unknown model means no-go: A real-data release requires the actual model or rule implementation, data lineage, validation, limitations, change history, and employer instructions to be reviewed. A polished interface is not evidence that the score is job related or fair.

03 / GDPR

# Candidate scores, rankings, and employer actions are personal data

The system combines observed and inferred professional information with employer preferences to evaluate suitability. That processing can affect access to work and should be treated as high risk.

| Data group | Examples | Required treatment |
|---|---|---|
| Recruiter/account | Name, email, credentials/session, company, title, invites, login and security records | Account notice, security, retention, access and deletion |
| Hiring context | Roles, preferred skills, searches, shortlists, dismissals, notes and team activity | Purpose, tenant scope, audit, retention and customer instructions |
| Candidate profile | Name, role, location, portfolio, projects, assessments, certifications, skills and dates | Source, lawful basis, notice, visibility, accuracy, expiry and correction |
| Inference/decision | Score, recommendation, capability level, risk, rank, shortlist and not a fit | Profiling safeguards, explanation, objection, correction, appeal and human review |
| Telemetry/security | Feature context, IP/device, events, errors, logs, source maps and request metadata | Minimize, pseudonymize, redact, retain briefly and govern vendors/transfers |

## Role and lawful-basis model

Legal role is assigned per purpose. Talent Marketplace is likely controller for account security, product analytics, discoverability design, and any matching system whose essential means it determines. An employer is controller for its hiring decision. Marketplace may be processor for employer-only shortlists, notes, and workflow processed solely under employer instructions. The relationship with Skills Validation must be documented rather than inferred from architecture.

> No blanket lawful basis: Contract may cover recruiter accounts. Legitimate interests may support limited adult professional discovery only after a balancing test. Consent can support voluntary discoverability but is not a reliable blanket basis in an employment setting. Special-category data is excluded from beta.

03 / GDPR CONTINUED

# Minimum privacy operating model for real candidates

- Complete the system/data map, Article 30 record, role matrix, lawful-basis analysis, and legitimate-interest assessment.
- Deliver Article 13 notice for directly collected data and Article 14 notice when data comes from Skills Validation, an employer, or another source.
- Complete a DPIA before real candidate profiling. Employment impact, systematic scoring, data combination, and scale make high risk a prudent assumption.
- Support access, correction, deletion, restriction, portability where applicable, objection, and automated-decision rights.
- Avoid solely automated legal or similarly significant decisions. Human review must have information, authority, and time to change the result.
- Document subprocessors, international transfers, SCCs and transfer-risk measures where required.

## Controller and processor obligations

| Role | Core responsibility |
|---|---|
| Marketplace controller | Own privacy notice, lawful basis, DPIA, security, rights, model governance, vendor/transfer controls and breach decisions for its purposes. |
| Employer controller | Own hiring purpose, candidate notice where applicable, lawful use, accommodation, human decision, record retention and discrimination compliance. |
| Marketplace processor | Process employer-only workflow data on documented instructions under DPA; support rights, deletion, security and breach response. |
| Skills Validation boundary | Maintain source evidence accuracy, dates and correction; document the transfer purpose and propagate changes to Marketplace. |

## Article 22 design rule

Do not ship automatic rejection, automatic ordering that determines who is seen, or a score treated as decisive. Where an employer reviews recommendations, the interface should expose source evidence, dates, score limits, and alternatives. The reviewer should record independent reasoning and be able to include an unranked candidate.

> Correction must follow lineage: Source-evidence disputes go to Skills Validation, while Marketplace corrects presentation, caches, matching, and employer copies. Both layers need a linked case record and completion evidence.

