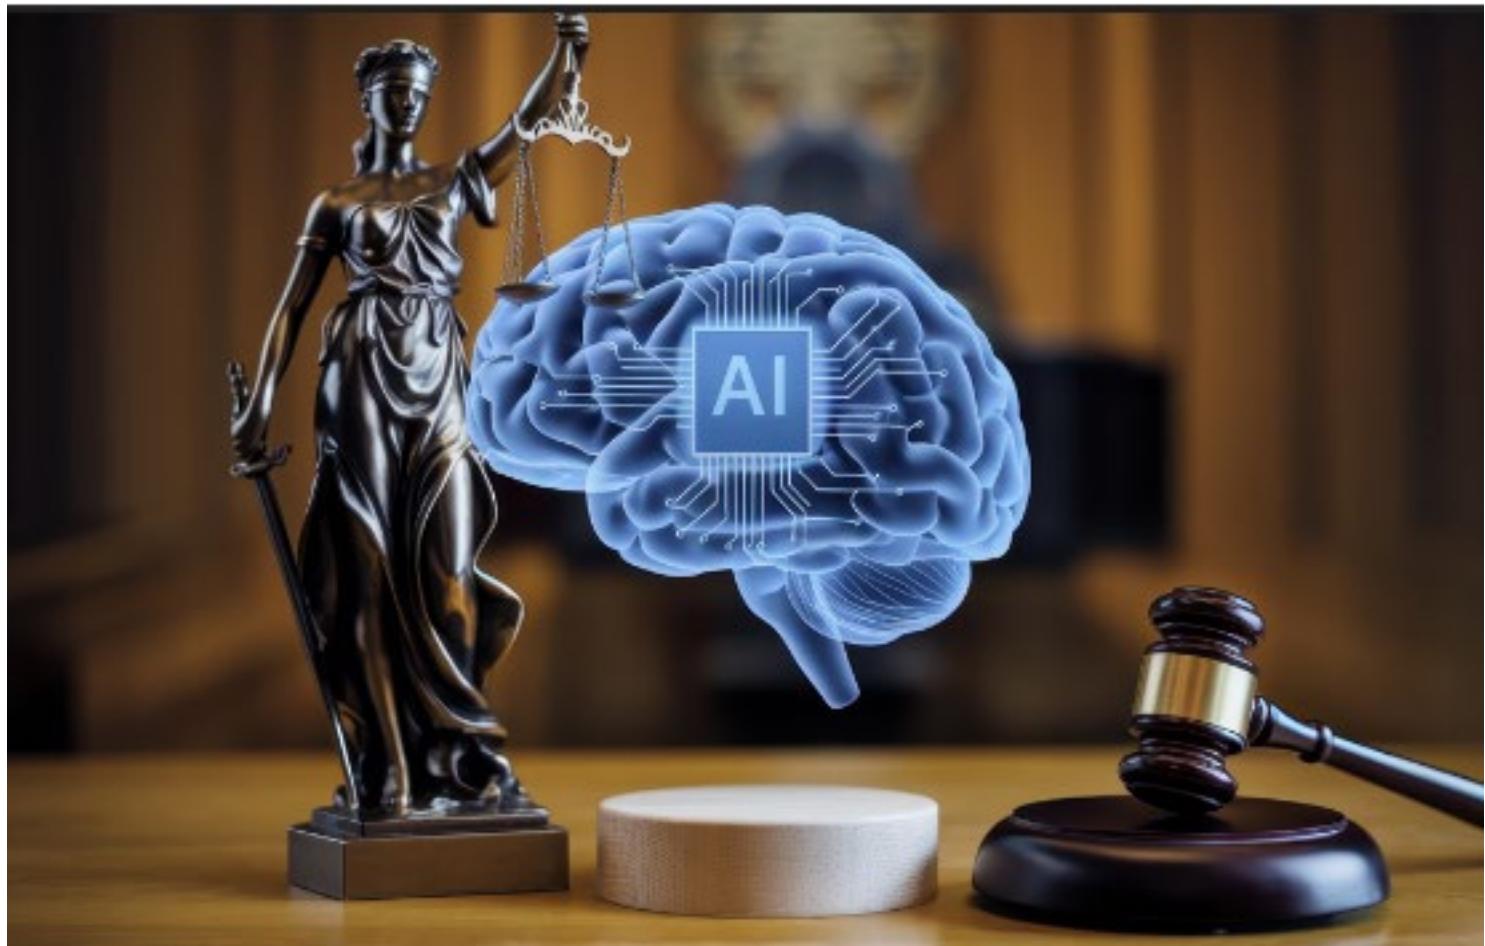


AI Governance Guide

For HR, OD, L&D, Talent Management Leaders



Practical templates to use AI safely in hiring, performance, mobility & L&D

How to use this guide

- **What this is:** A practical set of templates and checklists to help you deploy AI **safely** across talent processes.
- **What this isn't:** Legal advice or a substitute for your internal policies. Adapt to your context.

What's inside (quick list):

1. One-page AI Use Notices (candidate & employee)
2. Criteria-First Policy Worksheet (before any ranking/scoring)
3. Bias Audit Memo (lightweight, repeatable)
4. Risk Tiering & Human Oversight Matrix
5. AI Register (one sheet per tool)
6. Model/Update Log (stop silent drift)
7. Data Minimisation Checklist (by process)
8. Vendor Due-Diligence Questionnaire (+ acceptable evidence)
9. Performance Evidence Guardrails (+ evidence card schema)
10. L&D Recommendation Guardrails (+ approval form)
11. Decision Logs (Promotion & Mobility)
12. Human Review Request Form (candidate/employee)
13. 90-Day Implementation Plan
14. Metrics & Operating Rhythm
15. Glossary & References

Part 1 — Quick-Start Governance Blueprint

Principle: *AI drafts and surfaces options; people decide.*

Goal: Decisions that affect jobs (hire/promote/pay/mobility) are consistent, explainable, and defensible.

The 7 moves:

1. **Publish an AI Use Notice** (what's automated, what data, how to request human review).
2. **Criteria first.** Document job-related, business-necessary criteria **before** any AI ranking/scoring.
3. **Risk tiers + human oversight.** Low (drafting/summaries), Medium (recommendations), High (employment decisions) → require human approval with reasons for High.
4. **Bias checks you can explain.** Record selection rates/impact by group for high-stakes tools. Keep a one-page memo and act on findings.