04 / EMPLOYMENT AND AI LAW

# Hiring technology creates obligations beyond ordinary SaaS privacy

| Regime | Product implication |
|---|---|
| EU AI Act | If the system qualifies as AI, recruitment and candidate-selection use is Annex III high risk; profiling is especially hard to exempt. Current stand-alone high-risk application date is December 2, 2027 after the 2026 AI Omnibus. |
| GDPR Article 22 | Solely automated decisions with legal or similarly significant effects are restricted. Transparency, fairness, lawful basis, DPIA, accuracy and objection apply even when Article 22 is not triggered. |
| US civil rights | Title VII, ADA, ADEA and selection rules apply when software makes or informs hiring. Validate job relevance, test adverse impact, provide accommodations and consider less discriminatory alternatives. |
| NYC Local Law 144 | Covered AEDT use can require an independent annual bias audit, public results summary, candidate/employee notice and information about criteria. |
| Colorado and other overlays | Colorado repealed and reenacted its framework in 2026 and is rulemaking. Maintain a current state/local matrix before enabling employers. |

## Matching and fairness controls

- Define intended job, population, outcome, inputs, score meaning, thresholds, ordering logic and prohibited use.
- Prove that each signal is job related and supported by current validated evidence.
- Exclude protected traits and obvious proxies from ranking. Segregate any lawfully collected fairness-test data.
- Measure selection-rate and error disparities across relevant groups and intersections before launch and after material changes.
- Provide accessible alternatives and reasonable accommodations; do not infer disability, health, emotion, personality or protected status.
- Give candidates plain-language notice, correction, objection, appeal and meaningful human review.

04 / EMPLOYMENT AND AI LAW CONTINUED

# Required matching-system file before launch

| File section | Required evidence |
|---|---|
| Intended purpose | Jobs, candidate population, employer user, workflow stage, decision supported, prohibited use |
| Data lineage | Source fields, verification state, date/recency, missing-data handling, excluded data and proxies |
| Method | Rule/model code, version, training/reference data, objective, feature transformations, thresholds and ranking |
| Validation | Job-related criterion, population, metrics, calibration, robustness, comparison with alternatives and limitations |
| Fairness | Group and intersection tests, sample sizes, selection/error metrics, thresholds, investigation and remediation |
| Human oversight | What the recruiter sees, training, override, independent reasoning, explanation and candidate appeal |
| Change control | Owner, peer review, test gates, version approval, customer notice, rollback and post-deployment monitoring |
| Incidents | Quality/fairness thresholds, alert owner, disable switch, impacted-person analysis, correction and communication |

Four-fifths is not a safe harbor: Selection-rate comparison is a useful screen, but smaller statistically and practically meaningful differences can still matter. Model validity, job relevance, disability access, error distribution, intersectionality and less discriminatory alternatives must also be assessed.

## Candidate explanation standard

Provide the principal data and criteria in concise, intelligible language: which verified skills and evidence were used, their dates, what employer criteria mattered, what a score does and does not mean, and how to challenge incorrect data. A complex formula alone is not an adequate explanation.

05 / LAUNCH GATES

# Twelve gates for an invitation-only real-candidate beta

| Gate | Decision | Pass evidence |
|---|---|---|
| 01 | Scope and geography | Named employers, adult candidates, allowed roles, human-decision boundary, prohibited data/uses, approved jurisdictions |
| 02 | System and data map | Frontend, backend, Skills Validation, portfolio, identity, matching, storage, vendors, regions and legal roles |
| 03 | Candidate source and notice | Source, permission/visibility, lawful basis, notice, content rights, expiry, withdrawal and correction |
| 04 | DPIA and legal assessment | GDPR DPIA, Article 22, EU AI Act classification, US jurisdiction matrix and residual-risk approval |
| 05 | Matching validation | Implementation, lineage, job relevance, metrics, limitations, drift, versioning and employer instructions |
| 06 | Fairness and accessibility | Adverse-impact/error tests, alternatives, accommodation, ongoing monitoring and required independent audit |

Default geography containment: Exclude NYC and Colorado employers from the first real-data pilot unless counsel confirms coverage and every required audit, impact assessment, public statement, notice, correction and appeal step is complete.

## Release interpretation

Every gate is a product requirement, not a document-only exercise. A signed policy without functioning application controls, a bias report without a versioned model, or a DPA without an accurate system map does not close a gate.