5. **Data minimisation.** Use only inputs tied to the standard; avoid proxies (e.g., school prestige, postcode).
6. **Model stability.** Fix production randomness; version-lock updates where possible; keep a model/update log.
7. **AI Register.** One page per tool: purpose, data, owner, oversight step, last bias check, last update.

Part 2 — Templates & Checklists

2.1 AI Use Notice (Candidate-Facing)

Purpose: Transparency + contestability in recruitment.

Copy block (paste and fill):

Title: How we use AI in our hiring process

We use AI-assisted tools to [list steps: e.g., parse CVs, compare profiles to job criteria, schedule interviews]. These tools compare the information you provide to **documented, job-related criteria**.

- A **person reviews** important decisions such as shortlists and offers.
- You can **request a human review or raise a question** at [contact/email or portal].
- We do **not** use emotion detection or biometric categorisation.
- Data is handled according to our **Privacy Notice**: [link].

Footer: Last reviewed: [date] | Owner: [team]

2.2 AI Use Notice (Employee-Facing)

Copy block (paste and fill):

Title: How we use AI in talent and performance processes

We use AI-assisted tools to [list steps: e.g., summarise goals/feedback, surface project matches, recommend learning].

- Managers remain accountable for decisions (promotion, pay, mobility).
 - You can **view and correct** your profile data in [system] and **request a human review** of an AI-assisted outcome via [contact].
 - We do **not** use emotion detection or biometric categorisation.
 - See our **AI & Data Policy**: [link].
-

2.3 Criteria-First Policy Worksheet (1 page)

When to use: Before turning on any AI ranking/scoring for a role/process.

Fields (fill-in):

Process (eg: screening, readiness)	
Role & Level	
Business reason for AI use - e.g speed	
Job related criteria – must haves	
Job related criteria – nice to haves	
Evidence sources – e.g. assessments	
Excluded inputs (for proxy risk) e.g school	
Human approval (who & when)	
Effective date/next review	
Approvals (role, name & date)	

2.4 Bias Audit Memo (lightweight)

When to use: Annually and on major vendor updates for high-stakes tools.

Template:

Tool/Model name & version:	
Process: [e.g., screening]	
Population: [reqs, date range]	
Protected groups analysed: [list per your jurisdiction]	
Metric e.g. <i>Selection rate</i> = selected / applicants (per group) <i>Impact ratio</i> = group selection rate ÷ highest group rate	
Findings: [summarise by group; flag <0.80 impact ratio]	
Actions: [criteria review, feature removal, human overrides, vendor change]	
Owner	
Date	
Review Date	

Worked example (illustrative):

- Highest group selection rate = 24% (Group A)
- Group B = 18% → Impact ratio = 0.75 → **Action:** review criteria, add human review on borderline cases; re-run check in 30 days.
-
- **2.5 Risk Tiering & Human Oversight Matrix**

Process step	Risk Tier	Typical AI role	Human Oversight needed	Logs kept
Draft JDs & Competencies	Low	Generates draft text	HR check & update as needed	Version
Screening shortlist	High	Ranking versus criteria	Manager & HR approval with reasons	Decision log bias check records
Interviewing	Medium	Scheduling – question banks	Talent acquisition team oversight	Invite records
Performance conversation & feedback summaries	Medium	Themes and summary	Manager & employee review and approval	Source records maintained
Role matches (internal)	Medium/High	Suggests suitable roles for employee or suitable employees for role	Manager/HR approval and onboarding plan	Reasoning for match
Learning recommendations	Medium	Suggestions based on competency gaps	Manager approval	Pre and post competency assessment
Promotion readiness	High	Identify evidence	Manager/HR review & approval with reasons	Decision log

Adapt tiers for your organization

2.6 AI Register (one sheet per tool)

Field	Value
Tool name	
Owner (team & person)	
Process & purpose	
Decision impact (Low/Med/High)	
Data inputs (included/excluded)	
Data sources & storage location	
Explainability method shown to managers	
Human oversight step (who/when)	
Last bias check (date + summary)	
Last model update (date + version)	
Change log / release notes link	
Contact for review/challenge	

2.7 Model/Update Log

Tool/Model	Version	Change summary	Stability/QA check	Bias spot-check	Rollback link	Approver

Keep this simple but current.

2.8 Data Minimisation Checklist (by process)

Instructions: Tick only what is job-related. Strike or justify anything else.

Sourcing/Screening

- Verified experience in relevant tools/tech
- Observable competencies/indicators
- Work samples / portfolios
- Exclude:** school names/rank, postal codes, age, unrelated social links

Performance

- Goals & outputs within review period
- Customer outcomes (*use rate data see below*)
- Peer/manager observations (documented who, date)
- Exclude:** sentiment/emotion inference; unrelated personal data

Mobility/Readiness

- Verified must-have competencies
- Safety/clearance flags (if required)
- Availability & manager agreement
- Exclude:** personal circumstances unless employee opts-in

L&D Recommendations

- Documented gaps vs role standard
- Preferred learning modality/language
- Time/budget constraints
- Exclude:** resource popularity clicks as sole basis

2.9 Vendor Due-Diligence Questionnaire (+ evidence)

Ask vendors these (and capture links to their documentation):

Bias & Fairness

- Provide the **latest bias test** (groups, metrics, sample size) and remediation steps.
Acceptable evidence: PDF/portal report showing selection rates & impact ratios; dates; sample sizes.