05 / LAUNCH GATES CONTINUED

# Security, operations, and signed accountability

| Gate | Decision | Pass evidence |
|---|---|---|
| 07 | Human review and appeal | Evidence view, independent reasoning, no auto-reject, explanation, correction, pause and trained human appeal |
| 08 | Sessions and invites | No JS-readable bearer token, MFA, rotation/revocation, CSRF/XSS controls, single-use invite and URL/log redaction |
| 09 | Tenant decision records | Server-side employer context/shortlist/dismissal, user and tenant scope, audit, retention and negative tests |
| 10 | Privacy and legal package | Approved terms, DPA, notices, AUP, subprocessors/transfers, retention, rights, deletion and breach workflows |
| 11 | Production security and SDLC | Headers, encryption, secrets, scanning, least privilege, logs, alerts, rate limits, restore, incidents and CI tests |
| 12 | Release and monitoring | Legal/privacy/security/product/model sign-off, limitations, thresholds, rollback, disable plan and residual risks |

> Manual is acceptable only when it is real: Rights, notice, appeal, bias review, deletion and incident coordination may begin as documented human workflows for a small beta if they are named, rehearsed, logged and capable of meeting required timelines.

## One-page go record

The release owner links each gate to current evidence, names customer and candidate restrictions, records residual risk and approving owner, and confirms monitoring and rollback. Any open blocker keeps the decision at no-go.

06 / LEGAL PACKAGE

# Documents required before an employer sees real candidates

| Document | Minimum content |
|---|---|
| Beta terms/order form | Named users, beta limits, permitted hiring purpose, support, suspension, IP, confidentiality, liability, exit and shared duties |
| DPA | Article 28 schedule, TOMs, subprocessors, transfers, rights/breach help, return/delete and audit information |
| Recruiter notice | Account, security, product, flags, support and telemetry purposes, bases, recipients, retention, transfers and rights |
| Candidate privacy/AI notice | Source, visibility, employers, matching purpose/criteria, limits, retention, correction, objection, appeal and human review |
| Employer AUP and anti-discrimination schedule | No sole reliance, unlawful targeting/screening, scraping, re-identification, resale or non-hiring use; accommodation and notice |
| Matching factsheet/instructions | Purpose, inputs, outputs, validation population, metrics, limits, prohibited uses, oversight, version and monitoring |
| Candidate/portfolio terms | Visibility, display license, evidence accuracy, employer access, withdrawal, correction, IP/takedown and deletion effect |
| Subprocessor/transfer list | Vendor, purpose, data, region, retention, transfer mechanism, safeguards and change notice |
| Security overview/TOMs | Actual system boundary, access, tenancy, encryption, SDLC, vendors, monitoring, backups, incidents and customer duties |
| Jurisdiction disclosures | NYC audit summary/notices and current Colorado or other state/local assessments, notices and appeals |

Excluded features: Do not add background or consumer reports, criminal/credit data, health/disability questions, biometrics, video/emotion analysis, protected-trait inference, children or candidate-contact automation without a separate legal and product review.

07 / SECURITY

# Minimum technical controls for beta

## Identity and sessions

- HttpOnly, Secure, SameSite sessions or equivalent reviewed design
- Short access lifetime, rotation, revocation and session view
- MFA for staff and employer admins
- Single-use expiring invites, URL cleanup and generic errors
- Clear tenant caches on logout and account switch

## Authorization and tenancy

- Server enforcement on every candidate and employer action
- User, employer, purpose and role scopes
- Negative tests across API, cache, links, logs, exports and search
- Immutable events for view, shortlist, dismiss, export and admin
- Periodic access review and immediate offboarding

## Application and data

- CSP, HSTS, frame, MIME, referrer and permissions headers
- Encryption in transit/at rest and managed secrets
- Validation, encoding, rate limits, abuse and CSRF controls
- No candidate content, secrets or tokens in telemetry
- Production devtools removed; SBOM and scanning

## Operations and resilience

- Restricted logs, alert owners and defined retention
- Patch/vulnerability deadlines by severity
- Backups, recovery objectives and restore test
- Security and harmful-model incident playbooks
- Disable/rollback path for faulty or unfair model

Highest priority: Replace JavaScript-readable bearer-token storage before an external customer beta. An XSS issue should not automatically expose a reusable account token.