Oversight & Explainability (Runtime environment and + UI-level explainability)

- Human oversight: Show exactly where a human must approve/override; show how overrides capture reason codes and are exportable.
- Logging: Provide a sample audit log with user, time, inputs, model/version, suggestion, human decision, reason, attachments.
- Per-decision explainability: Screens/PDFs that reveal which factors/evidence drove each score/match.
- Stability/versioning: Changelog, version ID surfaced in UI/logs, rollback option.

Updates & Stability

- Show how Changelog, version ID shown in UI/logs, and existence of rollback option.

Data

- Training/runtime **data sources**; locations; **retention**.
- Runtime has whitelist/blacklist and data dictionary.
- Confirm **no emotion inference or biometric categorisation**.

Validity

- Evidence linking outputs to **job outcomes** (criterion validity), not just offline accuracy.
- White paper or study showing correlation/effect sizes vs performance.

Security & Privacy

- Pen test/ISO 27001/SOC 2 certifications; DPIA support.
- Data access controls, audit trails, tenant isolation.

Contractual

- Service boundaries (no secondary use of our data); support SLAs for audits.

2.10 Performance Evidence Guardrails

Scope: Evidence surfaced from approved systems (HRIS, CRM, ticketing, QA tools) to help managers write reviews. **No automated scoring.**

Whitelist

- Systems: [list]
- Allowed fields: [list]
- Excluded fields: [list; e.g., free-form sentiment without source, biometric data]

Normalisation

- Compare like-for-like (per-ticket/per-hour/per-account).
- Exclude system errors/outliers.
- Show date ranges (e.g., last 6–12 months).

Traceability

Every item must include **system, record ID, date, link**.

Evidence schema (ensure your system can capture this data):

- Person: [name/id]
- Source system: [CRM/Ticketing/HRIS]
- Record: [ID/link] Date: [yyyy-mm-dd]
- What happened (1–2 lines): [fact]
- Why it matters (indicator linked): [indicator id/name]
- Manager review: [kept/removed + reason]

Employee comment path

Employees can comment/correct contextual errors before final sign-off.

2.11 Learning Resource Recommendation Guardrails

Required metadata (per item): Competency tag(s) • Level • Learning objective • Duration • Modality • Language • Cost • Prerequisites

Policy:

- Recommend **only** against **verified gaps**.
- Offer **2–3 options** (not a single “best pick”).
- Manager approves or swaps.
- Measure **gap closure** (pre/post check).

Learning Path Approval Form (1 page):

- Person/Role/Level: []
- Verified gap(s): []
- Recommended items (2–3) with metadata: []
- Time/cost fit: []
- Manager approval & start date: []
- Post-completion check (date & measure): []

2.12 Decision Logs

Promotion Decision Log

Field	Value
Person	
From level → To level	
Criteria (next-level indicators met)	
Evidence attached (links)	
Panel/manager decision & rationale	
Bias check reference (if applicable)	
Approvals (names/dates)	
Effective date	

Mobility Decision Log (Lateral/Temporary)

Field	Value
Person	
From role/team → To role/team	
Must-have competencies verified	
Onboarding/ramp plan	
Duration (if temporary)	
Manager & employee agreement	
Outcome review at 30/90 days	

2.13 Human Review Request Form

Who can use: Candidates or employees impacted by an AI-assisted step.

Field	Value
Name & Contact	
Process & date	
Decision/outcome being reviewed	
Reason for request	
Additional information	
Received by (team)	
Reviewer (human) & date	
Outcome & rationale	
Follow-up actions	

Part 3 — 90-Day Implementation Plan

Days 1–15 — Foundations

- Publish **AI Use Notices** (candidate & employee).
- Run **Criteria-First** for 3–5 pivotal roles.
- Stand up **AI Register + Model/Update Log**.
- Configure **risk tiers** & approval workflows in your tools.

Days 16–45 — Pilot

- Pick one team/process (e.g., Screening).
- Complete a **Bias Audit Memo** after first month.
- Turn on **Performance Evidence Guardrails**; capture evidence
- Launch **L&D Guardrails** on one role's gap set.

Days 46–90 — Scale & Measure

- Extend to 3–4 teams.
- Add **Vendor Due Diligence Questions** to procurement checklists.
- Start **operating frequency** (see below).
- Publish a short **internal update** with early metrics.