07-08 / SECURITY AND PRIVACY OPERATIONS

# Repository-specific actions and candidate challenge workflow

## Engineering actions

- Replace _tmp.auth._token.local and expiration storage with the approved hardened session architecture.
- Move saved-candidates, dismissed-candidates and authoritative company-context-state to user- and tenant-scoped APIs.
- Exchange invite tokens immediately, remove them from browser URL/history, apply no-referrer on invite flows, and redact all telemetry.
- Verify LaunchDarkly contexts are private or pseudonymous and Cloudflare logs exclude candidate data, scores, tokens and request bodies.
- Remove Vue devtools from production; prove edge security headers in staging and production.
- Enable unit, lint, dependency/secret and a small auth/tenant/candidate smoke suite on pull requests.

## Candidate challenge workflow

| Step | Action |
|---|---|
| 1 | Log identity, employer/search, disputed fact or score, date and whether an employment decision is pending |
| 2 | Verify proportionately and pause affected recommendation or show the dispute where appropriate |
| 3 | Classify as source evidence, presentation, matching logic, employer action or combined issue |
| 4 | Correct source in Skills Validation and Marketplace display, caches, ranking and employer copies |
| 5 | Explain principal data and criteria in intelligible language, including dates and limitations |
| 6 | Provide trained human review with authority to change the result; record and notify |

Incident scope includes harmful outputs: Treat a confidentiality breach and a faulty or discriminatory ranking with the same ownership, containment, evidence preservation, affected-person analysis, rollback and corrective-action discipline.

09 / LIGHT SOC 2

# Build evidence habits now; do not buy an audit for beta

| Family | Light beta practice | Evidence |
|---|---|---|
| CC1-CC2 | Name security, privacy, product, model and employment-risk owners; approve concise policies and training | Approvals, responsibilities, training and customer instructions |
| CC3 | One risk register for privacy, discrimination, model, security, vendors and jurisdiction | Risks, DPIA/AI assessments, owners, dates and acceptance |
| CC6 | MFA, least privilege, tenancy, session hardening, approval, offboarding and review | Settings, approvals, negative tests, review exports and session evidence |
| CC7 | Monitor alerts, vulnerabilities, model quality and fairness; run both incident exercises | Scans, dashboards, incidents, tabletop results and remediation |
| CC8 | Peer review, automated tests, controlled deploy/model approval, rollback and emergency change | PR, CI, deploy, registry, approval and rollback records |
| CC9 | Tier identity, Cloudflare, LaunchDarkly, Skills Validation, portfolio, support and other vendors | Register, DPAs, transfers, reviews and exit decisions |
| PI design | Define score semantics, lineage, validation, completeness, correction and limits | Factsheet, tests, reconciliation, correction and appeal logs |

Later path: Type 1 when system boundary and control design are stable. Type 2 only after controls have operated consistently for a meaningful period. Security is required; add availability, confidentiality, processing integrity or privacy only when customer commitments justify them.

## Point-in-time dependency result

npm audit --package-lock-only --omit=dev reported zero known production vulnerabilities on August 13, 2026. This is not a clean bill of health: it does not prove private-package safety, application security, correct configuration, absence of secrets, or future advisory status. Retain ongoing scan evidence and remediation deadlines.

10 / DELIVERY PLAN

# 30/60/90-day path to a controlled beta

| Window | Outcome | Work |
|---|---|---|
| Days 0-30 | Contain and define | Keep candidates synthetic. Name owners. Map system/data/model. Decide source and discoverability. Classify roles/jurisdictions. Start DPIA, Article 22, AI Act and US assessments. Inventory vendors/telemetry. Draft prohibitions and risk register. |
| Days 31-60 | Implement and prove | Harden sessions/invites. Move employer context and decisions server-side. Prove tenancy. Implement notices, rights, retention and audit. Validate matching and adverse impact. Approve terms/DPA/notices/factsheet. Enable CI, restore and incident tests. |
| Days 61-90 | Named-employer pilot | Close jurisdiction requirements and independent audit where required. Sign documents. Activate approved people only. Rehearse correction, appeal and deletion. Monitor quality, disparity, access, incidents and complaints. Review every model change. |

## Release ownership

- Product owner: beta scope, prohibited uses, workflow and customer containment.
- Privacy/legal owner: roles, lawful bases, DPIA, notices, contracts, jurisdiction and rights.
- Security owner: sessions, tenant isolation, SDLC, monitoring, incident and evidence.
- Model owner: source lineage, validation, fairness, documentation, drift, changes and rollback.
- Support/operations owner: candidate correction, appeal, employer support, deletion and complaints.