Responsible, Accountable, Consulted, and Informed (suggested):

- **HR/Talent:** owner of criteria, oversight, decisions
- **People Analytics:** bias checks, logs, metrics
- **IT/Data:** data flows, access control, stability tests
- **Legal/Privacy:** notices, data minimisation, Data Protection Impact assessments
- **Exec sponsor:** support, adoption, reporting

Part 4 — Metrics (define once; review monthly/quarterly)

Adoption & control

- % high-stakes decisions with **human approval** logged
- % tools listed in **AI Register** with current entries
- **Update log** freshness (last 30 days)

Fairness & validity

- Selection rate **impact ratios** by group (screening)
- **Disagreement rate** (employee vs manager on evidence) trending down
- **Readiness accuracy** (decision vs subsequent performance)

Effectiveness

- **Time saved** per task (JD drafting, shortlist creation, review prep)
- **Gap closure** post-learning (pre/post)
- **Time-to-competence** for new roles

Part 5 — Operating Timetable

Frequency	Activity	Owner
Monthly	Model/Update Log review; spot bias check	People Analytics + HR
Quarterly	AI Register refresh; metrics pack to executives	HR/Talent
Biannually	Vendor DDQ updates; contract clauses review	Procurement + Legal
Annually	Full Bias Audit Memo for high-stakes tools	HR/People Analytics

Part 6 — Glossary (short)

- **AI-assisted tool:** Software that uses ML/GenAI to summarise, rank, or recommend in a talent process.
- **High-stakes decision:** Outcomes affecting hiring, promotion, pay, or termination.
- **Bias audit memo:** Brief record of selection rates/impact by group and actions taken.
- **Criteria-first:** Documenting job-related standards **before** running AI scoring.
- **Decision log:** A short, human-readable record of what was decided and why.
- **Emotion inference:** Attempt to detect emotions/personality from face/voice; avoid in work contexts.

References

- European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- European Commission. (2025). AI Act: Regulatory framework for artificial intelligence (overview & timeline). <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- Bird & Bird. (2025, Mar 12). AI & the workplace: Navigating prohibited AI practices in the EU (Art. 5—emotion recognition ban).
<https://www.twobirds.com/en/insights/2025/global/ai-and-the-workplace-navigating-prohibited-ai-practices-in-the-eu>
- Center for Democracy & Technology. (2025, Apr 14). EU AI Act brief—AI at work.
<https://cdt.org/insights/eu-ai-act-brief-pt-4-ai-at-work/>
- NYC Department of Consumer and Worker Protection. (n.d.). Automated Employment Decision Tools (AEDT). <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page>
- State of Colorado. (2024). Senate Bill 24-205 (enacted text).
https://leg.colorado.gov/sites/default/files/2024a_205_signed.pdf
- New Jersey Division on Civil Rights. (2025, Jan 8). Guidance on Algorithmic Discrimination and the New Jersey Law Against Discrimination.
https://www.nj.gov/oag/newsreleases25/2025-0108_DCR-Guidance-on-Algorithmic-Discrimination.pdf
- National Institute of Standards and Technology. (2023). AI Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). <https://doi.org/10.6028/NIST.AI.100-1>.
- National Institute of Standards and Technology. (2024). Generative AI Profile (NIST AI 600-1). <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- UK Information Commissioner's Office. (2024, Nov 6). ICO intervention into AI recruitment tools—recommendations for vendors & employers. <https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2024/11/ico-intervention-into-ai-recruitment-tools-leads-to-better-data-protection-for-job-seekers/>
- UK Information Commissioner's Office. (2024, Nov 6). Thinking of using AI to assist recruitment? Key data protection considerations. <https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2024/11/thinking-of-using-ai-to-assist-recruitment-our-key-data-protection-considerations/>
-

Disclaimer

This pack provides general information to help organisations apply practical governance to AI-assisted talent processes. It does not constitute legal advice. Always consult your legal and privacy teams to adapt these templates to your jurisdiction and policies.