> Go rule: Any open blocker keeps the real-data decision at no-go. Medium residual risk requires a named accepting owner, deadline and beta-specific containment.

11 / EVIDENCE CHECKLIST

# What the launch reviewer must be able to open

| Evidence | Owner | Pass condition |
|---|---|---|
| Scope/geography/prohibited use | Product / Legal | Named employers and people; approved jurisdictions and decision boundary |
| System/data map, RoPA, roles, LIA, DPIA | Privacy / Engineering | Dated and matches production fields, vendors, regions and flows |
| Candidate source, visibility, notice and rights | Product / Privacy | Every profile has source, basis, notice, expiry and correction |
| Model factsheet, code/version and validation | Model / Product | Job relevance, metrics, failure modes, lineage and changes approved |
| Fairness, access and jurisdiction results | Legal / Model | No unresolved material disparity; required audit/notices complete |
| Human review, correction and appeal rehearsal | Support / Product | Reviewer can explain, pause and change a result within target |
| Terms, DPA, notices, AUP, factsheet, TOMs | Legal / Privacy | Approved versions and signed employer acceptance |
| Session, invite, MFA, tenant and audit logs | Engineering / Security | Hardened settings and passing cross-tenant/account-switch tests |
| Telemetry, headers and vulnerability evidence | Security / Engineering | No tokens/candidate PII; no critical/high findings; headers proven |
| Restore and incident/model tabletop | Platform / Security | Successful tests with gaps owned and dated |
| Release and monitoring record | All owners | Gates closed; limits, thresholds, rollback and residual risk signed |

Reviewer standard: A hyperlink to current evidence is stronger than a statement that a control exists. Screenshots, configurations, tests, signed approvals, logs and completed rehearsals should all identify date, environment, owner and version.

12 / SOURCES

# Primary legal and assurance references


- EU General Data Protection Regulation, official text: https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng
- EDPB Guidelines on automated individual decision-making and profiling: https://www.edpb.europa.eu/documents/guideline/automated-decision-making-and-profiling_en
- EDPB endorsed guidance, including DPIA and transparency: https://www.edpb.europa.eu/endorsed-wp29-guidelines_en
- EU Artificial Intelligence Act, Regulation (EU) 2024/1689: https://eur-lex.europa.eu/eli/reg/2024/1689/oj?locale=en
- European Commission, 2026 AI Omnibus timeline: https://digital-strategy.ec.europa.eu/en/news/ai-omnibus-enters-force
- EEOC, Employment Tests and Selection Procedures: https://www.eeoc.gov/laws/guidance/employment-tests-and-selection-procedures
- EEOC Strategic Enforcement Plan 2024-2028: https://www.eeoc.gov/sites/default/files/2023-09/SEP%20FY%2020242028%20FINAL%20APPROVED.pdf
- EEOC and DOJ, disability discrimination and algorithmic hiring: https://www.eeoc.gov/newsroom/us-eeoc-and-us-department-justice-warn-against-disability-discrimination
- NYC DCWP, Automated Employment Decision Tools and Local Law 144: https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page
- Colorado Attorney General, Anti-Discrimination in AI Law rulemaking: https://coag.gov/ai/
- AICPA Trust Services Criteria: https://www.aicpa-cima.com/resources/download/2017-trust-services-criteria-with-revised-points-of-focus-2022


## Repository evidence

src/main.ts; router; auth service and signup/login; company context and saved-candidate composables; candidate services, types, cards, details and dashboard; onboarding and hiring context; architecture.md; Vite and HTML configuration; Wrangler, Terraform and Cloudflare workflows; E2E workflow and Playwright strategy; dependency manifests; and the production dependency audit.

> Important: This report is product, privacy and engineering guidance. It is not legal advice, a bias audit, a DPIA, a penetration test, a SOC 2 examination or an audit opinion. Qualified privacy and employment counsel should approve the legal model and jurisdictions. Qualified security and model-risk professionals should validate the deployed system before launch.